

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Binary Stars Across the Milky Way: Probes of Star Formation and Evolution

Permalink

<https://escholarship.org/uc/item/1bn3w3tc>

Author

El-Badry, Kareem Julian

Publication Date

2021

Peer reviewed|Thesis/dissertation

Binary Stars Across the Milky Way: Probes of Star Formation and Evolution

By

Kareem Julian El-Badry

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Astrophysics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Eliot Quataert, Co-chair
Associate Professor Daniel Weisz, Co-chair
Professor Martin White

Spring 2021

Binary Stars Across the Milky Way: Probes of Star Formation and Evolution

© Copyright 2021
by
Kareem Julian El-Badry
All Rights Reserved

Abstract

Binary Stars Across the Milky Way: Probes of Star Formation and Evolution

by

Kareem Julian El-Badry

Doctor of Philosophy in Astrophysics

University of California, Berkeley

Professor Eliot Quataert, Co-chair
Associate Professor Daniel Weisz, Co-chair

Binary stars are foundational to modern astrophysics. They underpin precision measurements of stellar structure, age, and composition; they provide stringent tests of general relativity; they make possible the study of faint and rare objects such as black holes and neutron stars; they are the progenitors of gravitational wave events. The components of binaries often interact, dramatically changing their evolution and giving rise to a bewildering zoo of astrophysical phenomenology. Binaries are not rare: about half of all Sun-like stars are binaries, and the binary fraction is even higher for massive stars. To understand stars – especially massive stars – it is thus critical to understand binaries. In the last decade, industrial-scale stellar surveys have provided exquisite spectroscopic, photometric, and astrometric data for hundreds of millions of stars in the Milky Way. The goal of this thesis has been to leverage this deluge of data toward new insight into the formation, evolution, and cosmic importance of binaries. On the one hand, binary population demographics represent the end state of the star formation process and thus inform theoretical star formation models. On the other, binary population demographics are an important input for stellar population models: stars in interacting binaries follow qualitatively different evolutionary histories from single stars of similar mass and composition. An improved census of the binary population thus has broad utility across astrophysics. Among other results, this thesis expands the total number of binaries known by more than a factor of 10. The methods developed here are expected to continue to bear fruit in the coming decades as panchromatic surveys look ever deeper into the Galaxy.

To my parents

Contents

List of Figures	viii
List of Tables	xiii
Acknowledgments	xiv
1 Introduction	1
1.1 Historical background	1
1.2 Thesis summary	4
1.3 Other work	5
2 Signatures of unresolved binaries in stellar spectra: implications for spectral fitting	10
2.1 Abstract	10
2.2 Introduction	11
2.3 Methods	11
2.4 Fitting binary spectra with single-star models	13
2.5 Fitting a binary spectral model	18
2.6 Conclusions	21
3 Discovery and Characterization of 3000+ Main-Sequence Binaries from APOGEE Spectra	23
3.1 Abstract	23
3.2 Introduction	24
3.3 Methods	25
3.3.1 Single-star spectral model	26
3.3.2 Binary spectral model	27
3.3.3 Model fitting	29
3.3.4 Fitting Multi-Epoch Spectra	31
3.4 Results	35
3.4.1 Effect of a velocity offset	36
3.4.2 Results for multi-epoch spectra with velocity variability	37
3.4.3 Color-magnitude diagram	44

3.4.4	Deriving Orbital Parameters	49
3.4.5	Are binaries gravitationally bound?	54
3.5	Discussion and Conclusions	55
3.5.1	Comparison to previous work	55
3.5.2	Future prospects	57
3.5.3	Summary	59
3.6	Neural network spectral model	61
3.7	Model selection	62
3.7.1	Tests with semi-empirical synthetic binary spectra	63
3.8	False positives	68
3.9	Orbit fitting convergence	71
3.10	Data products	72
4	Imprints of white dwarf recoil in the separation distribution of Gaia wide binaries	76
4.1	abstract	76
4.2	Introduction	77
4.3	Methods	78
4.3.1	General quality cuts	79
4.3.2	Initial Search	82
4.3.3	Removing clusters, moving groups, and higher-order multiples	86
4.3.4	Summary of the binary catalog	86
4.3.5	Contamination Rate	90
4.4	Separation distributions	93
4.4.1	Inferring intrinsic separation distributions	93
4.4.2	Fitting the separation distribution	95
4.4.3	Fitting results	95
4.5	Theoretical interpretation	97
4.5.1	Semi-analytic binary evolution models	97
4.5.2	Disruption from external perturbations	100
4.5.3	White dwarf recoil	101
4.6	Summary and Discussion	103
4.6.1	Discussion	106
4.7	Gaia spatial resolution and dependence on flux ratio	109
4.8	Relation between projected separation and semimajor axis	111
4.9	Distance dependence	112
4.10	ADQL query	112
4.11	Catalog	116
5	The wide binary fraction of solar-type stars: emergence of metallicity dependence at $a < 200$ AU	118
5.1	Abstract	118

5.2	Introduction	118
5.3	Methods	119
5.3.1	Wide binary catalog and extension	119
5.3.2	Spectroscopic metallicities	120
5.4	Results	121
5.4.1	Metallicity Distribution	121
5.4.2	Inferring the dependence of binarity on $[\text{Fe}/\text{H}]$	121
5.4.3	Separation distributions	125
5.5	Discussion and Conclusions	125
6	The geometric challenge of testing gravity with wide binaries	129
6.1	Abstract	129
6.2	Introduction	130
6.3	Simulations	132
6.3.1	Projection effects	132
6.3.2	Correcting for projection effects	134
6.4	Discussion	136
7	Discovery of an equal-mass “twin” binary population reaching 1000+ AU separations	138
7.1	Abstract	138
7.2	Introduction	139
7.3	Data	141
7.3.1	Is the twin excess real?	146
7.4	Modeling	151
7.4.1	Parameterization	152
7.4.2	Selection Function	155
7.4.3	Model inputs	155
7.5	Results	158
7.6	Discussion	167
7.6.1	Comparison to previous work	167
7.6.2	Origin of twin binaries	171
7.7	Summary and conclusions	179
7.8	Functional form of $p(q)$	182
7.8.1	How sharp is the twin feature?	182
7.8.2	Choice of q_{break}	182
7.8.3	Smoothly-broken power law	185
7.9	Sensitivity to systematics	187
7.10	Twins in other catalogs	191
7.11	Selection function	192
7.11.1	Single-star term	193
7.11.2	Contrast sensitivity for close pairs	198

7.12	Model validation	199
7.12.1	Single stars	199
7.12.2	Chance alignments	201
7.13	RV variability	201
7.14	Full fitting constraints	205
8	Not so fast: LB-1 is unlikely to contain a 70 M_{\odot} black hole	207
8.1	Abstract	207
8.2	Introduction	207
8.3	Basic issue	208
8.4	Contamination from absorption in LB-1	211
8.4.1	Fitting RVs	211
8.5	Discussion	213
8.5.1	Nature of the unseen companion	213
8.5.2	Source of the H α emission	214
9	A stripped-companion origin for Be stars: clues from the putative black holes HR 6819 and LB-1	218
9.1	Abstract	218
9.2	Introduction	219
9.3	Methods	220
9.3.1	Data	220
9.3.2	Spectral disentangling	221
9.3.3	Temperature and gravity	224
9.3.4	Radial velocities	227
9.3.5	Luminosity	231
9.3.6	Emission lines	233
9.3.7	Chemical abundances	235
9.3.8	Evolutionary history	237
9.4	Summary and Discussion	245
9.4.1	Implications for the formation of Be stars	246
9.5	Stellar parameters anchored to evolutionary tracks and distance	248
9.5.1	System parameters independent of evolutionary tracks	250
9.6	Spectral disentangling and determination of K_{Be}	251
9.7	B star variability	256
9.8	Rotation vs. macroturbulent broadening	258
9.9	Helium enrichment	260
10	A million binaries from Gaia eDR3: sample selection and validation of Gaia parallax uncertainties	262
10.1	Abstract	262
10.2	Introduction	263

10.3	Sample selection	264
10.3.1	Cleaning clusters, background pairs, and triples	267
10.4	Chance alignments	270
10.4.1	Chance alignment rate for subsets of the catalog	272
10.4.2	Estimating chance alignment probabilities	274
10.5	Basic properties of the catalog	277
10.5.1	Twin binaries and contrast sensitivity	281
10.5.2	Comparison to other catalogs	283
10.5.3	Space density	285
10.5.4	Orbital velocities from proper motion differences	287
10.5.5	Cross-match with LAMOST	291
10.6	Calibrating <i>Gaia</i> DR3 parallax uncertainties	295
10.6.1	Comparison to Gaia DR2	299
10.6.2	Parallax zeropoint corrections and 5-parameter vs 6-parameter solutions	300
10.6.3	Color dependence	301
10.6.4	Fitting function to inflate σ_{ϖ}	301
10.6.5	Angular correlations in parallaxes	303
10.6.6	Comparison to other work	304
10.7	Summary and discussion	304
10.7.1	Scientific uses for the catalog	307
10.7.2	Constraining the parallax zeropoint with binaries	308
10.7.3	Astrometric acceleration	309
10.8	Catalog description	311
10.9	Chance alignment probabilities	312
10.9.1	Sources with spurious astrometry	313
11	LAMOST J0140355+392651: An evolved cataclysmic variable donor transitioning to become an extremely low mass white dwarf	317
11.1	Abstract	317
11.2	Introduction	318
11.3	Data	320
11.3.1	LAMOST discovery spectra	320
11.3.2	ZTF light curve	320
11.3.3	Other photometry	320
11.3.4	Light curve analysis	323
11.3.5	KAST Spectra	326
11.3.6	Distance	331
11.3.7	Spectral energy distribution	332
11.4	Parameter constraints	336
11.4.1	Light curve model	336
11.4.2	Combined model	337
11.5	Outburst frequency limits	339

11.6	Evolutionary history and comparison to known CVs and ELM WDs	342
11.6.1	MESA models	345
11.6.2	Future evolution	351
11.6.3	Varying model parameters	353
11.7	Summary and Discussion	354
11.7.1	Is J0140 detached or mass-transferring?	356
11.8	Summary of LAMOST data	357
11.9	Summary of Kast spectra	358
11.10	Phase-dependent spectra	358
11.11	CMD position	358
	Bibliography	364

List of Figures

2.1	Simulated unresolved companion in an APOGEE-like spectrum	14
2.2	Effects of unrecognized binarity on labels recovered from spectral fitting	15
2.3	Binarity in the HR diagram with APOGEE-like spectra	16
2.4	Probability contours obtained by fitting a binary model	19
2.5	Difference in total χ^2 between binary and single-star models	20
3.1	Spectrum of unresolved binary compared to single star, with fit.	33
3.2	Single-star and binary model fits to binary systems	34
3.3	A binary system in which the stars' velocities change between visits.	38
3.4	Comparison of spectroscopically- and dynamically-inferred mass ratios for "SB2" binaries	39
3.5	Visit spectra and best-fit models for an SB1 system.	41
3.6	Visit spectra of a triple (SB3) system	43
3.7	Velocities for hierarchical triples	45
3.8	Visit spectra of a triple system with dark component	46
3.9	Color-magnitude diagram of binary candidates	47
3.10	Orbit fit to an SB2 systems with 15 epochs	49
3.11	Distribution of periods, semi-major axes, and masses	53
3.12	Range of binary periods and mass ratios that can be detected with different methods	57
3.13	Fitting semi-empirical binary spectra with single-star and binary models	64
3.14	Results for fitting semi-empirical synthetic binary spectra	66
3.15	χ^2 differences between best-fit single and binary models.	67
3.16	Label recovery diagnostic for semi-empirical binary spectra	69
3.17	Spectra of likely false positives	70
3.18	Minimum number of epochs for orbital solution	73
4.1	Joint distribution of projected separation and distance for all wide binaries	80
4.2	On-the-sky physical velocity difference between members of candidate binaries	81
4.3	Spatial distribution of wide binary candidates	84
4.4	Color-magnitude diagram for our final sample of wide binaries.	85
4.5	Distributions of observable properties of WD/MS, WD/WD, and MS/MS binaries	88
4.6	Radial velocity validation of binary candidates	89

4.7	Distributions of projected physical separation for MS/MS, WD/MS, and WD/WD binaries	92
4.8	Broken power-law fits to the intrinsic separation distributions of MS/MS, WD/MS, and WD/WD binaries	96
4.9	Separation distribution of MS/MS binaries for systems with higher- and lower-mass primaries	98
4.10	semi-major axis distributions from binary population synthesis simulations	99
4.11	Comparison of data to binary population synthesis separation distribution	104
4.12	Two-point autocorrelation function and <i>Gaia</i> contrast sensitivity	110
4.13	Ratio of semi-major axis to projected separation for simulated binaries	111
4.14	Separation distributions of nearby vs. more distant binaries	113
5.1	MDFs of binaries (red) and all stars with spectroscopic metallicities (black)	122
5.2	Metallicity-dependence of the binary fraction.	124
5.3	Normalized separation distributions for binaries of different metallicities,	126
6.1	Distributions of distance, proper motion, and radial velocity	131
6.2	One-dimensional velocity difference versus projected separation for simulated binaries.	133
6.3	Predicted ΔV recovered from precise proper motions when projection effects are corrected for	135
7.1	Color-magnitude diagram of all stars in the ER18 catalog of resolved wide binaries	142
7.2	SDSS images of a selection of binaries from our catalog.	144
7.3	Apparent magnitude difference between the two components of binaries	145
7.4	Difference in <i>G</i> -band magnitude for binaries and chance alignments	147
7.5	Ratio of the number of twins to non-twins	148
7.6	Magnitude difference as a function of separation for binaries with different primary masses	150
7.7	Parameterized mass ratio distribution.	153
7.8	Magnitude difference vs. mass ratio for binaries of different primary mass	157
7.9	Mass ratio distribution constraints.	159
7.10	Histograms of magnitude difference between the two components of a binary	161
7.11	Marginalized constraints on F_{twin}	162
7.12	Marginalized constraints on $\gamma_{\text{large}q}$	164
7.13	Constraints on q_{twin} , $\gamma_{\text{small}q}$, and γ_s	166
7.14	Cumulative distribution function of short-period double-lined spectroscopic binaries	169
7.15	F_{twin} and q_{twin} for solar-type binaries	170
7.16	Predictions of the toy model for orbit widening	176
7.17	Magnitude difference for binaries with unresolved hierarchies	178
7.18	Constraints on $p(q)$ near $q = 1$	183
7.19	Mass ratio distribution constraints for different q_{break}	184
7.20	Example smoothly broken power laws	186

7.21	Probability contours for one representative bin of primary mass and physical separation	188
7.22	Excess of twins in the SLoWPoKES-II catalog	189
7.23	Twin fraction vs. distance and angular separation	190
7.24	Effective section function	194
7.25	Lognormal fits to distributions of flux SNR	196
7.26	Fraction of suspected genuine sources within 200 pc that survive quality cuts . .	197
7.27	Contrast sensitivity given our quality cuts.	198
7.28	Validation of the single-star selection function and underlying Galactic model . .	200
7.29	Chance alignment magnitude difference distribution	202
7.30	Typical radial velocity scatter vs. magnitude	204
8.1	Schematic illustration of the effects of an RV-variable absorption line on the apparent velocity of a stationary emission line	209
8.2	Keck/HIRES spectra of LB-1 obtained at two different epochs, with and without subtraction of emission	210
8.3	Voigt profile fits to 7 epochs of Keck/HIRES data for LB-1	212
8.4	RVs for LB-1 with and without subtraction of stellar absorption	215
9.1	Disentangled spectra of the B star (top) and Be star (bottom) in HR 6819 . . .	223
9.2	Balmer lines as a surface gravity diagnostic in HR 6819	226
9.3	Temperature of the B star from ionization equilibrium	228
9.4	Radial velocities for the B star and Be star	230
9.5	Color-magnitude diagram showing HR 6819	234
9.6	Line emission from the disk and circumstellar envelope of the Be star	236
9.7	Models from BPASS binary evolution library that at any point in their evolution resemble HR 6819	238
9.8	Evolutionary history of the primary for stripped star models similar to HR 6819 and LB-1	239
9.9	Time-evolution of MESA model for HR 6819	241
9.10	Atmospheric parameters for the two components of HR 6819 compared to MIST evolutionary tracks	249
9.11	Total χ^2 versus RV semi-amplitude of the Be star assumed during spectral disentangling	253
9.12	Constraints on K_{Be} from the $\text{H}\beta$ line, wavelength vs. Fourier space disentangling	254
9.13	RVs for the Be star measured by fitting the He I absorption line with synthetic models	255
9.14	Line profile variations for the B star with time	257
9.15	Comparison of rotational and macroturbulent broadening for the B star	259
9.16	Diagnostics of He enhancement in the B star in HR 6819	261
10.1	Sky distribution of binary candidates in Galactic coordinates	268

10.2	Separation distribution of binary candidates at three stages of the selection procedure	269
10.3	Two methods for estimating the contamination rate from chance alignments . . .	271
10.4	Contamination rate from chance alignments in various subsets of the catalog . .	273
10.5	Ratio of the local density of chance alignments to binary candidates	275
10.6	RVS velocities of the components of binary candidates wider than 50,000 au . .	276
10.7	Basic properties of the high-bound probability binary sample	278
10.8	Color–absolute magnitude diagrams of high-confidence WD+MS and WD+WD binaries	281
10.9	Magnitude difference between the two stars in chance alignments and bound wide binaries	282
10.10	Comparison of the catalog produced in this work to other catalogs	283
10.11	Midplane number density of binaries in different cumulative distance bins	286
10.12	Plane-of-the-sky velocity difference between the components of wide binary candidates selected with different criteria	288
10.13	Binary candidates colored by ruwe	290
10.14	Wide binaries from the LAMOST survey	292
10.15	Distributions of uncertainty-normalized parallax difference between the two components of binaries	294
10.16	Fractional parallax uncertainty underestimate as a function of magnitude and separation	296
10.17	Fractional parallax uncertainty underestimate as a function of magnitude	296
10.18	Fraction of pairs in which at least one component fails an astrometric quality cut	297
10.19	Comparison of DR2 and DR3 parallax uncertainty underestimates	299
10.20	Median reported and corrected parallax uncertainty as a function of G magnitude	300
10.21	Parallax uncertainty underestimate factors for 5 and 6 parameter solutions . . .	302
10.22	Parallax uncertainty underestimate factors for binaries of different colors	303
10.23	Fitting function to inflate parallax uncertainties	305
10.24	Parameter distribution of chance alignments from the shifted catalog, binary candidates, and high-confidence binary candidates	314
10.25	Ratio of the local density of chance-alignments from the shifted catalog to that of binary candidates	315
11.1	ZTF, PTF, and CSS light curves of J0140	321
11.2	Light curve of J0140, showing data from several time-domain surveys	322
11.3	Difference between observed ZTF magnitudes and the mean magnitude at the corresponding orbital phase	324
11.4	Radial velocity of the proto-WD, phased to a period of 3.81 hours	327
11.5	Dynamically implied mass of the companion as a function of the proto WD’s mass and orbital inclination	330
11.6	Rest-frame normalized spectrum of the proto-WD and best-fit Kurucz model spectrum	332

11.7 Spectral energy distribution of J0140.	334
11.8 Best-fit Phoebe light curves when the inclination is fixed to 60, 75, and 85 degrees	338
11.9 Parameter constraints from joint fitting of the light curve, SED, parallax, orbital solution, and donor spectrum	340
11.10 Probability that at least one outburst would have been detected	343
11.11 Effective temperatures and orbital periods of J0140 and known CV donors that have been proposed to have evolved donors	344
11.12 MESA binary models for CVs with evolved donors	347
11.13 HR diagram for J0140	350
11.14 Future evolution of J0140 models	351
11.15 Radial velocities measured from individual LAMOST sub-exposures for two sep- arate visits	359
11.16 Coadded spectra of J0140 at different phases	361
11.17 J0140 on the Gaia color-magnitude diagram (CMD).	362

List of Tables

3.1	Orbital solutions for double-line spectroscopic binaries.	52
3.2	Minimum $\Delta\chi^2$ and improvement fraction f_{imp}	65
3.3	List of targets identified as single stars.	74
3.4	Best-fit labels for targets identified as SB1s.	74
3.5	Best-fit labels for targets identified as SB2s	75
3.6	Best fit labels for SB2s with a hidden dark component	75
3.7	Best-fit labels for SB3s	75
4.1	Catalog description	117
5.1	Sources of the spectroscopic abundances	120
7.1	Summary of priors adopted in each primary mass bin	154
7.2	Marginalized 2 sigma constraints on fitting parameters for all bins of primary mass and separation	206
9.1	Parameter constraints for HR 6819	225
9.2	Summary of our MESA binary models.	240
9.3	System parameters anchored to both distance and evolutionary tracks. Uncertainties are 1σ (middle 68%).	251
10.1	Contents of the binary candidate catalog	279
10.2	Features used by Gaussian KDE	313
11.1	Basic observables of J0140. Uncertainties are 1σ (middle 68%). Additional constraints are in Table 11.3.	328
11.2	Spectral energy distribution of J0140	333
11.3	Constraints from the combined fit of all data	341
11.4	LAMOST DR5 parameters for J0140, estimated from spectra obtained in two separate visits	357
11.5	Summary of Kast spectra of J0140	360

Acknowledgments

I am indebted to many people who have offered me advice, mentorship, and encouragement throughout my studies. This thesis would not have been possible without them.

First I would like to thank my advisors: Eliot Quataert, Dan Weisz, and Hans-Walter Rix. There is some luck involved in finding one good PhD advisor; fortune smiled on me when it gave me three. I am grateful to my advisors both for cultivating and guiding my scientific interests and for providing career advice throughout my development as a professional astrophysicist. Beyond excellent scientific training, I am grateful to Eliot for being the most loyal supporter of Friday Afternoon Cheese Time and co-star of the Campbell Hall anti-LB-1 photo shoot; to Dan for finding an excuse to invite me to four separate “Near/Far” workshops in Napa, even after I had effectively switched fields; and to Hans-Walter for cooking me 100+ meals over the last several years and taking me on regular trips to the mountains.

I also want to thank Marla Geha, my undergraduate research advisor, for launching my astronomy career and teaching me how to approach research problems. Without her patient mentorship early on, I likely would have ended up studying philosophy or (God forbid) becoming some kind of consultant.

I thank the many other mentors whom I have had the pleasure of working with during graduate school, particularly Joss Bland-Hawthorn, Charlie Conroy, David Hogg, Phil Hopkins, Changoo-Kim, Eve Ostriker, Yuan-Sen Ting, and Andrew Wetzel. My many excellent mentors provided a diverse array of lenses through which to view the Universe and toolkits with which to tackle physics problems, and my approach to doing research is basically an amalgamation of what I have learned from them.

I am also grateful to many Berkeley graduate students with whom I have learned, joked, and commiserated in the last five years, especially Fatima Abdurrahman, Ellie Abrahams, Sandra Albers, Nick Choksi, Drummond Fielding, Philipp Kempfski, David Khatami, Hannah Klion, Kara Kundert, Casey Lam, Michael Medford, Ned Molter, Sarafina Nance, Sean Ressler, Nathan Sandford, Wren Suess, and Tom Zick.

I am also indebted to Dexter Stewart, Nina Ruyemaker, Amber Banayat, and Yasasha Ridell, for making almost everything at Berkeley work; to the staff at Lick Observatory, especially Elinor Gates, Paul Lynam, Jon Rees, Matt Brooks, Shawn Stone, and Dan Espinosa, for their help in acquiring the data that made this thesis possible; to Geoff Tabin and In-Hei Hahn, for welcoming me in their homes for an unreasonably long period of time during the

COVID-19 pandemic; to Anna McLeod, Nick Choksi, Ben Margalit, and Yong Zheng, for organizing HICCUP and getting me into running; to Mary Beth Robinson, for hosting me in Tucson and Heidelberg throughout the last few years and always finding me peanut butter in Germany; and to Taylor Swift, Britney Spears, and Kesha, for unwitting aural support.

I thank my parents, Magdy and Sonja El-Badry, for fostering in me excitement about science and space from an early age, for making countless sacrifices in their own lives for my education, and for cheering for me from afar throughout college and graduate school. Sorry about the “garden to the Ivy League” incident. Thanks to my sister, Sophia, to my aunts, uncles, cousins, and grandparents, and to my adoptive family in the Tabin and Demarchis households, for all their love and support over the years. Finally, I thank Sara Tabin, my partner and best friend, for five years of friendship and unconditional support. Thank you for the good advice to start working on binary stars instead of galaxy formation, for being patient when I was grumpy about the physics not working, and for believing in me unwaveringly. I can’t wait for the rest of our lives together.

I thank the National Science Foundation for supporting my research through a graduate research fellowship, the Hellman Fellows program for awarding me a graduate fellowship, and the Berkeley Fellowship program for funding my first two years of graduate school.

Chapter 1

Introduction

1.1 Historical background

Binary stars have played a paramount role in the history of astronomy. Algol, the brightest eclipsing binary star in the sky, was observed to periodically fade every ~ 3 days by ancient Egyptian scribes more than 3000 years ago. Algol’s 2.85 day period¹ is clearly visible as a peak in a periodogram of the Cairo Calendar of lucky and unlucky days (Jetsu et al. 2013). This document is likely the oldest surviving record related to any variable star.

It is not known whether the ancient Egyptian understood the reason for Algol’s periodic dimming. The earliest surviving suggestion that eclipses were to blame was made by Goodricke (1783). Around the same time, Herschel (1782) began constructing the first known catalog of spatially resolved wide binary stars.² It took more than 100 more years until the first close binary was discovered spectroscopically (Pickering 1890). The first high-precision light curve for an eclipsing binary – once again Algol – was calculated by Stebbins (1910).

Outside the Solar system, binary stars provide the best measurements of stellar masses and radii (e.g. Andersen 1991; Torres et al. 2010). These measurements are (relatively) independent of modeling assumptions and depend on well-understood physics: Newtonian

¹In fact, the orbital period of Algol today is 2.87 days. However, mass transfer from a lower- to a higher-mass star widens a binary’s orbital period. For Algol’s mass transfer rate, $\dot{M} \approx 2 \times 10^{-7} M_{\odot} \text{yr}^{-1}$, the binary’s orbital period is predicted to have been 2.85 days in 1200 BCE, when the Cairo Calendar was produced (Jetsu et al. 2013).

²Herschel was not looking for binaries, and in fact did not know that binaries existed in 1782: he was attempting to measure stellar parallaxes, and was looking for chance alignments. In Herschel’s time, there was no universal coordinate frame (“ICRS”) against which to measure stellar positions. Herschel knew it would be easier to measure small angular displacements of a foreground star if there were a nearby background star with negligible parallax. With this goal in mind, Herschel identified more than 800 close pairs, down to about 11th magnitude, and carried out astrometric measurements of their components over 20 years. This was a clever idea, but it did not work, because almost all of the pairs in Herschel’s catalog were physically associated binaries, not chance alignments. This meant that the two stars had the same parallax, making it impossible to measure the parallax of one component relative to the other. Herschel eventually realized this (Herschel 1803).

mechanics and geometry. [McLaughlin \(1924\)](#) measured the masses and radii of the components of Algol by simultaneously fitting the light curve and spectroscopic radial velocity measurements, discovering the first evidence of the Rossiter-McLaughlin effect in the process. Precision measurements of masses and radii in a few dozen binaries allowed [Hertzsprung \(1923\)](#) and [Kuiper \(1938\)](#) to construct an empirical mass-luminosity relation for main sequence stars across more than two decades in mass and nearly eight decades in luminosity.

A mathematical theory of stellar structure and evolution was developed during the 1920s and 30s (e.g. [Eddington 1924](#); [Jeans 1925](#); [Eddington 1926](#); [Rosseland 1926](#); [Russell 1929](#); [Milne 1930](#); [Bethe 1939](#)). It was soon realized that the masses and evolutionary states of the components of Algol and other similar close binaries appeared to be inconsistent with the theory (e.g. [Kuiper 1941](#)). The apparent tension with the theory, dubbed the “Algol paradox”, can be summarized as follows: (a) it is assumed that the component stars in most binaries are coeval, (b) stellar evolution predicts that lower-mass stars have shorter lifetimes than higher-mass stars, and (c) in Algol and other similar binaries, the low-mass star is more evolved than the high-mass star, contradicting (b).

The Algol paradox was investigated by many astronomers in the next few decades, with a variety of incorrect explanations proposed (see [Pustyl'nik 1998](#), for a summary). The (presumably) correct solution was first proposed by [Crawford \(1955\)](#): the more evolved star was initially the more massive component, but it has since transferred most of its mass to the companion. This evolutionary scenario is now thought to be quite common in close binaries, but it was not widely accepted as the resolution of the Algol paradox for several decades. The upshot of the Algol story, which will become relevant again later in this thesis, is that *interactions between the components of binary stars make possible a much wider diversity of evolutionary histories than is normally allowed under single-star evolution.*

The utility of binary star *population demographics* – that is, their distributions of period, eccentricity, primary mass, and mass ratio – was recognized early. With the publication of a catalog of ~ 3000 wide binary candidates by [Struve \(1837\)](#), binary star studies firmly entered the realm of statistical astronomy. The purity and completeness of early samples were, however, not well understood. All binary star surveys are subject to selection effects, which for spatially resolved binaries introduce biases against unequal mass ratios and close separations. An attempt to correct for these biases was made by [Kuiper \(1935\)](#), who recognized the potential of binaries as a window into the star formation process. Using data from the recent Lick observatory spectroscopic survey of [Aitken & Doolittle \(1932\)](#) and various other catalogs, Kuiper combined samples of single-lined binaries (“SB1”; spectroscopic binaries in which only one component is detectable in the spectrum) and double-line binaries (“SB2”) to constrain the binary mass ratio distribution as a function of primary mass.

The methods he used to account for incompleteness were applied in many subsequent studies. [Abt & Levy \(1976\)](#) carried out a multi-epoch radial velocity survey with the goal of constraining the frequency of binaries among Solar-type stars. They found that the period distribution of Solar-type binaries is roughly lognormal, with a peak between 10 and 100 years and a total binary fraction of about 50%. This result was solidified in later studies by [Duquennoy & Mayor \(1991\)](#) and [Raghavan et al. \(2010\)](#); other works showed that the binary

fraction and peak separation vary with primary mass (e.g. [Duchêne & Kraus 2013](#); [Moe & Di Stefano 2017a](#)).

Beginning in the 1960s, it was recognized that mass transfer within close binaries gives rise to a broad range of astrophysical phenomenology. Mass-transferring white dwarf + main sequence binaries were identified as the cause of cataclysmic variables ([Kraft 1962](#)), and an extensive theory of a cataclysmic variable evolution was developed over the next few decades ([Warner 1995](#)). In the 1970s, many of the bright X-ray sources discovered by balloon and satellite missions were found to be associated with normal-star optical sources. These objects were soon identified as “X-ray binaries” in which a black hole or neutron star accretes mass from a close stellar companion (e.g. [Webster & Murdin 1972](#); [Clark 1975](#); [Margon et al. 1979](#)). The first radio pulsar discovered in a binary ([Hulse & Taylor 1975](#)) made possible precision measurements of the masses of neutron stars, and precision tests of general relativity. Increasingly sophisticated binary evolution models were developed to explain the formation of observed systems (e.g. [Kopal 1956](#); [Spruit & Ritter 1983a](#); [Bhattacharya & van den Heuvel 1991](#); [Kalogera & Webbink 1996](#); [Hurley et al. 2002a](#); [Paxton et al. 2011](#)). Binaries were also identified as the progenitors of gamma ray bursts ([Narayan et al. 1992](#)), type Ia supernovae ([Whelan & Iben 1973](#); [Iben & Tutukov 1984](#); [Branch et al. 1995](#); [Hachisu et al. 1999](#)), and stripped-envelope (Ib/Ic) supernovae (e.g. [Yoon et al. 2010](#)).

The last few decades have witnessed a meteoric rise of stellar surveys, which have obtained spectra, light curves, and astrometric solutions for millions of stars across the Milky Way. The revolution began with the Hipparcos mission ([ESA 1997](#)), which measured parallaxes and proper motions of $\sim 10^5$ bright stars within ~ 100 pc. Hipparcos also provided full binary astrometric solutions for a few hundred sources, fitting the plane-of-the-sky motion of these sources with an astrometric model that included elliptical motion of the photocenter, in addition to motion due to parallax and proper motion. Meanwhile, spectroscopic surveys including SDSS, RAVE, LAMOST, APOGEE, and GALAH obtained single- and multi-epoch spectra of several million stars, including thousands of short-period binaries ([York et al. 2000](#); [Steinmetz et al. 2006](#); [Allende Prieto et al. 2008](#); [De Silva et al. 2015](#); [Kunder et al. 2017](#); [Majewski et al. 2017](#)). The CoRoT, Kepler, and K2 mission produced light curves with millimag precision for a few $\times 10^5$ stars. The “highlight” of these missions was the discovery of transiting exoplanets, but large catalogs of close binaries, discovered through eclipses as well as ellipsoidal variation, doppler beaming, reflection, gravitational microlensing, and tidally-excited oscillations, were a happy byproduct for the binary field.

The genesis of this thesis in 2018 was motivated by several planned and just-released datasets. After decades of planning and commissioning, the LIGO collaboration had recently detected its first gravitational wave events: mergers of massive ($\sim (10-30) M_{\odot}$) binary black holes at cosmological distances (e.g. [Abbott et al. 2016](#)). Most of the black holes detected by LIGO were more massive than those known to exist in local X-ray binaries (e.g. [Remillard & McClintock 2006](#)). Their discovery raised broader questions about the black hole population: how many other black holes might lurk in detached binaries in the Milky Way? Within a few years, a binary neutron star merger produced the first gravitational wave event with an electromagnetic counterpart (e.g. [Abbott et al. 2017](#)), generating interest in the system’s

formation history and progenitors. These discoveries motivated me to search for binaries containing compact objects and/or interacting massive stars in the Milky Way.

In 2018, the 2nd data release of the *Gaia* mission (Gaia Collaboration et al. 2018b) revolutionized stellar and Galactic astronomy. Before *Gaia*, stellar distances were hard to come by and plagued by systematic errors. Suddenly, precise distances were available for most stars within about a kiloparsec. Poorly measured distances had been a dominant source of uncertainty in modeling several classes of interacting binaries (e.g. cataclysmic variables and X-ray binaries). Without well-measured distances, it had also been impossible to distinguish the widest gravitationally bound binaries from chance alignments.

Finally, the launch of the TESS satellite in 2018 (Ricker et al. 2015) promised to deliver high-precision light curves for ~ 20 million bright stars. Unlike Kepler, which during its primary mission monitored only about 0.25% of the sky, TESS surveys the whole sky every 24 months, with most sources receiving ~ 27 days of consecutive coverage. Particularly for bright sources (which are easier to follow-up spectroscopically), the TESS data were expected to be a treasure trove for study of close and interacting binaries.

This thesis represents an early step in leveraging large-scale surveys for binary star population studies. A host of upcoming surveys will shine new light on binaries and stellar astrophysics in general in the next decade. eROSITA will discover new populations of low-luminosity X-ray binaries, including wind-transfer BH binaries (Doroshenko et al. 2014). The Rubin observatory’s LSST survey will reveal millions of close binaries throughout the Local Group, including large samples of massive interacting binaries (Prša et al. 2011); it is also expected to detect thousands of isolated stellar-mass BH candidates via microlensing (Abrams & Takada 2020). The Roman telescope microlensing survey will also detect BH candidates, with better mass constraints than LSST (Lam et al. 2020). In the coming years, I will use these surveys both for binary population demographics and to select objects for detailed follow-up. I am particularly invested in the SDSS-V survey (Kollmeier et al. 2017), which will provide broad opportunities to advance the search for interesting binaries and BHs beginning this year. I am leading SDSS-V open fiber programs to obtain BOSS and APOGEE spectra of $\sim 50,000$ light-curve selected massive binaries and BH candidates and $\sim 5,000$ white dwarf + main sequence wide binaries, continuing the work described in Chapter 10. Stellar binarity is among the oldest fields in astrophysics, but the field continues to revolutionize itself. As one co-chair of this thesis told me, binary stars are the new black!

1.2 Thesis summary

This thesis is organized as follows. Chapters 2 and 3 concern the spectroscopic detection and characterization of binaries with wide-field surveys. Chapter 4 introduces a method to distinguish spatially resolved wide binaries from chance alignments using data from the *Gaia* mission. Chapter 5 analyzes the *Gaia* binary sample to constrain the metallicity-dependence of the binary fraction. Chapter 6 outlines a method for testing modifications to the gravitational force law in the low-acceleration regime with wide binaries. Chapter 7

is focused on the excess population of equal-mass “twin” binaries and on the binary mass ratio distribution in general. Chapters 8 and 9 examine two unusual binaries proposed to contain dormant black holes. Chapter 10 uses wide binaries to constrain the astrometric precision of the *Gaia* survey. Chapter 11 introduces a new type of mass-transfer binary that is transitional between the known populations of cataclysmic variables and low-mass white dwarfs.

1.3 Other work

It is said that a thesis which goes exactly as planned is a boring thesis. This one was not boring: I changed my focus more than once and had the opportunity to explore a variety of subfields across astrophysics. My other work focused primarily on galaxy formation and creative uses of survey data. A bibliography of papers I (co)-authored while I was Ph.D. student that are not included in this thesis is below.

41. **El-Badry, K.**, Ostriker, E. O., Kim, C.-G., Quataert, E., Weisz, D. R., 2019, “Evolution of supernovae-driven superbubbles with conduction and cooling”, arXiv:1902.09547, MNRAS, 490, 1961.
40. **El-Badry, K.**, Rix, H.-W., Weisz, D. R. 2018, “An empirical measurement of the initial-final mass relation with Gaia white dwarfs”, arXiv:1805.05849, ApJL, 860, 17.
39. **El-Badry, K.**, Quataert, E., Weisz, D. R., Choksi, N., Boylan-Kolchin, M. 2019, “The formation and hierarchical assembly of globular cluster populations”, arXiv:1805.03652, MNRAS, 482, 4528.
38. **El-Badry, K.**, Bland-Hawthorn, J., Wetzel, A., Quataert, E., Weisz, D. R., Boylan-Kolchin, M., Hopkins, P. F., Faucher-Giguère, C.-A., Kereš, D., Garrison-Kimmel, S. 2018, “Where are the most ancient stars in the Milky Way?”, arXiv:1804.00659, MNRAS, 480, 652.
37. **El-Badry, K.**, Bradford, J., Quataert, E., Geha, M., Boylan-Kolchin, M., Weisz, D. R., Wetzel, A., Hopkins, P. F., Chan, T. K., Fitts, A., Kereš, D., Faucher-Giguère, C.-A. 2018, “Gas kinematics in FIRE simulated galaxies compared to spatially unresolved HI observations”, arXiv:1801.03933, MNRAS, 477, 1536.
36. **El-Badry, K.**, Quataert, E., Wetzel, A., Hopkins, P. F., Weisz, D. R., Chan, T. K., Fitts, A., Boylan-Kolchin, M., Kereš, D., Faucher-Giguère, C.-A., Garrison-Kimmel, S. 2018, “Gas kinematics, morphology, and angular momentum in the FIRE simulations”, arXiv:1705.10321, MNRAS, 473, 1930.
35. **El-Badry, K.**, Weisz, D. R., Quataert, E. 2017, “The statistical challenge of constraining the low-mass IMF in Local Group dwarf galaxies”, arXiv:1701.02347, MNRAS, 468, 319.

34. **El-Badry, K.**, Wetzel, A., Geha, M., Quataert, E., Hopkins, P. F., Kereš, D., Chan, T. K., Faucher-Giguère, C.-A. 2017, “When the Jeans do not fit: How stellar feedback drives stellar kinematics and complicates dynamical modeling in low-mass galaxies”, arXiv:1610.04232, ApJ, 835, 193.
33. **El-Badry, K.**, Wetzel, A., Geha, M., Hopkins, P. F., Kereš, D., Chan, T. K., Faucher-Giguère, C.-A. 2016, “Breathing FIRE: How stellar feedback drives radial migration, rapid size fluctuations, and population gradients in low-mass galaxies”, arXiv:1512.01235, ApJ, 820, 131.
32. Nelson, T., Ting, Y.-S., Hawkins, K., Ji, A., Kamdar, H., **El-Badry, K.**, 2021, “Distant relatives: The chemical homogeneity of comoving pairs identified in Gaia”, arXiv:2104.12883, ApJ, submitted.
31. Stern, J., Sternberg, A., Faucher-Giguère, C.-A., Hafen, Z., Fielding, D., Quataert, E., Wetzel, A., Anglès-Alcàzar, D., **El-Badry, K.**, Kereš, D., Hopkins, P. F., 2021, “Neutral CGM as damped Ly α absorbers at high redshift”, MNRAS, submitted.
30. Santistevan, I., Wetzel, A., Sanderson, R., **El-Badry, K.**, Samuel, J., Faucher-Giguère, C.-A., “The origin of metal-poor stars on prograde disk orbits in FIRE simulations of Milky Way-mass galaxies”, MNRAS, submitted.
29. Martin, D. V., **El-Badry, K.**, Hodžić, V. K., Triaud, A. H. M. J., Angus, R., Birky, J., Foreman-Mackey, D., Hedges, C., Montet, B., Murphy, S. J., Santerne, A., Stassun, K. G., Stephan A. P., Wang, J., Benni, P., Krushinsky, V., Chazov, N., Mishevskiy, N., Ziegler, C., Soubkiou, A., Benkhaldoun, Z., Caldwell, D. A., Collins, K., Henze, C. E., Guerrero, N. M., Jenkins, J. M., Latham D. W., Levine, A., McDermott, S., Mullally, S. E., Ricker, G., Seager, S., Shporer, A., Vanderburg, A., Vanderspek, R., Winn, J. N., 2021, “TOI-1259Ab – a gas giant with 2.6% deep transits and a bound white dwarf companion”, arXiv:2101.02707, MNRAS, submitted.
28. Mercado, F. J., Bullock, J. S., Boylan-Kolchin, M., Moreno, J., Wetzel, A., **El-Badry, K.**, Graus, A. S., Fitts, A., Hopkins, P. F., Faucher-Giguère, C.-A., 2020, “Totally metal: A relationship between stellar metallicity gradients and galaxy age in dwarf galaxies”, arXiv:2009.01241, MNRAS, 501, 5121.
27. Velázquez, J. F., Gurvich, A. B., Faucher-Giguère, C.-A., Bullock, J. S., Starkenburg, T. K., Moreno, J., Lazar, A., Mercado, F. J., Stern, J., Sparre, M., Hayward, C., Wetzel, A., **El-Badry, K.**, 2020 “The time-scales probed by star formation rate indicators for realistic, bursty star formation histories from the FIRE simulations”, arXiv:2008.08582, MNRAS, 501, 4812.
26. Xiang, M.-S., Rix, H.-W., Ting, Y.-S., Zari, E., **El-Badry, K.**, Yuan, H.-B., Cui, W.-Y., 2020, “Data-driven spectroscopic estimates of absolute magnitude, distance,

- and binarity — method and catalog of 16,002 O- and B-type stars from LAMOST”, arXiv:2008.10637, ApJ, in press.
25. Irrgang, A., Geier, S., Heber, U., Kupfer, T., **El-Badry, K.**, Bloemen, S., 2020, “A proto-helium white dwarf stripped by a substellar companion via common-envelope ejection: Uncovering the true nature of a candidate hypervelocity B-star”, arXiv:2007.03350, A&A., in press.
 24. Kamdar, H., Conroy, C., Ting, Y.-S., **El-Badry, K.**, 2020, “Spatial and kinematic clustering of stars in the Galactic disk”, arXiv:2007.10990, ApJ, submitted.
 23. Stern, J., Faucher-Giguère, C.-A., Fielding, D., Quataert, E., Hafen, Z., Gurvich, A. B., Ma, X., Byrne, L., **El-Badry, K.**, Anglès-Alcàzar, D., Chan, T.-K., Feldmann, R., Kereš, D., Wetzel, A., Murray, N., Hopkins, P. F., 2020, “Virialization of the inner CGM in the FIRE simulations and implications for galaxy discs, star formation and feedback”, arXiv: 2006.13976, ApJ, in press.
 22. Li, F., Rahman, M., Murray, N., Hafen, Z., Faucher-Giguère, C.-A., Stern, J., Hummels, C. B., Hopkins, P. F., **El-Badry, K.**, Kereš, D., 2020, “Probing the CGM of low-redshift dwarf galaxies using FIRE simulations”, MNRAS, 500, 1038.
 21. Lazar, A., Bullock, J. S., Boylan-Kolchin, M., Chan, T.-K., Hopkins, P. F., Graus, A., Wetzel, A., **El-Badry, K.**, Wheeler, C., Straight, M. C., Kereš, D., Faucher-Giguère, C.-A., Fitts, A., Garrison-Kimmel, S., 2020, “A dark matter profile to model diverse feedback-induced core sizes of Λ CDM haloes”, arXiv:2004.10817, MNRAS, 497, 2393.
 20. Coronado, J., Rix, H.-W., Trick, W., **El-Badry, K.**, Rybizki, J., Xiang, M., 2020, “From birth associations to field stars: mapping the small-scale orbit distribution in the Galactic disc”, arXiv:2002.09496, MNRAS, 495, 4098.
 19. Santistevan, I. B., Wetzel, A., **El-Badry, K.**, Bland-Hawthorn, J., Boylan-Kolchin, M., Bailin, J., Faucher-Giguère, C.-A., Benincasa, S., 2020, “Growing pains: the formation times and building blocks of Milky Way-mass galaxies in the FIRE simulations”, arXiv:2001.03178, MNRAS, 497, 747.
 18. Pelliccia, D., Mobasher, B., Darvish, B., Lemaux, B. C., Lubin, L. M., Hirtenstein, J., Shen, L., Wu, P.-F., **El-Badry, K.**, Wetzel, A., Jones, T., 2020, “Effects of stellar feedback on stellar and gas kinematics of star-forming galaxies at $0.6 < z < 1.0$ ”, arXiv:2001.00590, ApJL, 896, 26.
 17. Hafen, Z., Faucher-Giguère, C.-A., Anglès-Alcàzar, D., Stern, J., Kereš, D., Esmerian, C., Wetzel, A., **El-Badry, K.**, Chan, T.-K., Murray, N., 2019, “The fates of the circumgalactic medium in the FIRE simulations”, arXiv:1910.01123, MNRAS, 494, 3581.

16. Tian, H.-J., **El-Badry, K.**, Rix, H.-W., Gould, A., 2019, “The separation distribution of ultrawide binaries across galactic populations”, arXiv:1909.04765, ApJS, 246, 4.
15. Hawkins, K., Lucey, M., Ting, Y.-S., Ji, A., Katzberg, D., Thompson, M., **El-Badry, K.**, Teske, J., Nelson, T., Carrillo, A., 2019, “Identical or fraternal twins?: The chemical homogeneity of wide binaries from *Gaia* DR2”, arXiv:1912.08895, MNRAS, 492, 1164.
14. Jahn, E. D., Sales, L. V., Wetzel, A., Boylan-Kolchin, M., Chan, T.K., **El-Badry, K.**, Lazar, A., Bullock, J. S., 2019, “Dark and luminous satellites of LMC-mass galaxies in the FIRE simulations”, MNRAS, 489, 5348.
13. Samuel, J., Wetzel, A., Tollerud, E., Garrison-Kimmel, S., Loebman, S., **El-Badry, K.**, Hopkins, P.F., Boylan-Kolchin, M., Faucher-Giguère, C.-A., Bullock, J., Benincasa, S., Bailin, J., 2019, “A profile in FIRE: resolving the radial distributions of satellite galaxies in the Local Group with simulations”, arXiv:1904.11508, MNRAS, 491, 1471
12. Garrison-Kimmel, S., Wetzel, A., Hopkins, P. F., Sanderson, R., **El-Badry, K.**, Graus, A., Chan, T.K., Feldmann, R., Boylan-Kolchin, M., Hayward, C., Bullock, J. S., Fitts, A., Samuel, J., Wheeler, C., Kereš, D., Faucher-Giguère, C.-A., 2019, “Star formation histories of dwarf galaxies in the FIRE simulations: dependence on mass and Local Group environment”, arXiv:1903.10515, MNRAS, 489, 4574.
11. Dickey, C. M., Geha, M., Wetzel, A., **El-Badry, K.**, 2019, “AGN all the way down? AGN-like line ratios are common in the lowest-mass isolated quiescent galaxies”, arXiv: 1902.01401, ApJ, 884, 180.
10. Emami, N., Siana, B., Weisz D. R., Johnson, B. D., Ma, X., **El-Badry, K.**, 2018, “A closer look at bursty star formation with $L_{H\alpha}$ and L_{UV} distributions”, arXiv:1809.06380, ApJ, 881, 71.
9. Fitts, A., Boylan-Kolchin, M., Bozek, B., Bullock, J. S., Graus, A., Robles, V., Hopkins P. F., **El-Badry, K.**, Garrison-Kimmel, S., Faucher-Giguère, C.-A., Wetzel, A., Kereš, D., 2018, “Dwarf galaxies in CDM, WDM, and SIDM: disentangling baryons and dark matter physics”, arXiv: 1811.11791, MNRAS, 490, 962.
8. Hafen, Z., Faucher-Giguère, C.-A., Anglès-Alcàzar, D., Stern, J., Kereš, D., Hummels, C., Esmerian, C., Garrison-Kimmel, S., **El-Badry, K.**, Wetzel, A., Chan, T. K., Hopkins, P. F., Murray, N., 2018, “The origins of the circumgalactic medium in the FIRE simulations”, arXiv:1811.11753, MNRAS, 488, 1.
7. Hirtenstein, J., Jones T., Wang, X., Wetzel, A., **El-Badry, K.**, Hoag, A., Treu, T., Bradač, M., Morishita, T., 2018, “The OSIRIS lens-amplified survey (OLAS) I: dynamical effects of stellar feedback in low mass galaxies at $z \sim 2$ ”, arXiv:1811.11768, ApJ, 880, 54.

6. Garrison-Kimmel, S., Hopkins, P. F., Wetzel, A., Bullock, J., Boylan-Kolchin, M., Kereš, D., Faucher-Giguère, C.-A., **El-Badry, K.**, Lamberts, A., Quataert, E., Sanderson R. E., 2018, “The Local Group on FIRE: Dwarf galaxy populations across a suite of hydrodynamic simulations”, arXiv:1806.04143, MNRAS, 487, 1380.
5. Debattista, V. P., Gonzalez O. A., Sanderson R. E., **El-Badry, K.**, Garrison-Kimmel, S., Wetzel, A., Faucher-Giguère, C.-A., Hopkins, P. F., 2018, “Formation, vertex deviation and age of the Milky Way’s bulge: input from a cosmological simulation with a late-forming bar”, arXiv:1805.12199, MNRAS, 485, 5073.
4. Fitts, A., Boylan-Kolchin, M., Bullock, J., Weisz, D. R., **El-Badry, K.**, Wheeler, C., Faucher-Giguère, C.-A., Quataert, E., Hopkins, P. F., Kereš, D., Wetzel, A., 2018, “No assembly required: mergers are mostly irrelevant for the growth of low-mass dwarf galaxies”, arXiv:1801.06187, MNRAS, 479, 319.
3. Garrison-Kimmel, S., Hopkins, P. F., Wetzel, A., **El-Badry, K.**, Sanderson R. E., Bullock, J., Ma, X., van de Voort, F., Hafen, Z., Faucher-Giguère, C.-A., Hayward, C. C., Quataert, E., Kereš, D., Boylan-Kolchin, M., 2018, “The origin of the diverse morphologies and kinematics of Milky Way-mass galaxies in the FIRE-2 simulations”, arXiv:1712.03966, MNRAS, 481, 4133.
2. Chan, T. K., Kereš, D., Wetzel, A., Hopkins, P. F., Faucher-Giguère, C.-A., **El-Badry, K.**, Garrison-Kimmel, S., Boylan-Kolchin, M. 2017, “The origin of ultra diffuse galaxies: stellar feedback and quenching”, arXiv:1711.04788, MNRAS, 478, 906.
1. Hopkins, P. F., Wetzel, A., Kereš, D., Faucher-Giguère, C.-A., Quataert, E., Boylan-Kolchin, M., Murray, N; Hayward, C. C., **El-Badry, K.** 2017, “How to model supernovae in simulations of star and galaxy formation”, arXiv:1707.07010, MNRAS, 477, 1578.

Chapter 2

Signatures of unresolved binaries in stellar spectra: implications for spectral fitting

An earlier version of this article was previously published as El-Badry, K., Rix, H.-W., Ting, Y.-S., Weisz, D. R., Bergemann, M., Cargile, P., Conroy, C., Eilers, A.-C., 2018, MNRAS, 473, 5043.

2.1 Abstract

The observable spectrum of an unresolved binary star system is a superposition of two single-star spectra. Even without a detectable velocity offset between the two stellar components, the combined spectrum of a binary system is in general different from that of either component, and fitting it with single-star models may yield inaccurate stellar parameters and abundances. We perform simple experiments with synthetic spectra to investigate the effect of unresolved main-sequence binaries on spectral fitting, modeling spectra similar to those collected by the APOGEE, GALAH, and LAMOST surveys. We find that fitting unresolved binaries with single-star models introduces systematic biases in the derived stellar parameters and abundances that are modest but certainly not negligible, with typical systematic errors of 300 K in T_{eff} , 0.1 dex in $\log g$, and 0.1 dex in $[\text{Fe}/\text{H}]$ for APOGEE-like spectra of solar-type stars. These biases are smaller for spectra at optical wavelengths than in the near-infrared. We show that biases can be corrected by fitting spectra with a binary model, which adds only two labels to the fit and includes single-star models as a special case. Our model provides a promising new method to constrain the Galactic binary population, including systems with single-epoch spectra and no detectable velocity offset between the two stars.

2.2 Introduction

Roughly half of all solar-type stars reside in binary systems (Raghavan et al. 2010; Moe & Di Stefano 2017a). Outside the solar neighborhood ($d \gtrsim 100$ pc), most binaries are spatially unresolved and will fall within the same fiber in spectroscopic surveys. Short-period binaries are easily identified spectroscopically by a line-of-sight velocity offset between the two components (Pourbaix et al. 2004; Matijević et al. 2010; Fernandez et al. 2017; Merle et al. 2017) or time-variable radial velocity measurements (Troup et al. 2016; Gao et al. 2017). However, most solar-type binary systems have orbital periods of hundreds to thousands of years (see Duchêne & Kraus 2013, for a review) and are missed by binarity detection methods based on the Doppler shift.

Absent a detectable velocity offset, the spectrum of an unresolved binary still contains signatures of the unseen companion. These are strongest when the two components have very different spectral features, as is the case with white dwarf – main sequence binaries (Rebassa-Mansergas et al. 2007; Ren et al. 2013). However, the presence of an unresolved companion changes the observable spectrum in *all* cases where the two stars are not identical.

This is on one hand useful for the study of binary systems, as it implies that many binaries should be detectable from detailed modeling of their spectra, even without a detectable velocity offset. However, it also presents a challenge for Galactic surveys aiming to measure the properties of stars by fitting their spectra. When the spectrum of an unresolved binary is fit with single-star models, contamination from the unseen companion may introduce biases in the inferred atmospheric parameters and abundances. Quantifying and remedying these biases is the subject of this paper.

2.3 Methods

In order to disentangle the effects of unresolved binaries from systematic errors in stellar models, we conduct our analysis with synthetic spectra.

We generate spectra for single stars using 1D-LTE atmosphere models computed by the ATLAS12 code (Kurucz 1970, 1979, 1992). Radiative transfer calculations are carried out with SYNTHE (Kurucz 1993). We vary (and later fit for) four stellar labels: T_{eff} , $\log g$, $[\text{Fe}/\text{H}]$, and $[\alpha/\text{Fe}]$, where Mg, Si, S, Ar, Ca, and Ti are included in α and are varied in lockstep.

To efficiently calculate the synthetic spectrum at an arbitrary point in label space (as is required during fitting) we interpolate between nearby models at each wavelength pixel using a neural network. This approach, which makes it possible to self-consistently fit for many stellar labels simultaneously, is similar to that developed in Ting et al. (2016a) and Rix et al. (2016). It was introduced in Ting et al. (2017b) and will be fully explained in Ting et al. (2017c, in prep). In brief, we first generate an irregular grid of ab-initio Kurucz model spectra throughout the region of label space within which a spectral model is desired. These spectra are then used as a training set to train an individual neural network to predict the flux at each wavelength pixel as a function of labels. The full spectral model then consists

of all the individual neural networks for all wavelength pixels stitched together. We have verified through cross-validation that typical interpolation errors in model fluxes are small ($\ll 1$ percent). We produce synthetic spectra with the wavelength coverage and typical S/N per pixel of three surveys:

- APOGEE, with $R = 22,500$, $S/N = 100$, and coverage across $1.51 - 1.70 \mu\text{m}$, with some gaps (Holtzman et al. 2015).
- GALAH, with $R = 28,000$, $S/N = 100$, and coverage over four $\sim 250 \text{ \AA}$ channels between $0.47 \mu\text{m}$ and $0.79 \mu\text{m}$ (De Silva et al. 2015).
- LAMOST, with $R = 1,800$, $S/N = 30$, and coverage across $0.37 - 0.90 \mu\text{m}$ (Cui et al. 2012a).

These surveys were chosen to span a range of resolution, S/N , and wavelength coverage representative of modern spectroscopic surveys. A primary goal of all three surveys is to derive precise stellar parameters and abundances for large samples of stars in order to study the formation and enrichment history of the Galaxy.

Synthetic spectra are calculated at $R = 300,000$, degraded to lower resolutions assuming a Gaussian line spread function (LSF) with $\text{FWHM} = \lambda/R$, and resampled on the wavelength grid of each survey. Best-fit labels are determined through full-spectrum fitting with χ^2 minimization on normalized spectra. We normalize APOGEE-like spectra using the *Cannon*-type continuum normalization routine from the APOGEE package (Bovy 2016a). For LAMOST- and GALAH-like spectra, we normalize by dividing the spectrum by a smoothed version of itself, as in Ho et al. (2017) and Ting et al. (2017b). For our idealized experiments, the choice of normalization method has a negligible effect on the best-fit stellar parameters. We emphasize that we make no attempt to closely reproduce the LSF, continuum normalization, or fitting procedure of the surveys we emulate. Rather, we seek to qualitatively measure how unresolved binarity changes stellar spectra at different resolutions and wavelengths, without making strict assumptions about implementation.

We have verified that our approach can recover the labels of single-star spectra with added noise with high fidelity. We also experimented with generating and fitting spectra using NLTE spectra computed using the 1D hydrostatic MAFAGS-OS model atmospheres (Grupp 2004; Bergemann et al. 2012) instead of ATLAS12 and SYNTHE, obtaining very similar results for the effects of unresolved binarity.

We emphasize that the choice of spectral model and normalization method have negligible effects on the best-fit stellar parameters for each survey only because the spectral model and the spectra to be fit are always generated self-consistently. For example, rather than predicting spectra normalized with the “true” continuum, the model for normalized spectra predicts spectra normalized in the same way as the spectra to be fit (for each mock-survey). In cases where there is a mismatch between the model and the spectra being fit, changes to the normalization procedure and other details of the spectral model will have more significant effects.

2.4 Fitting binary spectra with single-star models

Figure 2.1 illustrates the effects of unresolved binarity on a small wavelength range of an APOGEE-like spectrum for a typical K-dwarf star ($T_{\text{eff}} = 5000$ K, $\log g = 4.59$, $[\text{Fe}/\text{H}] = -0.2$). The top panel shows the spectrum of the primary (blue), secondary (red), and their sum (black). We always assume that the secondary has the same age and abundances as the primary (neglecting the subtle effects of atomic diffusion; see Dotter et al. 2017), so that its spectrum is completely determined by the mass ratio $q = m_2/m_1$. Atmospheric parameters of the secondary ($T_{\text{eff}} = 4100$ K, $\log g = 4.72$ for the star in Figure 2.1) are calculated for a given mass ratio using MIST isochrones (Choi et al. 2016).

In the middle panel, we show the continuum-normalized spectrum of the binary and primary star, as well as the result of fitting the binary spectrum with a single-star model. The corresponding residuals are shown in the bottom panel. Because the primary and secondary stars have different temperatures, they have noticeably different spectral features. For example, the absorption line at 16064 \AA is weaker in the secondary, while the lines at 16057 \AA and 16060 \AA are much weaker in the primary. As a result, the normalized binary spectrum deviates from that of the primary in several lines.

Fitting the unresolved binary spectrum with a single-star model produces a good fit, with significantly smaller residuals compared to the primary star spectrum. However, the labels of the best-fit single-star model differ significantly from the true labels of the primary. This case study shows that fitting a binary spectrum with a single-star model can introduce biases in the derived stellar parameters and abundances. We now investigate these systematically.

We construct mock spectra for 2000 unresolved binaries by combining pairs of main-sequence single stars with masses spanning $(0.2 - 1.7) M_{\odot}$ and ages spanning $(1 - 12)$ Gyr, with $3100 \text{ K} \leq T_{\text{eff}} \leq 7000 \text{ K}$ and $4.0 \text{ dex} \leq \log g \leq 5.05 \text{ dex}$. Masses for both the primary and secondary are drawn from a uniform distribution. Atmospheric parameters for both stars are calculated from MIST isochrones; cases where either star has left the main sequence are discarded. We assign metallicities and α -abundances to each pair simultaneously by drawing randomly from labels of stars in APOGEE DR13 (Holtzman et al. 2015). We draw orbital periods from the lognormal period distribution for solar-type stars found in Duchêne & Kraus (2013); we assume random orbital orientations and phases to calculate line-of-sight velocity offsets for each system, assuming circular orbits. We exclude the $\sim 20\%$ of orbital configurations with velocity offsets $|\Delta v| > 10 \text{ km s}^{-1}$, as these could potentially be identified as double-line binaries in high-resolution spectra and removed from the sample. We then sum the unnormalized single-star spectra, add Gaussian noise to bring the binary spectra to the fiducial S/N of each survey, and fit them with single-star models. The results of this experiment are shown in Figure 2.2. For each mock-survey, we compare the true labels of the primary star to those found by fitting a single-star model to the binary spectrum.

In the limit of low q , systematic errors in labels are small, because the secondary star contributes a negligible fraction of the total light. At $0.4 \lesssim q \lesssim 0.9$, fitting a single-star model leads to a best-fit T_{eff} that is on average too low, because the primary spectrum is contaminated by a cooler companion. For all surveys, systematic errors in all labels except




Figure 2.1: Effect of an unresolved companion on the observable APOGEE-like spectrum of a K-dwarf star. **Top:** synthetic spectrum of the primary, secondary (mass ratio $q = m_2/m_1 = 0.75$), and their sum, with no velocity offset. **Middle:** normalized spectrum of the primary star, unresolved binary, and the best-fit single-star model. **Bottom:** residuals of the spectra of the true primary star and the best-fit single star model with respect to the unresolved binary spectrum. The unresolved binary spectrum deviates subtly but nontrivially from the that of the primary (blue). It can be well-fit by a single-star model (orange), but the inferred stellar labels are biased toward lower temperature and metallicity.

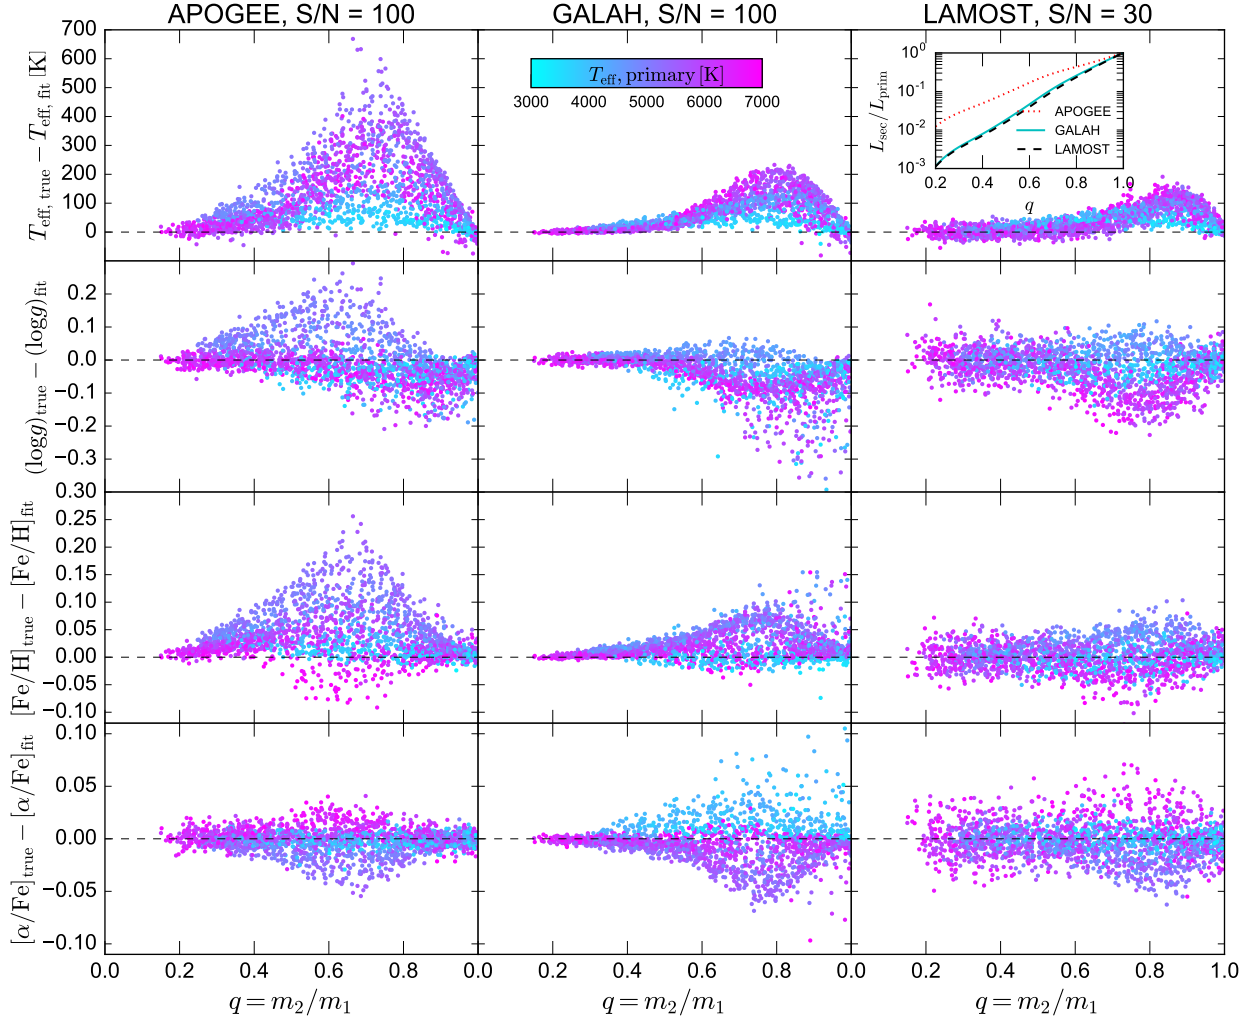


Figure 2.2: Effects of unresolved binarity on the accuracy of stellar labels recovered from spectral fitting. We model 2000 synthetic unresolved main-sequence binaries with a realistic distribution of orbits, mass ratios, and stellar parameters. We produce synthetic spectra with wavelength coverage, resolution, and signal-to-noise per pixel similar to three spectroscopic surveys (Section 2.3) and fit them with single-star models. Unresolved binarity introduces systematic biases in the inferred stellar parameters which are largest at $q \approx 0.7$. Inset shows the luminosity ratio at each survey’s wavelength coverage for a Sun-like primary. Binarity leads to larger systematics for APOGEE-like spectra than for optical spectra because lower-mass companions contribute a larger fraction of the total light in the near-infrared.

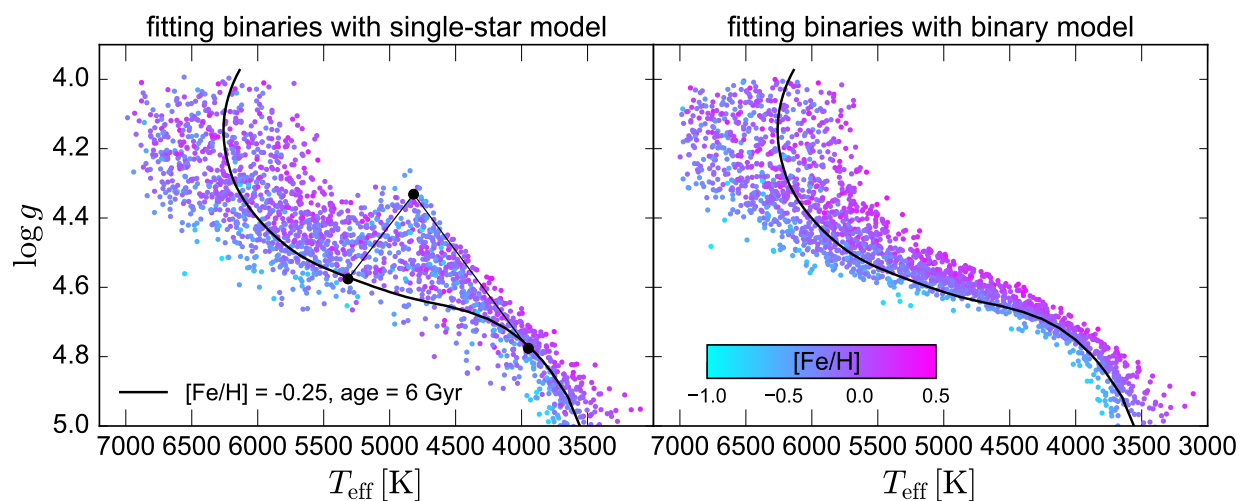


Figure 2.3: Effect of unrecognized binarity on the HR diagram for APOGEE-like spectra. **Left:** We generate synthetic spectra for a population of 2000 unresolved main-sequence binaries and fit them with single-star models, plotting the best-fit parameters. A MIST isochrone is shown for comparison. For one system, black points show the *true* T_{eff} and $\log g$ of the primary and secondary star, as well as the best-fit values obtained by fitting the binary spectrum with a single-star model. Unrecognized binarity causes the inferred stellar parameters of many binaries to scatter above the main sequence and disrupts the monotonic decrease in T_{eff} with metallicity. **Right:** Best-fit parameters for the primary found by fitting spectra with a binary model (see Section 2.5). Fitting a binary model recovers the true parameters without systematic biases.

$\log g$ are largest at $0.6 \lesssim q \lesssim 0.9$. The mass ratio at which binarity has the largest effect is somewhat lower for APOGEE-like spectra than for LAMOST- and GALAH-like spectra, because cooler secondaries contribute more significantly to the total spectrum in the near-infrared than at optical wavelengths. The biases caused by unresolved binarity are subtle but not negligible. For example, for APOGEE-like spectra of solar-type stars with $0.5 < q < 0.8$, the median errors in T_{eff} and $[\text{Fe}/\text{H}]$ are 330 K and 0.1 dex, respectively. In the limit of $q \sim 1$, the two components have identical spectra. However, since we include velocity offsets of up to 10 km s^{-1} (with a median $|\Delta v| = 1.2 \text{ km s}^{-1}$), some biases remain at $q \sim 1$. The largest errors in recovered parameters for $q \sim 1$ binaries are in $\log g$, which is most sensitive to changes in line shape. We have verified that biases at $q \sim 1$ disappear when $\Delta v = 0$.

The strength of the bias in best-fit single star labels varies with the temperature of the primary: systems with low T_{eff} have smaller systematic offsets on average. This is because T_{eff} is almost constant with mass along isochrones low on the main-sequence. In binaries where $T_{\text{eff, primary}} \lesssim 4000 \text{ K}$, the binary and secondary have similar temperatures, and thus, similar spectra.

The typical systematic biases at intermediate q are larger in APOGEE-like spectra than in GALAH- or LAMOST-like spectra due to APOGEE’s redder wavelength coverage. Lower-mass stars are always redder at fixed metallicity, and they therefore contribute a larger fraction of the total light at longer wavelengths. This is illustrated in the inset in the upper-right panel of Figure 2.2, which shows as a function of q the luminosity ratio of the secondary star to the primary, integrated over the wavelength coverage of each survey. At $q = 0.7$, the secondary contributes $\sim 30\%$ as much light as the primary at APOGEE wavelengths, compared to only $\sim 12\%$ and $\sim 11\%$ at the wavelengths probed by GALAH and LAMOST, respectively. Likely for this reason, previous work has found the biases introduced by unresolved binaries to be relatively small at optical wavelengths (Matijević et al. 2010).

Here we only model main-sequence binaries, as binarity will likely have a much smaller effect on stellar parameter estimates for red giant stars. This is because a dwarf secondary will contribute a minuscule fraction of the total light in giant-dwarf binaries, while in giant-giant binaries, the two components will necessarily have very similar masses, and consequently, similar spectra.

In Figure 2.3, we show how unresolved binarity changes the distribution of stars in the best-fit $T_{\text{eff}} - \log g$ plane for APOGEE-like spectra. The left panel shows the results of fitting the spectra of the binary population from Figure 2.2 with single-star models. For most binaries, the best-fit single-star parameters scatter off the main sequence. As one might expect, the spectral contributions of a cooler companion bias the inferred T_{eff} to lower temperatures than the true value for the primary. The best-fit $\log g$ is often lower than the true value for either star. This occurs because $\log g$ is significantly covariant with both T_{eff} and $[\text{Fe}/\text{H}]$ (see e.g. Ting et al. 2017a), and the single-star model can best accommodate the change in the spectral line profiles due to the cooler star by decreasing both T_{eff} and $\log g$. As a result, single-star fits to binary spectra often lie above the main sequence. This suggests that unresolved binarity may be partially responsible for the presence of “main sequence upturn” stars found in many spectroscopic surveys (e.g Kordopatis et al. 2013; Sharma et al.

2017), though limitations in atmosphere models likely also play a role. The right panel of Figure 2.3 shows that fitting spectra with a binary model, as described in the next section, can correct these biases.

2.5 Fitting a binary spectral model

We now demonstrate how the biases discussed in the previous section can be corrected by fitting a binary model; i.e., fitting the spectrum with a sum of two single-star spectra. We construct a binary spectral model by adding two additional labels to the single-star model: the mass ratio q , which sets T_{eff} and $\log g$ of the secondary from isochrones, and the line-of-sight velocity offset Δv . We note that the assumption that the primary and secondary star reside on a theoretical isochrone may be problematic for cool stars. There is in principle no need to enforce this restriction; fitting all parameters of the primary and secondary simultaneously simply comes at the cost of greater model complexity.

In Figure 2.4, we show the results of fitting the synthetic spectrum of the binary shown in Figure 2.1, as it would be observed by each survey, with a binary model. The posterior for each spectrum was sampled with `emcee` (Foreman-Mackey et al. 2013a) with non-informative flat priors, and marginalized projections are generated with `corner` (Foreman-Mackey 2016); constraints on Δv are marginalized over the uncertainties in the velocities of both stars, which are fit independently. Unsurprisingly (since the data are generated and fit with the same model), the binary model is able to recover the true input parameters without systematic biases. The constraining power of spectra from the three surveys is in large part a function of S/N and pixel count – so LAMOST-like spectra provide weaker constraints for most labels than APOGEE- and GALAH-like spectra – but it also depends on the sensitivity of different parts of the spectrum to different labels. Thus, the covariances between different labels in Figure 2.4 are somewhat different for different surveys.

APOGEE- and GALAH like spectra have comparable resolution and adopted S/N , but constraints on most labels are stronger for GALAH-like spectra, in large part due to the higher spectral information content at short wavelengths, where more lines contribute in fixed-length wavelength interval. For the same reason, constraints on T_{eff} and $\log g$ for LAMOST-like spectra are nearly as strong as those from APOGEE-like spectra despite the significantly lower adopted S/N . The 68 and 95% probability contours show that, at the fiducial S/N assumed in our analysis, the surveys we consider are sufficiently informative to constrain the parameters of a binary model with high fidelity (at least at $q \sim 0.75$). The additional labels introduced by the binary model (q and Δv) are not strongly covariant with other labels, though all three surveys show a negative covariance between q and $\log g$; i.e., decreasing q increases $\log g$ of the secondary, which has a similar effect on the total spectrum as increasing $\log g$ of the primary.

This fitting approach can be used to constrain the Galactic binary population: probable binaries can in principle be identified as stars which are significantly better-fit in a χ^2 sense by a binary model than by a single star. An important question is whether binary spectra can

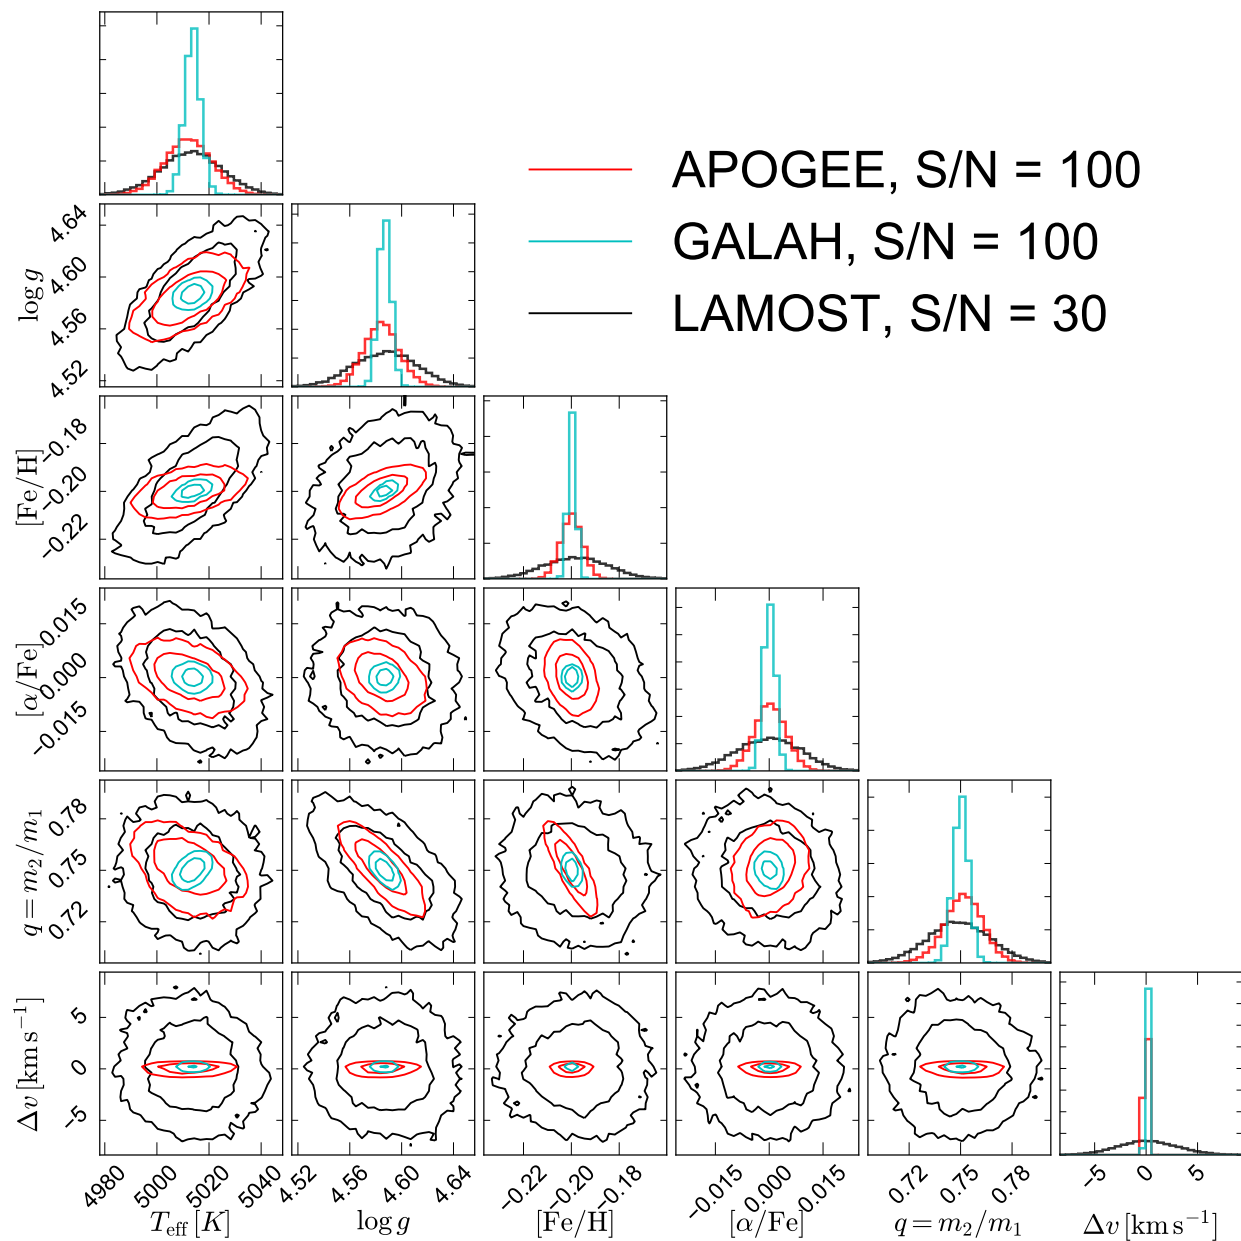


Figure 2.4: 68 and 95% probability contours obtained by fitting a binary model to the synthetic spectrum of an unresolved binary as it would be observed by three spectroscopic surveys. Fitting a binary model corrects the bias that is present when the spectrum is fit by a single-star model and adds only two additional labels: the mass ratio, q , and the velocity offset, Δv . A binary model makes it possible to identify unresolved binaries even in systems where no velocity offset is measurable.

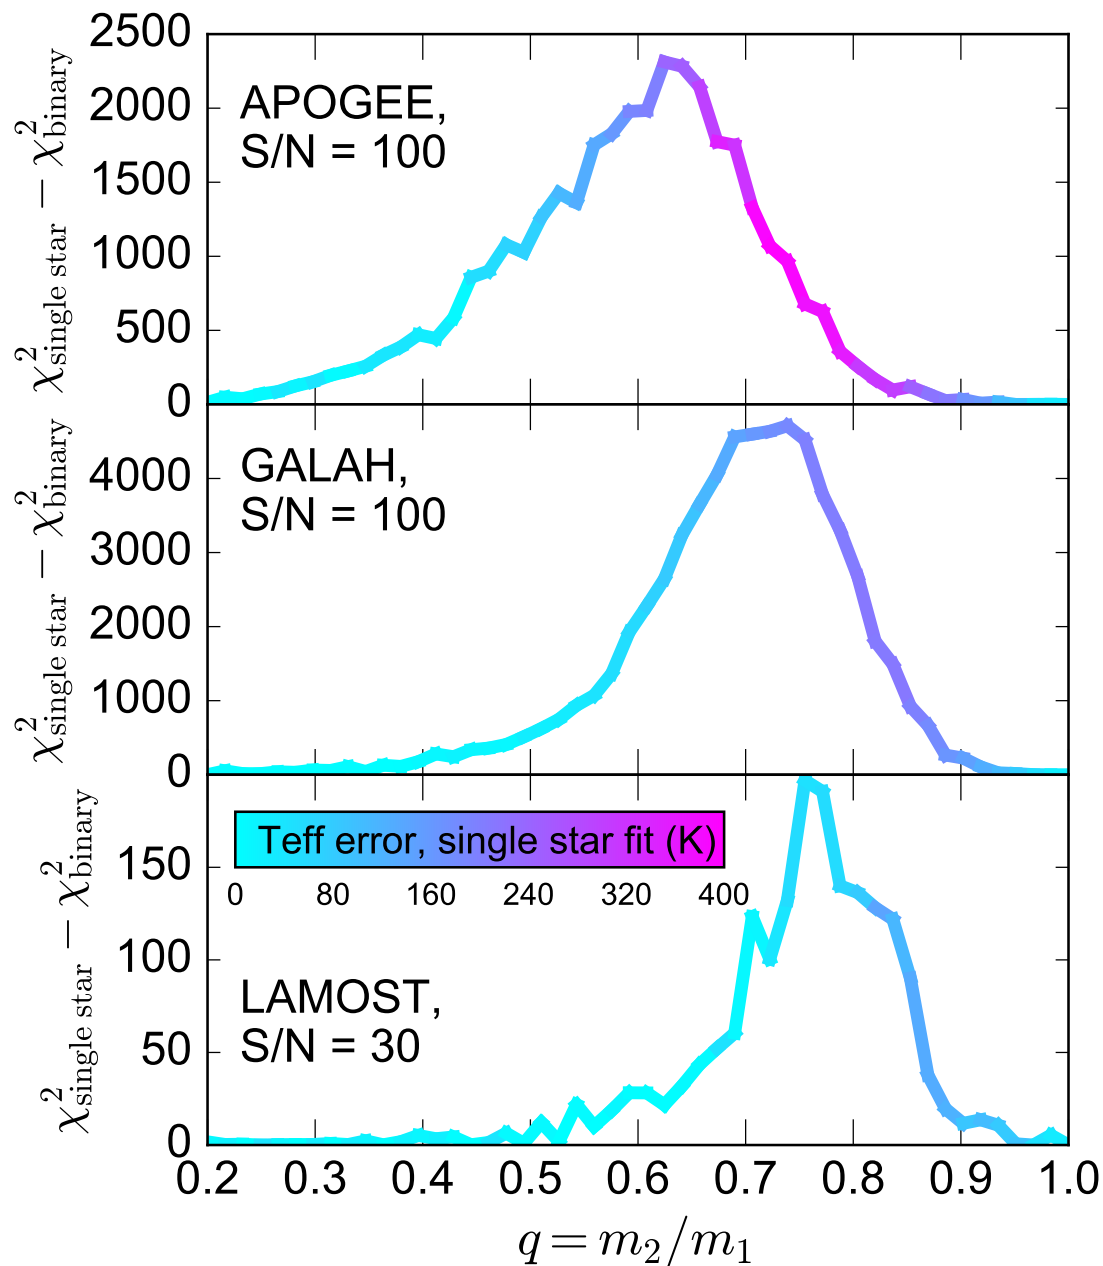


Figure 2.5: Difference in total χ^2 when fitting the synthetic spectrum of an unresolved binary with a Sun-like primary with a single-star model and a binary model (i.e., a sum of two single-star models). At intermediate mass ratios, unresolved binaries can be identified as systems which are significantly better fit by a binary model than a single-star model.

reliably be distinguished from single-star spectra with different parameters. We investigate this in Figure 2.5.

For a Sun-like primary star, we generate synthetic binary spectra with q varying from 0.2 to 1. We set $\Delta v = 0$, since this is the limit in which detecting binarity is most challenging. We then fit each spectrum with both a single-star and a binary model. The difference in χ^2 values between these fits quantifies “how much better” the spectrum is fit with a binary model. At intermediate mass ratios, the χ^2 difference between the best-fit binary and single-star models is significant: $\Delta\chi^2 \gtrsim 100$ for LAMOST-like spectra, and $\Delta\chi^2 \gtrsim 1000$ for APOGEE- and GALAH-like spectra. The q range where binaries are detectable is similar (though not identical) to the range in which fitting a single-star model introduces the largest systematics.

We find that fitting a single-star spectrum with a binary model always converges to $q = 1$ and $\Delta v = 0$, (which is completely degenerate with the single-star case) and recovers the correct labels for the primary star, with $\Delta\chi^2$ of order unity. Cases where the binary model provides a better fit but $\Delta\chi^2$ is small should be interpreted as consistent with being single stars. It is thus straightforward (in the limit of perfect models) to detect binaries with $0.6 \lesssim q \lesssim 0.85$ in LAMOST- and GALAH-like spectra; in APOGEE-like spectra, binaries are in principle detectable at mass ratios as low as $q \sim 0.3$.

These experiments are of course idealized, as they assume no systematic shortcomings in the model spectra. Detecting unresolved binaries in real spectra will present challenges which do not arise with synthetic data. For example, systematic errors in stellar models may cause some single star spectra to be better-fit by a binary model. In this case, care must be taken in formulating a heuristic to determine what $\Delta\chi^2$ constitutes reliable evidence of binarity. A useful approach for fitting real spectra may be to use a data-driven model like *The Cannon* (Ness et al. 2015a) as opposed to synthetic spectra. Another complication not considered in this work is the effect of stellar rotation, which can broaden spectral lines in a manner superficially similar to a velocity offset in a binary system with $q \sim 1$. This effect is small for most cool stars but becomes important at $T_{\text{eff}} \gtrsim 6500$ K (Glebocki et al. 2000). Because the effective kernel for rotational broadening changes spectral line profiles differently from a velocity offset in a binary system, a straightforward solution is to incorporate the projected rotation velocity as another label in the spectral model.

2.6 Conclusions

We have shown that unrecognized binarity in main-sequence stars can introduce biases in stellar labels found by fitting single-star models. These can be significant compared to the target measurement precision of high-resolution spectroscopic surveys (which is $\sigma_{T_{\text{eff}}} \sim 50$ K, $\sigma_{[\text{Fe}/\text{H}]} \sim 0.05$ dex for APOGEE and GALAH) and present a potential challenge for works requiring precise measurements of stellar parameters and abundances of main-sequence stars, e.g., for chemical tagging (Freeman & Bland-Hawthorn 2002; Ting et al. 2015; Hogg et al. 2016; Ting et al. 2016b). These biases can be eliminated by fitting a binary model,

with a modest increase in complexity.

Fitting a binary model also makes it possible to spectroscopically identify main-sequence binary systems, even absent a velocity offset between the two stars: if the two components of a binary system have different temperatures, their combined spectrum will be better-fit by a binary model than a single-star model. The spectral signatures of binarity are stronger in the near-infrared than at optical wavelengths, because a cooler secondary contributes a larger fraction of the total light at longer wavelengths. For APOGEE-like spectra, the signatures of binarity are strongest for systems with $0.4 \lesssim q \lesssim 0.8$. We will present results of fitting APOGEE spectra with a data-driven binary model in forthcoming work.

An important limitation of fitting a binary model is that, without distance information, our method cannot detect binaries with mass ratios near $q = 1$ and small velocity offsets: because the two components of a $q = 1$ binary system have essentially identical spectra, their normalized combined spectrum will be indistinguishable from that of a single star. However, in this case there is also no bias in the inferred stellar parameters and abundances. The degeneracy between a single star and a $q = 1$ binary can also be broken with moderate-accuracy distance information, since at fixed distance, a $q = 1$ binary will be twice as bright as a single star.

Acknowledgements

We thank the anonymous referee for a helpful report. We thank Eliot Quataert, Chao Liu, David Hogg, and Andy Gould for useful conversations and comments. K.E. acknowledges support from the SFB 881 program (A3), a Berkeley Fellowship, a Hellman award for graduate study, and an NSF graduate research fellowship. H.W.R. received support from the European Research Council under the European Union’s Seventh Framework Programme (FP 7) ERC Grant Agreement n. [321035]. Y.S.T is supported by the Australian Research Council Discovery Program DP160103747. D.R.W. is supported by a fellowship from the Alfred P. Sloan Foundation. C.C. acknowledges support from NASA grant NNX15AK14G, NSF grant AST-1313280, and the Packard Foundation.

Chapter 3

Discovery and Characterization of 3000+ Main-Sequence Binaries from APOGEE Spectra

An earlier version of this article was previously published as El-Badry, K., Ting, Y.-S., Rix, H.-W., Quataert, E., Weisz, D. R., Cargile, P., Conroy, C., Hogg, D. W., Bergemann, M., Liu, C., 2018, MNRAS, 476, 528.

3.1 Abstract

We develop a data-driven spectral model for identifying and characterizing spatially unresolved multiple-star systems and apply it to APOGEE DR13 spectra of main-sequence stars. Binaries and triples are identified as targets whose spectra can be significantly better fit by a superposition of two or three model spectra, drawn from the same isochrone, than any single-star model. From an initial sample of $\sim 20,000$ main-sequence targets, we identify $\sim 2,500$ binaries in which both the primary and secondary star contribute detectably to the spectrum, simultaneously fitting for the velocities and stellar parameters of both components. We additionally identify and fit ~ 200 triple systems, as well as ~ 700 velocity-variable systems in which the secondary does not contribute detectably to the spectrum. Our model simplifies the process of simultaneously fitting single- or multi-epoch spectra with composite models and does not depend on a velocity offset between the two components of a binary, making it sensitive to traditionally undetectable systems with periods of hundreds or thousands of years. In agreement with conventional expectations, almost all the spectrally-identified binaries with measured parallaxes fall above the main sequence in the color-magnitude diagram. We find excellent agreement between spectrally and dynamically inferred mass ratios for the ~ 600 binaries in which a dynamical mass ratio can be measured from multi-epoch radial velocities. We obtain full orbital solutions for 64 systems, including 14 close binaries within hierarchical triples. We make available catalogs of stellar parameters, abundances, mass

ratios, and orbital parameters.

3.2 Introduction

About half of solar-type stars are in binary or higher order multiple-star systems (Raghavan et al. 2010; Moe & Di Stefano 2017a). Beyond the Solar neighborhood, most binaries are too close on the sky to be spatially resolved; they appear as single photometric point sources, and both components of binary systems contribute to the spectra observed by spectroscopic surveys.

Spectroscopically identifying such unresolved binaries is straightforward only if the period is relatively short ($P \lesssim 5$ years). In this case, spectra exhibit split or “double” lines if the two components have comparable luminosities (so-called “SB2” systems), and two peaks can be identified in the cross-correlation function (Pourbaix et al. 2004; Fernandez et al. 2017; Merle et al. 2017). Even if the secondary is faint and does not contribute significantly to the spectrum, short-period binaries can be identified from radial velocity variability when multi-epoch spectra are available (“SB1” systems; Minor 2013; Troup et al. 2016; Price-Whelan et al. 2017; Badenes et al. 2017).

However, about half of solar-type binaries have periods exceeding 200 years (Duquennoy & Mayor 1991; Duchêne & Kraus 2013). The typical line-of-sight velocity separation between the two stars in such systems is of order 1 km s^{-1} , while the typical change in the stars’ individual velocities over a one year baseline is of order 0.01 km s^{-1} . Such systems will be missed by binary-detection methods based on the Doppler shift.

Unresolved binarity in main-sequence stars presents both a nuisance and an opportunity for spectroscopic surveys of the Milky Way. Because spectral morphology is a strong function of effective temperature, contamination from a cooler secondary star¹ makes the observable spectrum of an unresolved binary different from that of the primary, and in many cases, different from that of any single star. This means that, if binarity is ignored and all spectra are simply fit with single-star models, biases can be introduced in the stellar parameters and abundances inferred for unrecognized binaries (El-Badry et al. 2018a).

On the other hand, binarity-induced features in stellar spectra can be exploited to detect binaries that could not be detected based on velocity shifts alone: binaries can be identified as systems whose spectrum can be significantly better fit by a binary spectral model (i.e., a sum of two single-star models) than any single-star model. This approach, if it can successfully be applied to large spectroscopic surveys, will make possible systematic study of the Galactic binary population on an unprecedented scale.

El-Badry et al. (2018a, hereafter E18) recently demonstrated that fitting a binary model to *synthetic* APOGEE-like spectra makes it possible to spectroscopically identify many binaries and to simultaneously recover the atmospheric parameters and abundances of both

¹We adopt the convention that the secondary is the less-massive of the two stars (e.g. Duchêne & Kraus 2013). For the equal-age, equal-composition main-sequence binaries that we model, the secondary is always cooler and less luminous.

component stars. In this paper, we apply the method described in E18 to real spectra from DR13 of the APOGEE survey (Majewski et al. 2017). We focus on main-sequence stars, for which the effects of unresolved binarity on the spectrum are typically larger than in giants. We demonstrate that, although the spectral signatures of binarity are strongest in close systems with a large velocity offset between the two stars, binaries with mass ratios $0.4 \leq q \leq 0.8$ can be detected with high fidelity even in the absence of any detectable velocity offset (where $q = m_2/m_1$).

This paper is organized as follows. We describe our spectral model for single and binary stars in Section 3.3 and its application to the combined APOGEE spectra in Section 3.4. In Section 3.3.4, we extend the model to fit multi-epoch spectra of close binaries with detectable velocity changes between visits, calculating dynamical mass ratios from the relative velocities of the two components. We identify and derive parameters for close binaries, triples, and systems with unusual velocity shifts in Section 3.4.2 and derive orbital solutions for the subset of binaries with sufficient visits and phase coverage in Section 3.4.4. We discuss our results and conclude in Section 3.5.

We provide many of the underlying model details in the Appendices. Specifically, in Appendix 3.6, we describe the spectral model; model selection and tests with semi-empirical synthetic binary spectra are described in Appendix 3.7; shortcomings of the model and false positives are discussed in Appendix 3.8, and diagnostics of orbit-fitting convergence are presented in Appendix 3.9. Available catalogs are described in Appendix 3.10.

3.3 Methods

Our binary spectral model depends on two steps: (a) creating a data-driven generative model for single-star spectra (Section 3.3.1), and (b) combining the spectra of two single star models, with a suitable velocity offset (Section 3.3.2). To find candidate binaries, we fit spectra with both single-star and binary models (Section 3.3.3) and identify systems that can be significantly better fit by a binary model (Appendix 3.7).

In this work, we only attempt to fit main sequence stars; i.e., targets with $\log g \geq 4$. We do not attempt to identify binaries in which one star is a giant because in most giant-dwarf binaries, the dwarf secondary will contribute a negligible fraction of the total light, while in giant-giant binaries, two components with the same age will necessarily have similar masses, and thus, quite similar spectra. We note that short-period binaries containing giants can be straightforwardly detected from radial velocity variability (Troup et al. 2016; Badenes et al. 2017), and some giant-subgiant binaries can likely be detected spectroscopically (Section 3.5.1).

3.3.1 Single-star spectral model

We model APOGEE spectra of single stars using a data-driven² generative model to predict the rest-frame normalized flux density at a given wavelength as a function of a set of “labels,” $\vec{\ell}$, which determine the spectrum. Our approach is very similar to that employed by *The Cannon* (Ness et al. 2015a): the spectral model is a fitting function that maps labels to normalized spectra, and the free parameters of this fitting function are determined by optimization on a training set, whose labels are obtained separately or known a priori, e.g., from ab-initio fitting.

The primary difference between our method and existing implementations of *The Cannon* is that, as in Ting et al. (2017b) and E18, we model the normalized flux density at a particular wavelength pixel using an artificial neural network rather than a polynomial function. We find a neural network model to be more flexible than a polynomial and to typically produce smaller errors in model spectra during cross-validation; this formalism, which we refer to as *The Payne*, is described further in Appendix 3.6 and will be explained in detail in Ting et al. (in prep). The full spectral model then consists of all the individual neural networks for all wavelength pixels stitched together.

We predict rest-frame spectra with a single-star model that depends on five labels,

$$\vec{\ell} = (T_{\text{eff}}, \log g, [\text{Fe}/\text{H}], [\text{Mg}/\text{Fe}], v_{\text{macro}}). \quad (3.1)$$

We use $[\text{Mg}/\text{Fe}]$ as a proxy for all “ α -elements.” We experimented with including more elemental abundances as labels, including C, N, O, and Si. We found that this did not substantially change our identification of likely binary targets or their inferred mass ratios, so we opted to use a relatively simple model in the interests of reduced complexity. v_{macro} primarily accounts for the effects of stellar rotation, and is small ($< 10 \text{ km s}^{-1}$) for most stars with $T_{\text{eff}} \lesssim 6000 \text{ K}$. In practice, spectra are not observed in the rest frame, so an additional label v_{Helio} also determines the model spectrum and must be included in fitting. However, our neural network model always predicts spectra in the rest frame; Doppler shifts are applied subsequently.

An ideal training set would contain only stars known to be single a priori. Unfortunately, it is nearly impossible to conclusively rule out the possibility that an unresolved system is a binary.³ We therefore construct a training set by beginning with a random sample of main sequence APOGEE stars and then iteratively removing stars whose spectra can be significantly better fit by the binary model described in Section 3.3.2. The ASPCAP pipeline does not derive reliably calibrated abundances for dwarfs. “Ground truth” labels for

²We also experimented with using synthetic, ab-initio spectral models, but we found them ill-suited for identifying binaries because systematic shortcomings in synthetic models cause almost *all* spectra to be significantly better fit by a sum of two models than a single model.

³The only exception is in the immediate Solar neighborhood ($d \lesssim 8 \text{ pc}$), where a combination of direct imaging and speckle interferometry can resolve nearly all systems where a velocity offset is not detectable (Simons et al. 1996; Reid & Gizis 1997). However, there are only ~ 66 stars in the Solar neighborhood for which binarity can be ruled out with high confidence; of these, only the Sun and Arcturus have been observed by APOGEE.

stars in the training set were derived from ab-initio fitting with single-star models, following a procedure similar to that used by Ting et al. (2017b); see Ting et al. (in prep) for details. For the initial training set, we randomly selected 2000 targets distributed throughout the region of label space within which a spectral model was desired, namely $4200 \text{ K} < T_{\text{eff}} < 7000 \text{ K}$, $4.0 < \log g < 5.0$, $-1 < [\text{Fe}/\text{H}] < 0.5$, $-0.4 < [\text{Mg}/\text{Fe}] < 0.6$, and $0 \text{ km s}^{-1} < v_{\text{macro}} < 45 \text{ km s}^{-1}$. We only attempt to fit targets for which the labels determined from ab-initio fitting lie within this region of parameter space, as (a) we are only interested in main-sequence stars, and (b) the labels determined from ab-initio fitting are less reliable outside this range (Ting et al., in prep).

There is of course no guarantee that the targets in our initial training set are actually single stars. After training the initial model, we therefore fit all spectra in the training set both with the initial single-star model and a binary model (as described in Section 3.3.2) based on this single-star model. We then removed from the training set the ~ 300 targets that could be significantly better fit by a binary-model than a single-star model⁴ and re-trained the single-star model on the resulting “cleaned” training set. We repeated this cleaning and re-training procedure until none of the targets in the training set could be significantly better fit by a binary model. This approach converges quickly: after the second iteration, fewer than 10 targets in the cleaned training set could be significantly better fit by a binary model; after the third iteration, no additional targets in the training set could be significantly better fit by a binary model.

This iterative cleaning procedure likely does not remove all unresolved binaries from the training set: only binaries whose combined spectrum is significantly different from any single-star star spectrum can be identified. For APOGEE-like spectra of solar-type stars with negligible velocity offsets, the range of mass ratios over which binarity is detectable is $0.4 \lesssim q \lesssim 0.85$ (E18). Binaries in the training set with mass ratios outside this range will not contaminate the spectral model, since their spectra are not significantly different from the spectrum of a single star with the labels of the primary.

Our approach would likely not work if binaries dominated the training set, or if the functional form of the spectral model were sufficiently complex to incorporate spectral features due to binarity in the single-star model. Because binaries with spectra that are significantly better fit by a binary model constitute only $\sim 15\%$ of the initial training set and the spectral model is not very complex (we use a small neural network with only 1 hidden layer of 5 neurons), detectable binary spectra are essentially treated as outliers and removed during iterative cleaning, preventing the model from overfitting the signature of unresolved binarity into the single-star model.

3.3.2 Binary spectral model

We assume that both components of a binary system have the same age and composition. Fitting a binary model thus adds three free parameters compared to the single-star model:

⁴Here, we quantified “significantly better fit” as having $\chi_{\text{single star}}^2 - \chi_{\text{binary}}^2 > 1000$. We develop a more detailed threshold for model selection in Appendix 3.7.

the mass ratio, $q = m_2/m_1$, which determines T_{eff} and $\log g$ of the secondary, and v_{macro} and v_{Helio} of the secondary. To model the normalized spectrum of a binary with a particular mass ratio, we estimate T_{eff} and $\log g$ of the secondary using MIST isochrones (Choi et al. 2016),⁵ predict the single-star spectra of the primary and secondary in unnormalized space, apply a Doppler shift, add the two spectra, and finally pseudo-continuum normalize the total spectrum; see E18 for details.

Since the data-driven model for single stars operates on normalized spectra, predicting unnormalized spectra for the primary and secondary requires a model for the pseudo-continuum by which the normalized spectra can be multiplied. We obtain the pseudo-continuum for a single star at a particular point in label space by applying our pseudo-continuum fitting procedure (see Section 3.3.3) to a spectrum produced by a synthetic spectral model trained on Kurucz spectra (Kurucz 1970, 1979, 1993). Synthetic spectra are first produced with units of surface flux density and are then multiplied by the surface area of the star in question, using radii estimated from MIST isochrones. The unnormalized flux density of an unresolved binary system viewed from a distance D is then given by

$$f_{\lambda,\text{binary}} = \frac{1}{D^2} (R_1^2 f_{\lambda,1} + R_2^2 f_{\lambda,2}), \quad (3.2)$$

where R_1 and R_2 represent the radii of the primary and secondary star, and $f_{\lambda,1}$ and $f_{\lambda,2}$ represent their individual flux densities. Because we subsequently normalize $f_{\lambda,\text{binary}}$ prior to fitting, the distance D is an arbitrary scaling factor and does not enter our analysis. In practice, R_1 and R_2 are estimated from T_{eff} , $\log g$, and $[\text{Fe}/\text{H}]$ using a neural network trained on a large grid of MIST isochrones.

Our results are not sensitive to the choice of synthetic model spectra, which sets only the relative flux contribution of the primary and the secondary, because the total binary spectrum is again normalized prior to fitting. We have verified that we obtain similar results by simply defining a continuum for each star as a blackbody with appropriate T_{eff} scaled by the surface area of the star.

For long-period systems with negligible velocity shifts, our model cannot detect binaries with mass ratios $q \lesssim 0.4$, because the secondary contributes a negligible fraction of the total light, or $q \gtrsim 0.85$, because the spectra of the primary and secondary are too similar. In practice, another, often more stringent limit on the lowest detectable mass ratio is set by our spectral model’s minimum T_{eff} of 4200 K. For systems with a hot primary star ($T_{\text{eff}} \gtrsim 6500$ K), this limit is not important, since a secondary with $T_{\text{eff}} < 4200$ K would be too faint to contribute significantly to the spectrum anyway. However, the model’s minimum T_{eff} reduces the range of detectable mass ratios for systems with cooler primaries: for a primary with $T_{\text{eff}} = 5800$ K, the effective minimum q that can be modeled is $q_{\text{min}} \approx 0.62$, while for a primary with $T_{\text{eff}} = 5000$ K, $q_{\text{min}} \approx 0.75$. We discuss this further in Appendix 3.7.1.

⁵In practice, we predict T_{eff} and $\log g$ of the secondary from T_{eff} and $\log g$ of the primary, $[\text{Fe}/\text{H}]$, and q using a neural network trained on a large grid of binary isochrones with $0.01 \leq (\text{age}/\text{Gyr}) \leq 13.5$ and $-1 \leq [\text{Fe}/\text{H}] \leq 0.5$. We have verified through cross validation that typical errors in the thus-estimated parameters of the secondary are small (~ 20 K in T_{eff} and ~ 0.01 dex in $\log g$).

3.3.3 Model fitting

Best-fit labels for binary and single-star models are determined through full-spectrum fitting of normalized spectra in vacuum wavelengths. Pseudo-continuum normalization is carried out using the *Cannon*-type normalization routine from the APOGEE package (Bovy 2016b), which fits a 4th order Chebyshev polynomial to pixels in which the gradient of the data-driven spectral model with respect to the labels is small. Bad pixels and pixels with poor sky subtraction, as flagged in the bitmasks produced by the APRED pipeline (Nidever et al. 2015), are masked during normalization and fitting.

Fitting is carried out using the Scipy `curve_fit` routine, which implements the “trust region reflective” algorithm (Branch et al. 1999) for χ^2 minimization. When fitting a single spectrum with a single-star model, we find that the optimization essentially always converges on the true global minimum, irrespective of the location in label space where it is initialized. However, for the binary model there is an obvious degeneracy: the normalized spectrum of a $q = 1$ binary model is identical to that of a $q = 0$ model in the limit of no velocity offset. Hence, the posterior for a binary model is often bimodal in q , and minimization can sometimes converge on a false local minimum. We therefore initialize ~ 10 separate optimizers with different initial values of q when fitting a binary model. If these do not all converge to the same model, we take as the best model the one that reaches the lowest global χ^2 . We have verified by fitting semi-empirical synthetic binary spectra that this approach converges on the true global minimum in $\sim 99\%$ of all cases (see Appendix 3.7.1).

Most APOGEE targets are observed more than once, with time baselines between individual visits ranging from ~ 1 hour to ~ 1200 days.⁶ Spectra from individual visits are shifted to rest frame and coadded to produce a single combined spectrum with higher S/N than the individual visit spectra by the APSTAR pipeline (Nidever et al. 2015). It is these combined spectra that are fit by the ASPCAP pipeline to derive the stellar parameters and abundances published for the main survey (Holtzman et al. 2015; García Pérez et al. 2016), but the reduced spectra from individual visits are also made publicly available.

Combined spectra are easier to work with than individual visit spectra both because they have higher S/N and because stars are often observed with a different fiber and with a different barycentric velocity at each visit, so that the combined spectrum is less affected by bad pixels, poor sky subtraction, and telluric absorption than the individual visit spectra. We therefore fit the combined spectra rather than spectra from individual visits when possible. However, if a system is an unresolved close binary, the orbital configuration and relative radial velocities of the primary and secondary will change between visits, so that the morphology of the total binary spectrum is different in each visit. In such cases, the combined spectrum does not represent any real physical system, and fitting it can yield biased labels.

⁶The APOGEE observing strategy aims to observe most targets 3 times, over a minimum time baseline of 1 month. Some targets, primarily faint stars, are visited more often to accumulate S/N; some targets in unfavorable locations, such as the Galactic Bulge, are visited only once (Zasowski et al. 2013). Most targets with baselines longer than 1 year, as well as those with multiple visits within 1 night, are targets which were observed initially during the survey commissioning period and again during the main survey.

For this reason, we attempt to fit all targets that may be close binaries using the individual visit spectra rather than the combined spectrum. We identify potential close binaries as targets for which (a) the best-fit model to the combined spectrum is a binary model in which the line-of-sight velocity separation of the two components, Δv_{los} , is greater than 10 km s^{-1} , or (b) the V_{scatter} term calculated from the radial velocities determined by the APSTAR pipeline (Nidever et al. 2015) is greater than 1 km s^{-1} , indicating potential radial velocity variability. Some of these targets, particularly stars with high T_{eff} or low S/N, are single stars with poorly constrained radial velocities, but many are close binary systems. Fitting individual visit spectra for targets with $V_{\text{scatter}} > 1 \text{ km s}^{-1}$ also protects against the possibility of a single star erroneously appearing to be a binary if the radial velocities are calculated incorrectly while creating the combined spectrum; otherwise, coadding two visit spectra with different Doppler shifts could produce a combined spectrum bearing erroneous signatures binarity with $q = 1$.

The number of free parameters to be optimized increases substantially when we fit spectra from many visits simultaneously, since the radial velocities at each visit are all free parameters. This can make the fit more susceptible to convergence on an erroneous local minimum in χ^2 ; we discuss the measures taken to ensure global convergence in this case in Section 3.3.4.

For both single-visit and combined spectra, we inflate the uncertainties of pixels with $\text{S/N} > 200$ to 0.5% (i.e., S/N of 200) during fitting because empirical S/N diagnostics based on the variation in a given pixel across visits show that the noise model underestimates uncertainties for bright stars and is likely limited by systematics at this level (Nidever et al. 2015). We also find that our fitting approach often performs poorly at low S/N, primarily due to poor continuum normalization. We therefore do not attempt to fit any visit spectra with median $\text{S/N} < 30 \text{ pixel}^{-1}$. Since most APOGEE targets are bright, this restriction excludes less than 20% of the targets in our sample; for these targets, we report labels obtained by fitting the combined spectrum, which has higher S/N, but we caution that results for targets with large V_{scatter} and low S/N are likely less reliable.

We do not report uncertainties on labels for individual targets. Formal fitting uncertainties based on the concavity of the likelihood function in the vicinity of the global maximum can be computed with `curve_fit` (e.g. Ness et al. 2015a; Ho et al. 2017), and comparable uncertainties can be obtained by MCMC sampling. However, the thus-obtained uncertainties are typically unrealistically small for high- S/N spectra (e.g., $\sigma(T_{\text{eff}}) < 10 \text{ K}$ for typical APOGEE spectra) because they do not properly account for systematic errors in the spectral model. Systematic errors can arise if (a) the spectral model is not sufficiently complex to account for all the variance in the dataset, (b) there are unaccounted-for errors in the labels assigned to the training set, or (c) the adopted set of labels does not fully characterize all the variance in the dataset. We investigate the typical precision of our best-fit labels in Appendix 3.7.1.

3.3.4 Fitting Multi-Epoch Spectra

We attempt to fit the individual visit spectra rather than the combined spectra of all stars that were visited more than once and are flagged as potential close binaries. In order to fully exploit the information contained in the spectra, we fit all single-visit spectra for each system *simultaneously*, requiring the physical parameters of the component stars to be the same at all epochs. Because we fit all visit spectra with the same spectral model, we implicitly treat the instrumental line spread function as constant across all fibers and visits. For the single-star model, we also require the line-of-sight velocity to be the same at each epoch; in this case, the model is no more complex than when fitting a single combined spectrum.

For an isolated binary system, the line-of-sight velocities of the two components are not independent: in the center-of-mass frame, conservation of linear momentum requires that the radial velocity of the primary along any line of sight, v_1 , and that of the secondary, v_2 , are related by $v_2 = -v_1/q_{\text{dyn}}$, where q_{dyn} is the dynamical mass ratio of the system. If the center-of-mass heliocentric velocity of the binary is γ , then

$$v_{\text{Helio},2} = \gamma + (\gamma - v_{\text{Helio},1})/q_{\text{dyn}}. \quad (3.3)$$

Here v_{Helio} denotes a velocity at a single epoch, measured in the frame of the center of mass of the Solar system.

For true, isolated binary systems containing two main sequence stars, q_{dyn} should be equal to the *spectral* mass ratio q , which determines the contribution of the secondary star to the binary spectrum. We will use q and q_{spec} interchangeably in the rest of this paper. However, we fit q_{dyn} and q_{spec} separately to allow for the possibility of companions whose contribution to the spectrum is different from what is predicted by the dynamical mass ratio. This could occur, for example, if there are biases in the isochrones used in the spectral model, if the secondary falls near the edge of the APOGEE fiber and only a fraction of its flux contributes to the spectrum, or if a third object is present in the system. Comparing the best-fit q_{dyn} and q_{spec} provides a useful diagnostic of the accuracy of our spectral model.

Our basic ‘‘SB2’’ binary model does not allow the velocities of both stars to vary freely, but instead enforces the restriction that the velocities at all epochs follow Equation 3.3 when two or more visit spectra are fit simultaneously. In most cases, this leads to best-fit velocities that are similar (within $\sim 200 \text{ m s}^{-1}$ on average, and nearly always within a few km s^{-1}) to those obtained when Equation 3.3 is not enforced. However, there are some targets for which the best-fit velocities are very different – and produce a much better fit – when Equation 3.3 is not enforced than when it is. Such systems have velocities inconsistent with being a simple two-body system and likely contain a third component. To avoid mischaracterizing these systems, we also fit all targets with a binary model in which the velocities of both components are allowed to vary freely; systems that are significantly better-fit by this model are classified as SB2s with an unseen third component (see Section 3.3.4 for details).

We also find systems in which there is a clear radial velocity shift in the spectrum between different visits but no individual visit spectrum is better-fit by a binary model; i.e., the

existence of a companion can be inferred from its gravitational effects on the primary, but the companion does not significantly contribute to the observed spectrum. Most of these single-line binary (“SB1”) systems are probably ordinary main-sequence binaries with low mass ratios and relatively short periods; some are likely binaries in which the companion is a stellar remnant. To distinguish between SB1s and SB2s, we fit all potential close binary systems with an SB1 model, which is identical to the single-star model, except that the radial velocity is allowed to vary between visits. We designate systems as SB1s if the SB1 and SB2 models converge on essentially the same fit; i.e., if there is no detectable contribution to the spectrum from the secondary.

Finally, we find some systems whose visit spectra cannot be well-fit by any single star or binary model: the binary model provides a better fit than the single-star model, but many lines are poorly fit or are missing entirely from the best-fit binary model. We find that many of these systems can be much better fit by a *triple* model: i.e. three stars with independent velocities and masses, restricted to lie on the same isochrone.

Summary of models fit to visit spectra

We simultaneously fit the N visit spectra for each object in the “potential close binary” subsample with a total of five different models, which we summarize here. We classify systems based on the total χ^2 of each model, preferring the least complex model when different models have similar χ^2 .

1. *Single-star*: The single-star model has 6 free parameters, regardless of the number of visit spectra:

$$\vec{\ell}_{\text{single star}} = (T_{\text{eff}}, \log g, [\text{Fe}/\text{H}], [\text{Mg}/\text{Fe}], v_{\text{macro}}, v_{\text{Helio}}). \quad (3.4)$$

In particular, this model forces the heliocentric velocity of the star to be the same in all visits.

2. *SB1*: The SB1 model is identical the single-star model, except that the heliocentric velocity is allowed to vary between the N visits. The $5 + N$ free parameters are:

$$\vec{\ell}_{\text{SB1}} = (T_{\text{eff}}, \log g, [\text{Fe}/\text{H}], [\text{Mg}/\text{Fe}], v_{\text{macro}}, v_{\text{Helio},i}), \quad (3.5)$$

where i enumerates the visits.

3. *SB2*: The SB2 model fits two stars, with different velocities at each visit, but with the restriction that the velocity satisfy Equation 3.3. The $9 + N$ free parameters are

$$\vec{\ell}_{\text{SB2}} = (T_{\text{eff}}, \log g, [\text{Fe}/\text{H}], [\text{Mg}/\text{Fe}], q, v_{\text{macro1}}, v_{\text{macro2}}, q_{\text{dyn}}, \gamma, v_{\text{Helio},i}). \quad (3.6)$$

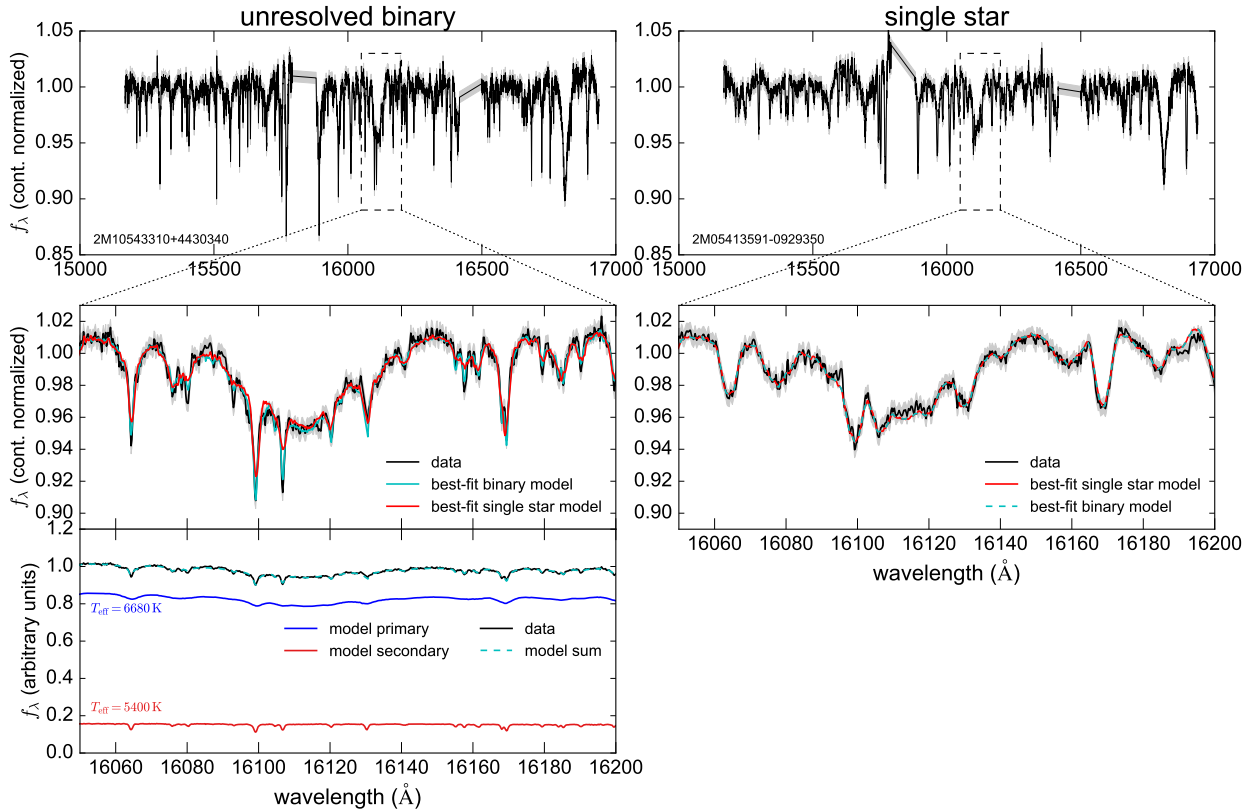


Figure 3.1: **Left:** Spectrum of an unresolved main-sequence binary with $q = m_2/m_1 \approx 0.7$ as observed by APOGEE. Top panel shows the full normalized spectrum. Middle panel shows the spectrum and best-fit binary and single-star models, zoomed-in on a narrow wavelength range enclosing a hydrogen Brackett line. The binary model fits the data significantly better than the single-star model. Bottom panel shows the two components of the best-fit binary model. The spectrum’s broad features are due primarily to the hotter star, which contributes $> 80\%$ of the total light but has no strong narrow lines; the shape of the sharp line profiles is primarily due to the cooler star. Our method makes it possible to identify many long-period binaries like this one, in which the velocity offset between the two stars is negligible. **Right:** Spectrum of a presumed single star with similar parameters to the primary in the system shown in the left panels. In this case, the best-fit binary and single-star models are identical.




Figure 3.2: Examples of single-star and binary model fits to binary systems with a negligible (top), intermediate (middle), and large (bottom) line-of-sight velocity offset between the two stars. All three systems have a mass ratio $q = m_2/m_1 \sim 0.7$, a primary star with $T_{\text{eff}} \sim 5400$ K and $\log g \sim 4.5$, and $[\text{Fe}/\text{H}] \sim 0$. Detecting binarity in systems with a large velocity offset ($\Delta v_{\text{los}} \gtrsim 15 \text{ km s}^{-1}$) is straightforward, because the two stars' lines become separated in velocity space. However, binarity can also be detected in many systems where the line-of-sight velocity offset is negligible, as in the top panel, because the two component stars have different temperatures and ionization states, so their combined spectrum cannot be well-fit by any single-star model.

4. *SB2 with unseen 3rd object*: This model fits two stars but allows their velocities to vary freely, without enforcing Equation 3.3. If it provides a significantly better fit than the SB2 model, the relative radial velocity shifts are inconsistent with being a simple Keplerian two-body system. The $7 + 2N$ free parameters are:

$$\vec{\ell}_{\text{SB2, unseen 3rd object}} = (T_{\text{eff}}, \log g, [\text{Fe}/\text{H}], [\text{Mg}/\text{Fe}], q, v_{\text{macro1}}, v_{\text{macro2}}, v_{\text{Helio1},i}, v_{\text{Helio2},i}). \quad (3.7)$$

5. *SB3*: The SB3 model fits three stars and imposes no restrictions on their velocities. The $9 + 3N$ free parameters are:

$$\vec{\ell}_{\text{SB3}} = (T_{\text{eff}}, \log g, [\text{Fe}/\text{H}], [\text{Mg}/\text{Fe}], q_2, q_3, v_{\text{macro1}}, v_{\text{macro2}}, v_{\text{macro3}}, v_{\text{Helio1},i}, v_{\text{Helio2},i}, v_{\text{Helio3},i}), \quad (3.8)$$

where $q_2 = m_2/m_1$ and $q_3 = m_3/m_1$.

We note that the SB2 models are in principle identical to the SB1 model (and the SB3 model to the SB2 models) in the limit where $q = 0$. We keep these models separate in practice because our model does not transition smoothly from the minimum possible q that can be modeled (corresponding to $T_{\text{eff}} = 4200$ K) to $q = 0$.

Fitting many visits simultaneously increases the number of labels to be fit, increasing the risk of the optimizer's convergence on a local minimum. For example, for a target with 10 visits, fitting the SB2 (SB3) model requires optimization of the likelihood in 19 (39) dimensions, which is computationally demanding. In tests with synthetic binary spectra, we find that convergence on the global best-fit is nearly always achieved *as long as the optimizer is initialized reasonably close to the global minimum*; i.e., with all velocities within $\pm \sim 20 \text{ km s}^{-1}$ of their true values at all epochs. We therefore first fit individual visit spectra one at a time to estimate the velocity of each component at each epoch, and then use the resulting best-fit labels to initialize the global optimizer during simultaneous fitting. Because the velocity offsets at each epoch are nearly uncorrelated with those in other epochs – i.e., changing $v_{\text{Helio},1,i}$ only shifts the spectrum predicted for the i th visit – the optimization remains tractable in many dimensions.

3.4 Results

We fit the spectra of 20,142 targets from APOGEE DR13 that ab-initio fitting with single-star models (Ting et al., in prep) found to (a) lie on the main sequence ($\log g > 4$), (b) fall within the region of label space where the synthetic spectral model is reliable ($4200 \text{ K} \leq T_{\text{eff}} \leq 7000 \text{ K}$ and $-1 \leq [\text{Fe}/\text{H}] \leq 0.5$), and (c) be acceptably fit, in a χ^2 sense, by synthetic spectral models. From this initial sample, we identify 2645 targets in which more than one star contributes significantly to the spectrum and an additional 663 targets with time-variable radial velocities but no detectable spectral contribution from the

secondary. Catalogs of targets classified as single stars, binaries, and triples are presented in Appendix 3.10.

Figure 3.1 illustrates how our model identifies systems that are likely binaries but show no significant radial velocity variability or split lines due to a velocity offset between the two components. Panels on the left show the spectrum of a target that can be significantly better fit by a binary model than a single-star model; those on the right show one that cannot.

We fit the full spectrum simultaneously, but we zoom-in on a small region to show the qualitative signatures of binarity. The spectrum in the left-hand panels contains features of both hot and cool stars: wide hydrogen lines and rotationally-broadened line profiles at the wing of all lines, and deeper, narrow line cores that do not show rotational broadening. No single-star model can achieve a good fit: the absorption lines in the best-fit single star model are too shallow, and some lines in the data spectrum are blended in the best-fit single star model or are missing altogether. On the other hand, the binary model can provide a good fit and reproduces the line profiles of the observed spectrum. The decomposition of the binary model spectrum in the bottom panel shows that the broad features are all due to the hot primary star, while the sharper features originate in the spectrum of the cooler secondary.

In the right panels of Figure 3.1, we show the spectrum of a typical single star with stellar parameters and abundances similar to the primary in the left panels; as expected, it is similar to the spectrum in the left panels with the sharper, narrow lines removed. In this case, the binary and single-star models converge on what is essentially the same spectrum, so there is no reason to prefer the binary model.

The binary spectrum in the left panels of Figure 3.1 illustrates why it is often possible to spectrally identify binaries even when one star is much brighter than the other: although the secondary star in the binary system contributes less than 20% of the total *light*, it contributes a large fraction of the total *absorption* because lines in hotter stars are often intrinsically weaker than those in cool stars. For many binaries containing a hot primary and cool secondary, the spectrum and binary model exhibit lines that are completely absent from the spectrum of the primary because the relevant species are ionized at its higher T_{eff} .

3.4.1 Effect of a velocity offset

Although a line-of-sight velocity difference between the primary and secondary star is in many cases not required to identify binaries with our model, a velocity offset makes the signatures of unresolved binarity more obvious and extends the range of detectable mass ratios. This is illustrated in Figure 3.2, which compares the spectra of three binary systems with similar stellar parameters, abundances, and mass ratios, but a range of velocity offsets between the primary and secondary component. The system shown in the top panel has a small line-of-sight velocity offset, similar to the system in the left panels of Figure 3.1. In this case, the effects of binarity are quite subtle, and binarity can likely only be detected with detailed spectral modeling. As the velocity offset increases (middle and bottom panels), binarity-induced changes to the spectrum become more obvious. In all three panels of Figure 3.2, we plot the APSTAR combined spectrum, not the spectra from individual visits.

However, the target in the bottom panel, which is the only target of the three for which we might expect a large velocity change between visits, was only visited once.

For APOGEE spectra with $R = 22,500$, one resolution element corresponds to a radial velocity difference of $\delta v \sim c/R \sim 13.5 \text{ km s}^{-1}$. The traditional method of identifying binaries as systems in which the cross-correlation function of an observed spectrum with a synthetic template exhibits two peaks can only reliably detect binaries in which the line-of-sight velocity offset is of order 1-3 resolution elements; such systems are usually referred to as ‘‘SB2’’ systems. For example, [Fernandez et al. \(2017\)](#) found that binaries could only be reliably detected in APOGEE spectra when the maximum line-of-sight velocity separation exceeded $\Delta v_{\text{los}} = 30 \text{ km s}^{-1}$.⁷ Figure 3.2 shows that even a small velocity offset can substantially strengthen the signatures of binarity. How much a velocity offset improves detectability for our method depends on the stellar parameters and abundances of the primary, because it is easier to detect velocity offsets in stars with many deep, narrow lines. For most stars with $T_{\text{eff}} \lesssim 6500 \text{ K}$, a velocity offset of $\Delta v_{\text{los}} \gtrsim (5 - 10) \text{ km s}^{-1}$ makes it possible to identify binaries from single-epoch spectra even when the mass ratio is close to $q = 1$; such systems are not otherwise detectable with our method (see Appendix 3.7.1).

3.4.2 Results for multi-epoch spectra with velocity variability

Examples of targets whose spectra are best-fit by SB2, SB1, SB2 with an unseen 3rd object, and SB3 models are shown in Figures 3.3, 3.5, 3.8, and 3.6.

Figure 3.3 shows a system that is best-fit by the SB2 model (i.e., case (iii) from Section 3.3.4) and exhibits spectra that change substantially from one epoch to the next. In the upper panel, we plot the combined spectrum and the best-fit binary and single-star models obtained by fitting it. Although the binary model is a better fit (and our initial fit to the combined spectrum did flag the system as a likely binary), the fit is not very good: some features in the combined spectrum cannot be accommodated by either the single-star or the binary model. In the lower panel, we show the spectra obtained in the three individual visits, which are coadded to produce the combined spectrum, and the binary model obtained by simultaneously fitting them. The fit to the individual visit spectra is good. The poor fit to the combined spectrum is a consequence of the fact that the components’ velocities change between visits, meaning that the combined spectrum is an unphysical superposition of different spectra.

The inset in Figure 3.3 shows the heliocentric velocities of the primary and secondary star at each visit for the best-fit SB2 model. The slope and intercept of the line on which these velocities fall can be used to calculate the dynamical mass ratio, q_{dyn} , and the center-of-mass velocity, γ . For binary systems in which the velocities of the two stars change significantly between visits, it is therefore possible to obtain a constraint on the mass ratio

⁷Other surveys find similar sensitivity to spectroscopic binaries; e.g., [Merle et al. \(2017\)](#) found that binaries could be detected down to $\Delta v_{\text{los}} = 15 \text{ km s}^{-1}$ in UVES spectra ($R = 47000$) from the Gaia-ESO survey, and [Matijevic et al. \(2010\)](#) found a minimum Δv_{los} for reliable detection of 50 km s^{-1} in the RAVE survey ($R = 7500$).




Figure 3.3: A binary system in which the stars’ velocities change between visits. Top panel shows a small portion of the *combined* spectrum (black), which is produced by coadding spectra from different visits, best-fit single-star model (red), and best-fit binary model (cyan). The binary model provides a better fit than the single-star model, but it cannot fully reproduce the combined spectrum. Bottom panel shows the individual visit and the best-fit SB2 model, which produces an excellent match to all the individual visit spectra. Inset shows heliocentric velocities of the primary and secondary star at each epoch; momentum conservation requires that these lie on a line with slope $-1/q_{\text{dyn}}$, where q_{dyn} is the dynamical mass ratio. The spectrally-inferred mass ratio, $q_{\text{spec}} = 0.93$, is in good agreement with the dynamical mass ratio, $q_{\text{dyn}} = 0.91$.

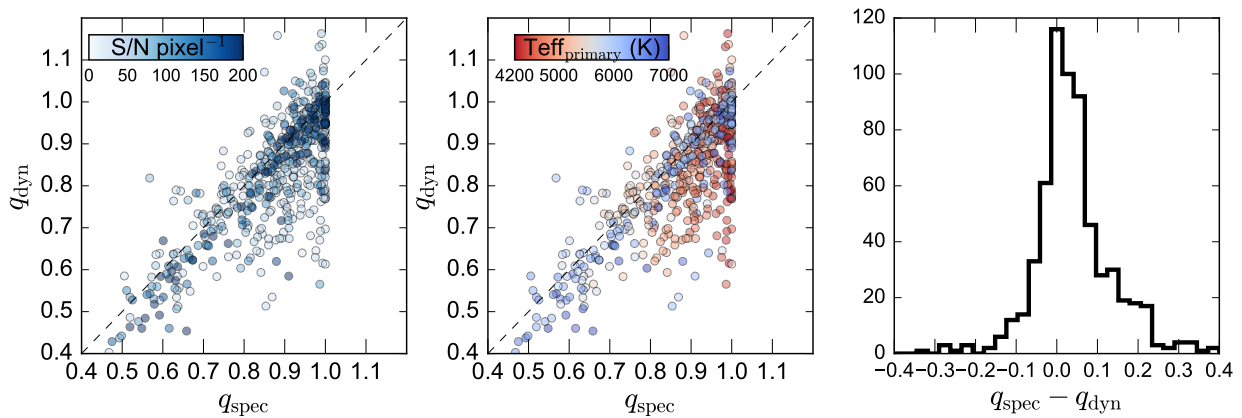


Figure 3.4: Comparison of spectroscopically- and dynamically-inferred mass ratios for “SB2” binary systems in which a dynamical mass ratio can be measured. q_{spec} is measured from the relative contribution of each star to the spectrum, and q_{dyn} , from the relative changes of the radial velocities of the primary and secondary across multiple epochs (see Figure 3.3). The designation of primary and secondary components is based on their relative contribution to the spectrum: q_{spec} is bounded by 1, but q_{dyn} is not. 623 systems have sufficiently short periods to allow measurement of q_{dyn} . Most systems for which the disagreement between q_{spec} and q_{dyn} is large have low S/N (left) and cool primaries (middle). Due to the spectral model’s minimum T_{eff} of 4200 K, low mass ratio systems can only be detected if the primary is hot, and mass ratios are less accurate for cooler systems. The median absolute difference between q_{dyn} and q_{spec} is 0.048.

that is independent of the spectral label q . Such constraints will of course not be reliable if the orbital configuration does not change significantly between visits: in this case, all measurements of $v_{\text{Helio},1}$ and $v_{\text{Helio},2}$ will be clustered around one point, and the slope of the line is ill-constrained. We also emphasize that linear momentum conservation requires that the slope of the line on which $v_{\text{Helio},2}$ and $v_{\text{Helio},1}$ fall must be *negative* for a true binary system. [Fernandez et al. \(2017\)](#) attempted to infer q_{dyn} also from systems in which the slope of this line is positive or zero (e.g. their Figure 6), but mass ratios inferred in this way have no physical interpretation and indicate either inaccurate radial velocity measurements or the presence of a third, unseen component.

In Figure 3.4, we compare the best-fit values of q_{spec} and q_{dyn} obtained for SB2 systems in which q_{dyn} can be reliably measured; we identify such systems as those in which the range of v_{Helio} spanned across visits is at least 10 km s^{-1} for both stars, corresponding to a velocity shift of slightly less than 1 resolution element. We color points by the median of S/N per pixel as reported in the allVisit catalog, where the median is over all visit spectra used in the fit.

The agreement between q_{spec} and q_{dyn} is in general quite good, with a median absolute difference between q_{spec} and q_{dyn} of $\text{med}(|q_{\text{spec}} - q_{\text{dyn}}|) = 0.048$ and a corresponding middle 68% range of $(0.012 - 0.14)$. The agreement is on average better for targets whose spectra have higher S/N; most systems with significantly different q_{spec} and q_{dyn} have $\text{S/N} \lesssim 50$. Particularly at lower mass ratios, q_{dyn} is on average slightly lower than q_{spec} ; i.e., assuming q_{dyn} is usually more accurate than q_{spec} , the latter is biased to slightly higher q . This can be understood as a consequence of the minimum T_{eff} of our spectral model, which sets an effective minimum q_{spec} . If a cool primary has a companion with T_{eff} cooler than 4200 K that cannot be fully accommodated by the spectral model, a better fit can often still be achieved by a binary model with $T_{\text{eff}} = 4200 \text{ K}$ and a too-high q_{spec} than a single star model which ignores the secondary entirely. This in part explains the substantial number of cool systems with q_{spec} near 1 and lower q_{dyn} , though we note that most cool systems also have lower S/N.

If the secondary is very faint compared to the primary, its contribution to the spectrum may be completely undetectable, in which case binary and single-star models will converge to the same model spectrum as long as the velocities are allowed to vary between visits. Such systems can be distinguished from isolated single stars by the fact that the ‘‘SB1’’ model provides a better fit than the single-star model, which requires a target’s velocity to be the same at all visits. Figure 3.5 shows an example of such a system. Our model makes it possible to set an upper limit on the mass ratio, under the assumption that the companion is a main sequence star: in this case, the SB2 model would provide a better fit than the SB1 model if the secondary had $T_{\text{eff}} \gtrsim 4200 \text{ K}$. This limit is likely conservative in practice for main-sequence secondaries, as discussed above. However, it will not apply for binaries in which the companion is a stellar remnant.

We note that most SB1s and some close SB2s can be qualitatively identified as unlikely to be single based on the scatter across visits in the radial velocities measured by the APSTAR pipeline (e.g. [Badenes et al. 2017](#)). However, we find a nontrivial number of SB2 systems (~ 100 systems out of the ~ 20000 targets studied in this work) that show clearly time-variable




Figure 3.5: Visit spectra and best-fit models for an SB1 system. The SB1 model contains only a single star contributing to the spectrum, but its radial velocity can vary across visits. The SB2 model includes the possibility of a second star contributing to the spectrum. In this case, the best-fit SB1 and SB2 models are identical, indicating that there is no detectable contribution to the spectrum from the secondary. However, radial velocity variability of the primary clearly indicates that a companion is present. Assuming that the companion is a main-sequence star, an upper-limit of $q \lesssim 0.45$ can be derived; if q were larger, the secondary would contribute detectably to the spectrum, and the SB2 model would provide a better fit.

spectra, with changes in the velocities of both components larger than 30 km s^{-1} , for which the APSTAR-derived v_{Helio} measurements change at the $< 1 \text{ km s}^{-1}$ level. This indicates that APOGEE radial velocity measurements are likely problematic for these systems, and studies that flag short-period binaries based on velocity variability will miss some SB2 systems.

Figure 3.6 shows an example of a spectrum classified as a triple. The SB2 model (cyan) clearly cannot provide a good fit to the observed spectra, which simply have too many lines; on the other hand, the triple model is a good fit to all visits. The inset shows the velocities of the three components at each epoch; note that these are all allowed to vary freely and are not restricted to follow any equivalent of Equation 3.3. One component, the spectral primary, has effectively constant velocity (within $\pm 0.5 \text{ km s}^{-1}$) across all visits. On the other hand, the velocities of the secondary and tertiary components vary a great deal between visits and fall on a line with negative slope, just as in the case of close binaries (Figure 3.3). The most straightforward explanation for these kinematics is that the system is a hierarchical triple (e.g. Ford et al. 2000; Toonen et al. 2016) consisting of a close binary orbiting a third system with a period much longer than that of the close binary, so that the velocity of the primary and the center-of-mass velocity of the close binary do not change significantly over the temporal baseline between visits (which is ~ 54 days for this target). This type of hierarchical orbital is illustrated schematically in Figure 3.6. Consistent with this interpretation, the spectrally inferred mass ratio between the two components of the close binary is similar to the dynamical mass ratio inferred from the slope of the line on which their velocities fall.

We find 114 triple systems, most of which have the same qualitative velocity configuration as the system in Figure 3.6: they contain one component with effectively constant velocity over all visits and two components with variable velocities that fall on a line as expected by a close binary. This is not surprising, as hierarchical configurations are the natural stable end state of the dynamical evolution of (otherwise chaotic) triple systems (Naoz et al. 2013). We also find systems in which the velocity of the third (long-period) component is not constant but changes approximately linearly with time; this is expected if the system’s outer period is long compared to the observation baseline but not so long that no change can be observed. In such cases, the heliocentric velocities of the other two components do not fall on a straight line but exhibit some intrinsic scatter; this scatter can be reduced if a constant multiple of the linear trend of the lone star is subtracted from the velocities of the other two stars.

Such systems are almost certainly gravitationally bound triples, since the velocities of all three components are correlated. However, for triples in which the velocity of one component is consistent with being constant over the time baseline spanned by observations, there is no guarantee that the three stars are actually gravitationally bound: the observed velocities could also be explained by a chance alignment between a close binary system and a background or foreground star. Whether such chance alignments constitute a substantial fraction of the targets we identify as hierarchical binaries can be diagnosed by comparing the center-of-mass velocity of the close binaries to the velocity of the third component. For gravitationally bound triples, these should be reasonably similar, with offsets of order the orbital velocity of the long-period component. The typical offsets should be larger (at mini-

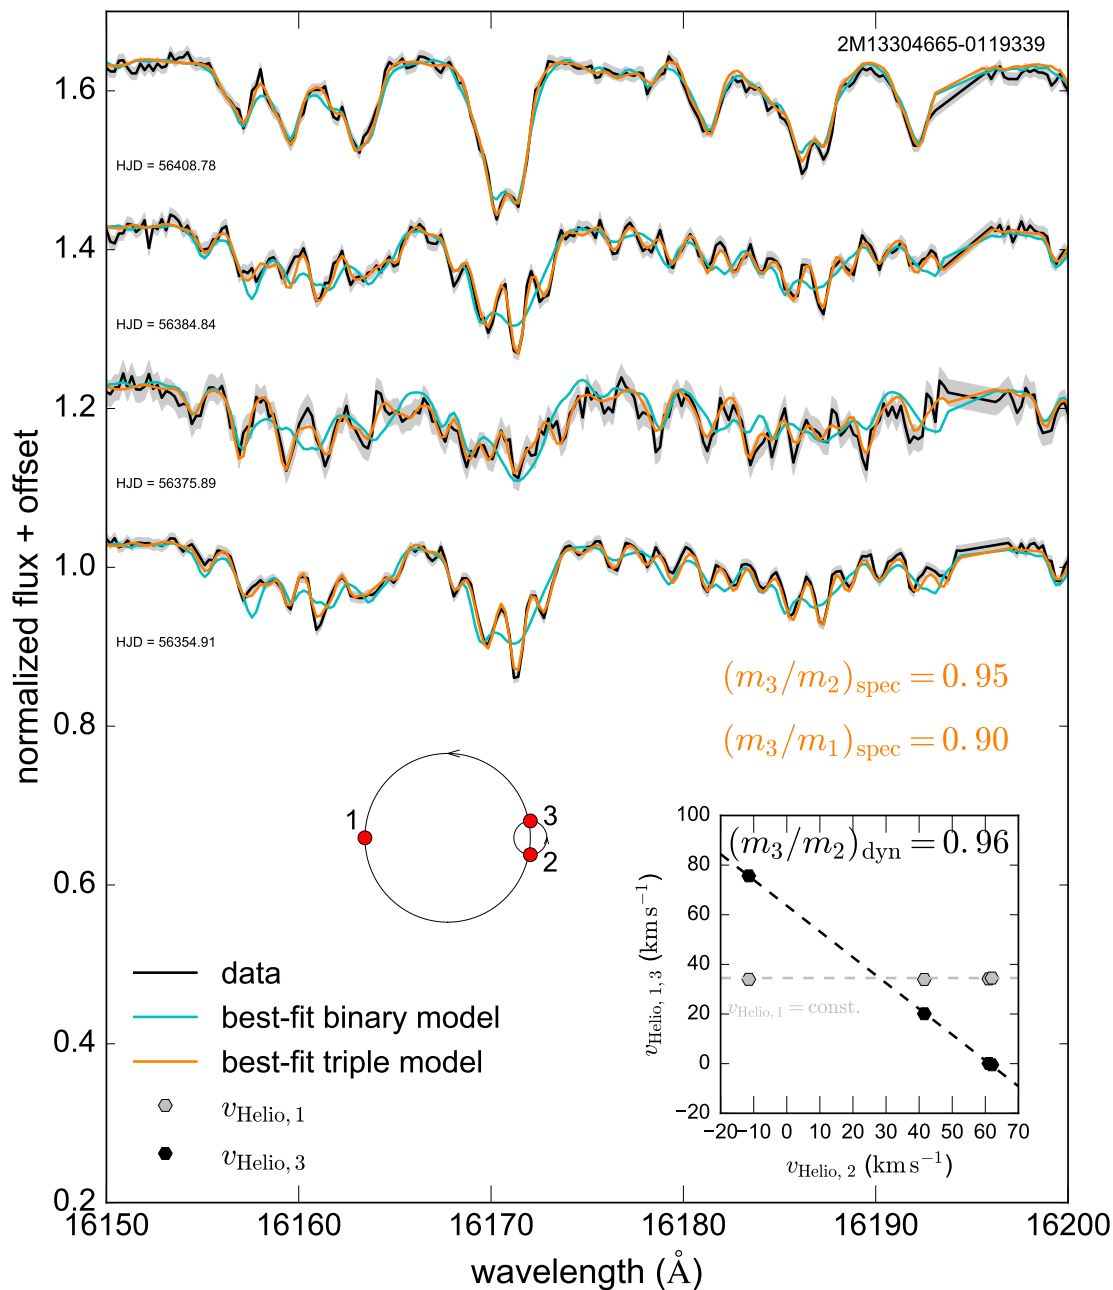


Figure 3.6: Visit spectra of a target identified as a triple (SB3) system. The three components have different line-of-sight velocities, so many lines can be seen in triple, and an SB2 model cannot provide a good fit. Inset shows the line-of-sight velocities of each component at each epoch. The heliocentric velocity of the primary is consistent with being constant at $v_{\text{Helio},1} \approx 34.5$ km s⁻¹, so no dynamical mass ratio can be estimated for m_2/m_1 or m_3/m_1 . However, $v_{\text{Helio},2}$ and $v_{\text{Helio},3}$ fall on a line with an implied mass ratio m_3/m_2 consistent with the spectrally inferred one. This implies that the system is a hierarchical triple, as is illustrated in the orbit schematic.

mum, of order 30 km s^{-1} , the velocity dispersion of the Milky Way’s stellar disk) for chance alignments.

We investigate this explicitly in Figure 3.7. Here we only plot systems that are consistent with the velocity of the of the long-period component being fixed over all epochs; we identify such cases as systems in which the change in the velocity of the long-period component across epochs is less than 2 km s^{-1} when all velocities are allowed to vary freely. Consistent with the expectation for bound triples, the system velocity of the close binary is in most cases within 10 km s^{-1} of that of the third component. There are 5 systems in which the offset is larger, but due to the relatively short observational baselines, we find that none of these velocity offsets are large enough to rule out the possibility that all three stars are gravitationally bound. We discuss the possibility of contamination due to chance-alignments of stars further in Section 3.4.5.

Along with SB1s, SB2s, and SB3s, we also identify a class of systems in which the presence of a third component can be deduced from radial velocity measurements, but only two star contribute significantly to the spectrum. Figure 3.8 shows an example of such a target. The standard SB2 model, which enforces Equation 3.3 with $q_{\text{dyn}} \geq 0.2$, cannot satisfactorily fit the spectrum. However, the “SB2 with unseen third component” model, which allows the velocities of both components to vary freely, provides a good fit, converging on a solution in which the velocity of one component is consistent with being fixed across epochs while that of the other varies.

As illustrated in the orbital schematic in Figure 3.8, such a radial velocity pattern can be explained straightforwardly if the system is a hierarchical triple in which the close binary is an SB1; i.e., one component of the close binary does not contribute to the spectrum, either because its mass is low or because it is a compact remnant. No dynamical mass ratio can be inferred for these systems, because the acceleration of the variable velocity component is due primarily to the unseen component. We identify 108 SB2 systems in which the presence of a third component can be inferred dynamically; the majority of these systems have velocity configurations similar to Figure 3.8, with one component’s velocity essentially constant.

3.4.3 Color-magnitude diagram

A straightforward diagnostic to verify that the targets we spectroscopically identify as binaries are primarily true binaries, as opposed to single stars whose spectra contain unusual features that are not well accounted for in the spectral model, is to examine their distribution in a color-magnitude diagram (CMD). True binaries are expected to lie above the single-star main sequence of a CMD (Hurley & Tout 1998; Li et al. 2013): binaries with $q \sim 1$ will have the same color as would either single star but will be twice as luminous, while binaries with $0.5 \lesssim q \lesssim 0.9$ will be both brighter and redder than a single star with the parameters of the primary.

Accurate measurements of absolute magnitude (and hence distance) are required to construct the CMD. To identify stars in our sample with accurate distance measurements, we cross-matched it with the Tycho-Gaia astrometric solution (TGAS) catalog (Michalik et al.

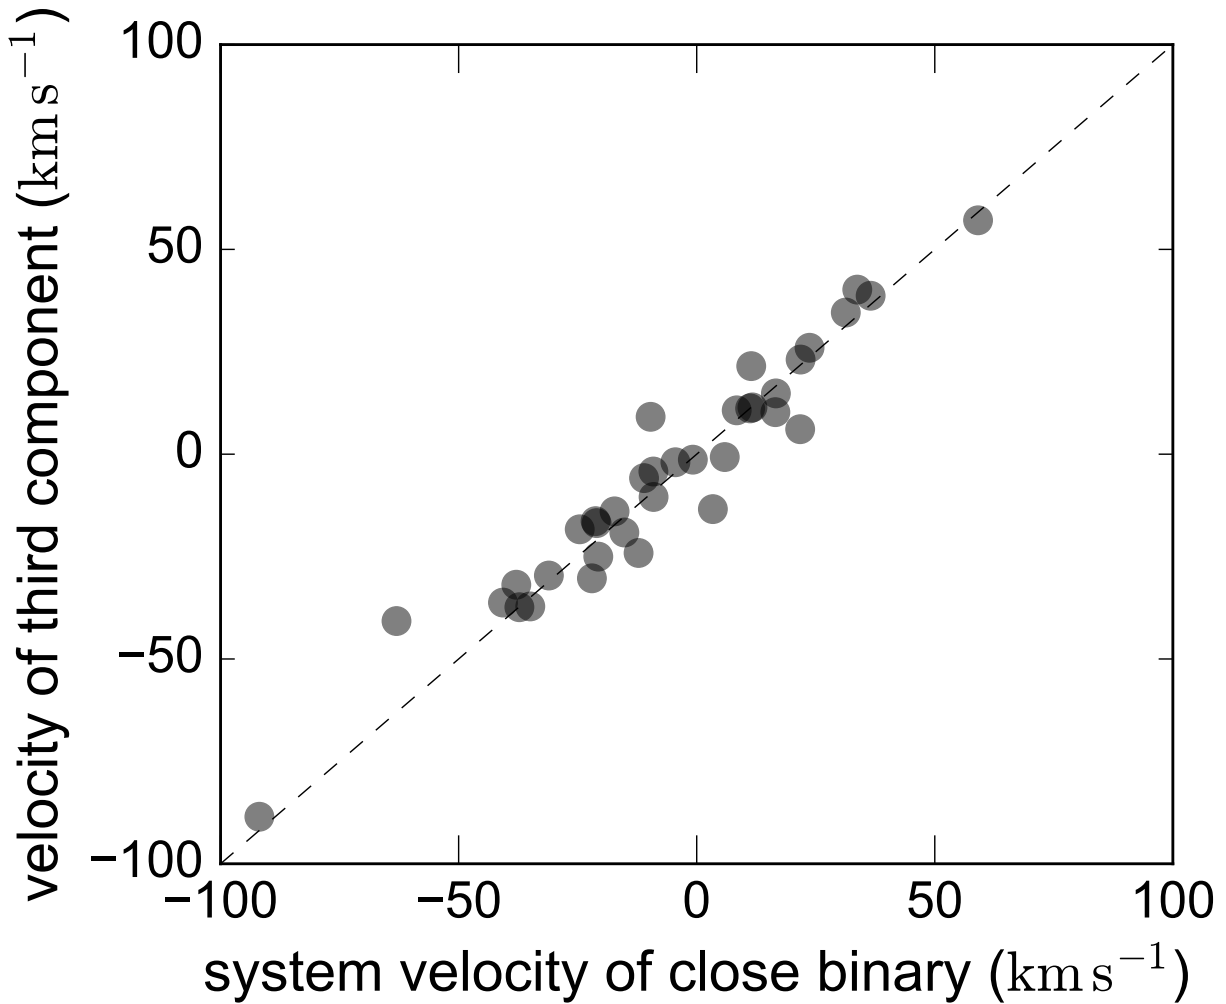


Figure 3.7: Line-of-sight velocities for hierarchical triples containing a close binary and third component with a much larger separation (see schematic in Figure 3.6). The tight correlation between the center-of-mass velocity of the close sub-binary and the velocity of the long-period component indicates that most systems are bona-fide gravitationally bound triples, not chance alignments between a close binary and a background or foreground star.

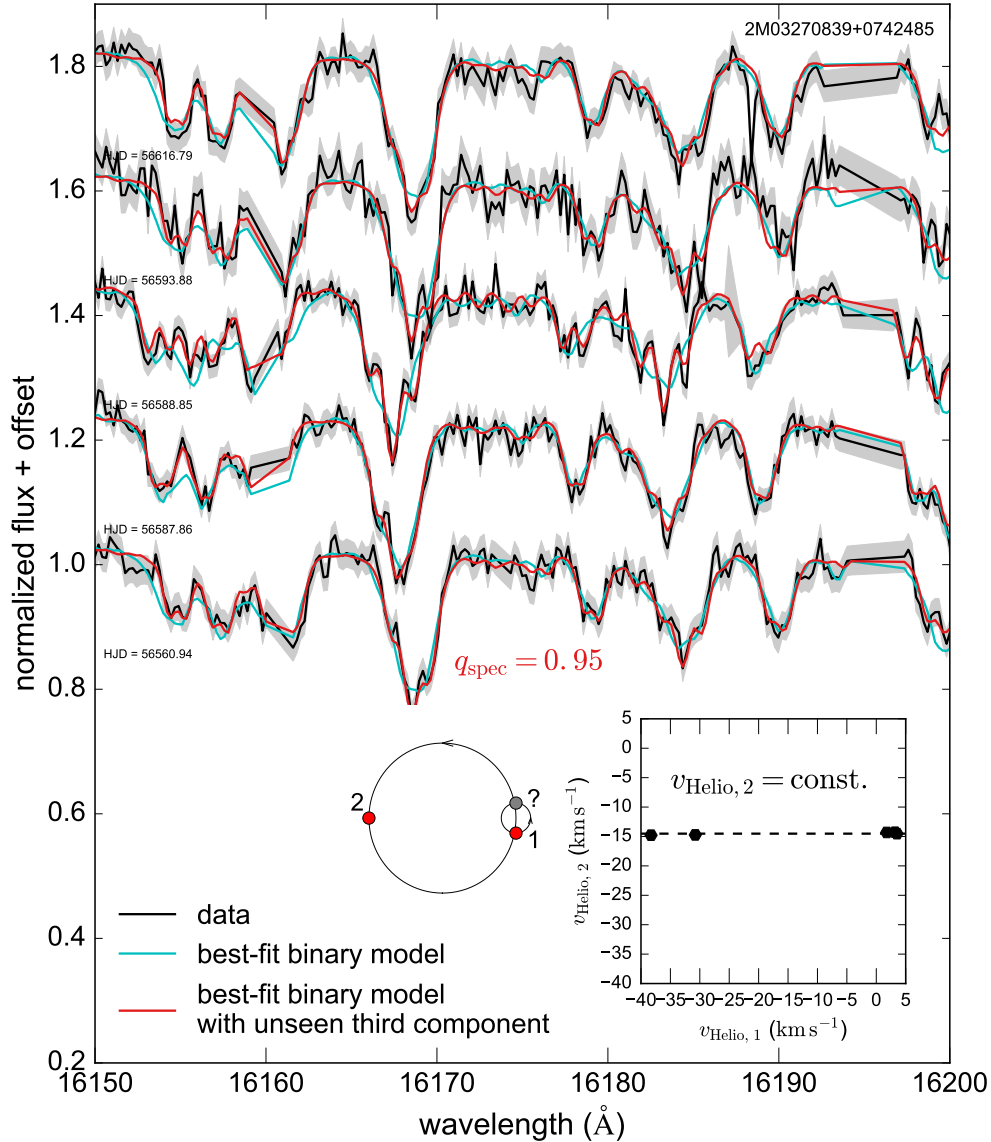


Figure 3.8: Visit spectra of a triple system in which the third component does not contribute significantly to the spectrum but can be detected gravitationally. Cyan line shows best-fit SB2 model with the restriction that $v_{\text{Helio},1}$ and $v_{\text{Helio},2}$ fall on a line with negative slope (Equation 3.3). Red line shows the best-fit binary model in which the velocities of the primary and secondary are allowed to vary freely. Inset shows the line-of-sight velocities corresponding to the red model. The velocity of the secondary is consistent with being constant at $v_{\text{Helio},2} \approx -14.5 \text{ km s}^{-1}$, while that of the primary varies substantially. This implies that the system is a hierarchical triple in which one component of the close binary does not contribute to the spectrum (i.e., it is a stellar remnant or faint M-dwarf); this is shown schematically in the orbital diagram.




Figure 3.9: Left: Color-magnitude diagram of all stars in our sample with parallax errors of 10% or less (gray circles). Black and red symbols represent spectroscopically identified binaries, with maximum line-of-sight velocity offsets between the two components greater (short-period) and less than (long-period) 30 km s^{-1} . Binaries with mass ratios $q < 0.6$ are marked with hexagons; those with $q > 0.6$, with stars. Triples are marked with blue points. Most spectroscopically-identified binaries lie above the main sequence. Traditional SB2 identification methods can only identify close binaries with large velocity offsets (black symbols); our method can identify many more long-period binaries with negligible velocity offsets (red symbols). **Right:** Schematic effect of unresolved binarity on the CMD. Black line shows a MIST isochrone for single stars with a single age and metallicity; colored loci show where binaries with different mass ratios and primary T_{eff} fall when they are spatially unresolved.

2015) using the `gaia_tools.xmatch` routine written by Jo Bovy. This revealed 1925 stars in our sample with parallax errors of 10% or better,⁸ 217 of which were spectroscopically identified as multiple-star systems in which at least two components contribute detectably to the spectrum.

We plot the CMD for these objects, based on 2MASS photometry, in the left panel of Figure 3.9. As expected, the majority of spectroscopically-identified binaries are scattered above the main sequence. We stress that our model for identifying binaries operates exclusively on normalized spectra and does not rely whatsoever on photometry; the fact that nearly all of the spectroscopically identified binaries populate the expected locus of the CMD above the main sequence is therefore a robust confirmation that our model is finding real binaries.

In the right panel of Figure 3.9, we show schematically how the presence of an unresolved companion is expected to change a star’s position on the CMD. For a single stellar population, binaries with mass ratios $0.6 \lesssim q \lesssim 1$ form a coherent second sequence ~ 0.75 magnitudes above the main sequence; i.e., it is *not* the case over this range of mass ratios that binaries with higher q fall farther above the main sequence. This occurs because as q is increased from 0.6 to 1 and T_{eff} of the secondary increases, unresolved binaries move blueward parallel to the main sequence. On the other hand, in binaries with $q \lesssim 0.4$, the secondary contribute so little light that the change in the unresolved system’s location on the CMD is negligible.

The lowest mass-ratio to which we are sensitive is $q \sim 0.4$, so the majority the binaries we identify should scatter above the main sequence for their respective isochrone. In the left panel of Figure 3.9, we plot separately binaries with $q \leq 0.6$ (which are only detectable around primaries with $T_{\text{eff}} \gtrsim 5800$ K; see Appendix 3.7.1) and those with higher mass ratios. As expected, the lower mass ratio systems are on average below the higher mass ratio systems on the CMD. With one exception,⁹ systems identified as triples (SB3) fall above the binary main sequence. We do not mark SB1s in Figure 3.9; their distribution on the CMD is similar to that of single stars, likely because most have low mass ratios.

The sample of stars for which accurate parallaxes are available spans a wide range of metallicities and ages, so significant intrinsic scatter is expected in the distribution of both single stars and binaries on the CMD. We note that there are some stars that our model does not identify as binaries but which still scatter well above the main sequence. We suspect that most of these systems are binaries with $q \sim 1$ and small velocity offsets; these are not detectable in our current framework.

We also divide suspected binaries into subsamples with large and small velocity offsets, corresponding approximately to systems which could and could not be detected with tradi-

⁸This corresponds to a magnitude error of ± 0.22 mag, plus typical 2MASS photometric errors of ± 0.03 mag. We do not attempt to correct for extinction or reddening, which is expected to be modest in the near-infrared at the distances of the stars with accurate parallaxes (which have a median distance of 200 pc).

⁹We have investigated the spectra of this target (2M07212735+2342096) in detail and find it to be an unambiguous triple, with clear changes in spectral morphology between visits. Why it falls below the main sequence is not clear; one possibility is that marginally resolved multiplicity led to an overestimate of its parallax.

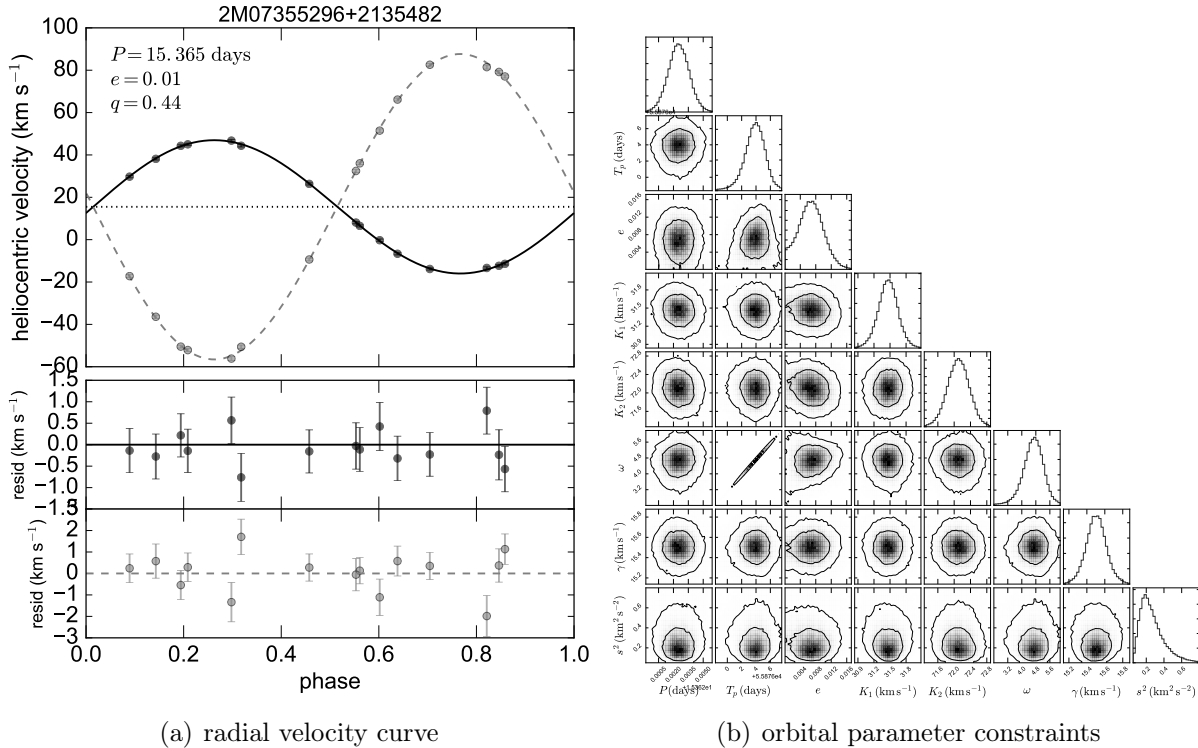


Figure 3.10: Left: Orbit fit to an SB2 systems with 15 epochs. In this relatively low-mass ratio system, the secondary contributes only $\sim 3\%$ of the light, but we can still determine its velocity to $\pm \sim 1 \text{ km s}^{-1}$. Black and gray lines show the heliocentric velocity of the primary and secondary star. *Right:* 68 and 95% marginalized probability regions for the system’s orbital parameters. Because the eccentricity is very nearly 0, T_p and ω are highly degenerate and are individually not well constrained, reflecting the fact that these quantities are meaningless for a circular orbit. We derive similar orbital solutions, shown in Table 3.1, for 64 systems.

tional CCF-based binary detection methods (Fernandez et al. 2017). Only $\sim 30\%$ of systems have velocity offsets that are large enough to be detected with traditional methods. This highlights one of the primary advantages of the method introduced in this work: it is sensitive to a substantially larger fraction of the binary population than methods based on radial velocity separation or variability alone.

3.4.4 Deriving Orbital Parameters

We derive full orbital solutions for 64 binary systems which have sufficient visits and phase coverage to constrain the orbit. Our criteria for determining whether the available velocity data are sufficient to constrain a system’s orbit are discussed in Appendix 3.9. We only attempt to derive orbital solutions for systems in which at least two stars contribute to the spectrum; orbital solutions for SB1 systems in APOGEE can be found in Troup et al.

(2016).

Velocities for both components are returned as labels for the best-fit spectral model. An initial estimate of the radial velocity uncertainty at each visit is obtained through bootstrapping: Gaussian noise proportional to the best-fit model residual at each pixel is added to the data spectrum and the fit is repeated; the uncertainty on each radial velocity is taken to be the standard deviation of the best-fit velocity at each epoch when this procedure is repeated many times.

Fitting a Keplerian orbit amounts to simultaneously maximizing the likelihood of the radial velocity curves of the primary and secondary, where the model radial velocity for a given set of orbital parameters is obtained by solving the two-body problem (Murray & Correia 2010). We use a custom Python implementation of the adaptive simulated annealing algorithm (Iglesias-Marzoa et al. 2015) to obtain an initial maximum-likelihood estimate of the best-fit orbital parameters and then sample the parameter space in the vicinity of the maximum likelihood with `emcee` (Foreman-Mackey et al. 2013a) to estimate parameter uncertainties. We use non-informative, flat priors throughout. In addition to the 7 standard Keplerian orbital parameters,¹⁰ we fit a “jitter” term, s^2 , to allow for the possibility of intrinsic scatter in the radial velocities due to e.g. stellar pulsation or underestimated radial velocity uncertainties (see e.g. Baluev 2009; Price-Whelan et al. 2017). The effective total uncertainties in the radial velocities used in fitting are then $\sigma_{\text{tot},i}^2 = \sigma_i^2 + s^2$, where σ_i are the radial velocity uncertainties at each epoch found from bootstrapping. Explicitly, the log-likelihood function is

$$\ln \mathcal{L} = -\frac{1}{2} \sum_i^N \left\{ \frac{[v_r(t_i, \vec{\theta}_1) - v_{\text{Helio1},i}]^2}{\sigma_{1,i}^2 + s^2} + \ln[2\pi(\sigma_{1,i}^2 + s^2)] \right. \\ \left. + \frac{[v_r(t_i, \vec{\theta}_2) - v_{\text{Helio2},i}]^2}{\sigma_{2,i}^2 + s^2} + \ln[2\pi(\sigma_{2,i}^2 + s^2)] \right\}. \quad (3.9)$$

Here the sum is over N epochs at times t_i , $v_r(t, \vec{\theta})$ represents the predicted radial velocity at time t for a system with orbital parameters $\vec{\theta}$ (Murray & Correia 2010), and $\vec{\theta}_1 = (P, T_p, e, \omega, K_1, \gamma)$ and $\vec{\theta}_2 = (P, T_p, e, \omega + \pi, K_2, \gamma)$ are the orbital parameters for each component. For most systems, the best-fit jitter is small ($s^2 \sim 0.1 \text{ km}^2 \text{ s}^{-2}$), indicating that our estimates of σ_i are reasonably accurate. However, for stars with large v_{macro} , which indicates significant rotation, jitter is sometimes of order 1 km s^{-1} . This suggests that our velocity measurements are less accurate for rapidly rotating stars, which is also supported by our experiments with semi-empirical binary spectra (Appendix 3.7.1).

For some systems with mass ratios near 1, we found that it was initially impossible to obtain a good fit to the measured radial velocities because the velocity assignments of the two stars were switched in the fits for one or more visits. We attempted to fit these systems by allowing the fitting algorithm to switch the assigned velocities for individual visits if doing

¹⁰These include the period, P , periastron time, T_p , eccentricity, e , argument of periastron, ω , center-of-mass velocity, γ , and the velocity semi-amplitudes, K_1 and K_2 .

so would improve the fit. In most cases, this solved the problem. A few systems ($\sim 5\%$ of those with sufficient coverage) remained with radial velocity curves that could not be well-fit by a Keplerian orbit, even with the possibility of switching the assigned velocities; these systems may have poorly measured radial velocities or contain an unseen component.

Figure 3.10(a) shows an example orbital solution for a system with typical phase coverage, radial velocity errors, and number of epochs. This system has the lowest mass ratio, $q = 0.44$, of the systems for which we derive orbital solutions. Because of the system’s low mass ratio, the secondary contributes only a small fraction ($\sim 3\%$) of the total light in the spectrum; it is not obvious from visual inspection that more than one star contributes to the spectrum. However, the secondary is detected unambiguously by our model, and the fact that the primary and secondary velocities all fall on a Keplerian orbit confirms the validity of the detection.

Marginalized probabilities for the orbital parameters of this system are shown in Figure 3.10(b). Most orbital parameters are well-constrained for this system, without strong parameter covariances. However, the periastron time, T_p , and argument of periastron, ω , are highly degenerate, because the orbit is nearly circular (the eccentricity, e , is consistent with 0); the orbit has no well-defined periastron, and hence T_p and ω are undefined. All systems with low eccentricities therefore have large uncertainties in ω and in T_p , even when the meaningful parameters of the orbit are well-constrained.

In Table 3.1, we provide best-fit orbital parameters and marginalized uncertainties for 64 systems for which an orbital solution could be obtained. Most of these systems are ordinary double-lined binaries (SB2s), similar to the system shown in Figure 3.3. However, we also include solutions for several close SB2s within hierarchical triples (similar to Figure 3.6), as well as SB1s within hierarchical triples with hidden third components (similar to Figure 3.8). These are fit very similarly to pure SB2 systems, with the only difference being that v_{Helio} measurements for individual stars at each visit are obtained from the “SB3” and “SB2 with an unseen 3rd object” models. We only attempt to fit such systems if they are consistent with the third component having constant velocity over the observed baseline. For “SB2 with an unseen 3rd object” systems, the fit is to a single radial velocity curve, so K_2 is not measurable. The rms velocity residual for all orbital solutions ranges from 0.04 and 1 km s^{-1} . The systems with the largest velocity errors have a) lower average S/N and b) higher T_{eff} and v_{macro} , both of which make it more difficult to accurately measure radial velocities.

The statistics U_N and V_N in the last column of Table 3.1 quantify the uniformity of coverage in orbital phase and velocity by the measured radial velocity data. We calculate these statistics following Troup et al. (2016) (their Equations 22 and 23); both U_N and V_N are bounded between 0 and 1, with values near 1 corresponding to uniform phase and velocity coverage. Troup et al. (2016) estimated that orbital parameters are unreliable for SB1s if $U_N V_N < 0.5$; of course, the probability of recovering the correct orbit also depends on the number of radial velocity measurements and their uncertainties. In Appendix 3.9, we carry out tests with synthetic radial velocity data to determine the number of epochs and phase + velocity coverage required for reliable orbit recovery of SB2s with radial velocity data similar to that obtained for real binaries.

2MASS ID	N_{epochs}	P [days]	T_p [BMJD]	e	ω [radians]	K_1 [km s ⁻¹]	K_2 [km s ⁻¹]	γ [km s ⁻¹]	$U_N V_N$
06212323+1701485	8	42.24 $\pm_{-0.17}^{+0.19}$	56657.15 $\pm_{-0.23}^{+0.21}$	0.3166 $\pm_{-0.0092}^{+0.0091}$	0.684 $\pm_{-0.034}^{+0.029}$	35.27 $\pm_{-0.29}^{+0.31}$	41.40 $\pm_{-0.31}^{+0.32}$	1.37 $\pm_{-0.18}^{+0.19}$	0.79
08544465+1130053	21	39.23855 $\pm_{-0.00093}^{+0.00097}$	55933.6200 $\pm_{-0.016}^{+0.016}$	0.6932 $\pm_{-0.0014}^{+0.0014}$	1.0931 $\pm_{-0.0035}^{+0.0035}$	59.79 $\pm_{-0.19}^{+0.20}$	62.93 $\pm_{-0.21}^{+0.20}$	-6.818 $\pm_{-0.069}^{+0.069}$	0.92
04030722+5150045	9	69.973 $\pm_{-0.093}^{+0.076}$	55906.20 $\pm_{-0.33}^{+0.38}$	0.569 $\pm_{-0.036}^{+0.045}$	4.026 $\pm_{-0.026}^{+0.024}$	31.5 $\pm_{-1.9}^{+2.9}$	33.4 $\pm_{-2.1}^{+3.1}$	-7.176 $\pm_{-0.091}^{+0.089}$	0.76
21313924+1307507	41	1.5567964 $\pm_{-0.0000015}^{+0.0000016}$	55731.18 $\pm_{-0.19}^{+1.10}$	0.0016 $\pm_{-0.0012}^{+0.0016}$	1.03 $\pm_{-0.76}^{+4.30}$	59.51 $\pm_{-0.12}^{+0.12}$	71.84 $\pm_{-0.13}^{+0.13}$	-52.199 $\pm_{-0.061}^{+0.069}$	0.92
18470667-0226077 ^a	32	7.52676 $\pm_{-0.00014}^{+0.00015}$	55823.81 $\pm_{-0.37}^{+0.39}$	0.0136 $\pm_{-0.0050}^{+0.0046}$	5.38 $\pm_{-0.30}^{+0.32}$	45.86 $\pm_{-0.26}^{+0.27}$	55.85 $\pm_{-0.26}^{+0.26}$	15.99 $\pm_{-0.15}^{+0.15}$	0.93
07355296+2135482	15	15.3645 $\pm_{-0.0011}^{+0.0010}$	55879.8 $\pm_{-1.5}^{+1.5}$	0.0065 $\pm_{-0.0034}^{+0.0033}$	4.62 $\pm_{-0.62}^{+0.60}$	31.47 $\pm_{-0.18}^{+0.17}$	72.11 $\pm_{-0.27}^{+0.27}$	15.50 $\pm_{-0.10}^{+0.11}$	0.87
15010903+3702218	7	17.5079 $\pm_{-0.0011}^{+0.0011}$	56090.8540 $\pm_{-0.045}^{+0.046}$	0.2996 $\pm_{-0.0013}^{+0.0014}$	2.5573 $\pm_{-0.0088}^{+0.0092}$	41.962 $\pm_{-0.092}^{+0.100}$	57.13 $\pm_{-0.14}^{+0.14}$	-47.938 $\pm_{-0.036}^{+0.040}$	0.61
08541894+1239291	22	1.30467925 $\pm_{-0.0000089}^{+0.0000087}$	55904.366 $\pm_{-0.052}^{+0.054}$	0.042 $\pm_{-0.010}^{+0.010}$	2.96 $\pm_{-0.24}^{+0.25}$	130.2 $\pm_{-1.7}^{+1.7}$	130.2 $\pm_{-1.8}^{+1.8}$	7.87 $\pm_{-0.80}^{+0.86}$	0.91
19303146+4210508 ^b	24	5.55412 $\pm_{-0.00014}^{+0.00013}$	56444.01 $\pm_{-0.22}^{+0.25}$	0.0120 $\pm_{-0.0030}^{+0.0029}$	6.05 $\pm_{-0.24}^{+0.28}$	43.20 $\pm_{-0.13}^{+0.13}$	–	-58.76 $\pm_{-0.100}^{+0.092}$	–
08464223+1205302	17	15.0232 $\pm_{-0.0024}^{+0.0023}$	56654.1460 $\pm_{-0.023}^{+0.023}$	0.2980 $\pm_{-0.0035}^{+0.0034}$	0.290 $\pm_{-0.011}^{+0.010}$	24.76 $\pm_{-0.12}^{+0.13}$	26.54 $\pm_{-0.14}^{+0.14}$	-6.493 $\pm_{-0.055}^{+0.058}$	0.93
...

Table 3.1: Orbital solutions for double-line spectroscopic binaries. We report the median and middle 68% of the marginalized posterior samples for each parameter. U_N and V_N quantify the phase and velocity coverage of the observations (see Appendix 3.9); systems with $U_N V_N \lesssim 0.5 - 0.6$ are susceptible to erroneous bad fits. ^aSystem is an SB2 within a hierarchical triple (10 systems). ^bSystem is an SB1 within a hierarchical triple (3 systems). This table is available in its entirety (with orbital solutions for 64 systems) in machine-readable form.




Figure 3.11: Top: Distribution of periods, semi-major axes, and masses for the 64 double-lined binary systems for which we derive an orbital solution. Due to APOGEE’s relatively rapid cadence (most targets have maximum baselines of a few months), these systems are heavily biased toward short periods. **Bottom:** Period-eccentricity distribution. Most systems with $P \lesssim 10$ days have low eccentricity due to tidal circularization.

In the top panel of Figure 3.11, we plot constraints on the semimajor axes and component masses derived from the orbital parameters of all systems for which we present an orbital solution.¹¹ Our sample contains systems with periods ranging from 0.6 days (short enough that the two stars are nearly touching, with $a \sin i \sim 0.8 R_\odot$) to ~ 600 days. Dynamical constraints on the absolute masses of stars in individual binaries are weak due to the degeneracy with $\sin i$, but the highest lower limit on the mass of an individual component is $\sim 1.5 M_\odot$. This corresponds to $T_{\text{eff}} \gtrsim 6600$ K for a solar-metallicity star on our adopted MIST isochrones; reassuringly, none of our dynamical mass constraints imply $T_{\text{eff}} > 7000$ K, which is the upper limit adopted for our spectral model.

We plot the periods and eccentricities of binary systems in the bottom panel of Figure 3.11. Most of the short period ($P \lesssim 10$ days) systems have nearly circular orbits, likely due tidal dissipation processes (e.g. Koch & Hrivnak 1981). However, a few systems with short periods do have eccentricity constraints that are inconsistent with zero. Some of these systems are short-period binaries within a hierarchical triple; such systems are known to be susceptible to eccentricity boosts via three-body interactions (Kozai 1962).¹²

3.4.5 Are binaries gravitationally bound?

Our method finds binaries and triples by identifying targets in which more than one star falls within a single APOGEE fiber and contributes to the observed spectrum. In all cases where the velocities of the components of a suspected binary or triple system are not observed to vary in a correlated way, there is no guarantee that all the components are gravitationally bound: chance alignments of stars at different distances that fall within the same fiber can produce spectra consistent with binarity.

To estimate the false-positive rate due to such “optical binaries” that are not gravitationally bound, we analyze mock photometric catalogs created with *Galaxia* (Sharma et al. 2011). *Galaxia* implements the Besançon model of stellar population synthesis (Robin et al. 2003) to populate the Galactic distribution function and produce realistic mock surveys along arbitrary lines of sight. Using a Galactic dust extinction map computed by *mw dust* (Bovy 2016a), we produced mock catalogs complete to $J = 14$ mag along 3 lines of sight representative of the range of stellar densities spanned by different APOGEE fields: one towards the bulge with $(\ell, b) = (0 \text{ deg}, 5 \text{ deg})$, one toward the Galactic anticenter with $(\ell, b) = (180 \text{ deg}, 0 \text{ deg})$, and one at high latitude with $(\ell, b) = (0 \text{ deg}, 60 \text{ deg})$, where ℓ and b represent Galactic longitude and latitude. We then checked, for each star in a mock catalog,

¹¹We calculate $a \sin i$, $M_1 \sin^3 i$, and $M_2 \sin^3 i$ using the standard formulas from Cox (2000). It is not possible to measure a or M directly from radial velocity data alone; we note that future astrometric constraints can break the degeneracy between these quantities and orbital inclination for nearby systems (Halbwachs et al. 2017b).

¹²The only short-period system which is distinctly noncircular and is not best fit by a triple model is 2M21320320+1107560, with $P = 6.70$ days, $e = 0.41$, and 41 epochs. It may well also be part of a hierarchical triple in which the third (long-period) component is too faint to appreciably contribute to the spectrum. Such a system would not be identifiable as having an unseen companion, as only systems in which the unseen companion is in the short-period sub-binary have velocities inconsistent with being an isolated binary.

whether any other stars fall within a circular aperture of diameter 2 arcseconds centered on that star, representing a single APOGEE fiber. If more than one star was found in a given aperture (including both dwarfs and giants), we classified all stars in that aperture as a single optical binary.

Toward the Galactic Bulge, we find a 0.2% probability that a star is an optical binary. The same probability is 0.05% toward the Galactic anticenter and 0.005% at high latitude. As these probabilities are all much smaller than our detected binary fraction of $\sim 15\%$, we conclude that optical binaries are unlikely to be a large source of false positives, though a small fraction of the systems we detected in fields toward the Bulge may be chance alignments masquerading as binaries.

Our model requires all components of a multiple-star system to have identical distances and abundances and fall on a single isochrone. This is likely a reasonable assumption for true, gravitationally bound binaries (Desidera et al. 2004; Andrews et al. 2017), but it is unlikely to hold for chance alignments. One could thus distinguish between true binaries and chance alignments by allowing the stellar parameters and abundances, and/or relative distance, of the secondary to vary freely and identifying cases where the best-fit model assigns significantly different abundances or distances to the different components. We defer such analysis to future work.

3.5 Discussion and Conclusions

3.5.1 Comparison to previous work

Chojnowski et al. (2015) compiled a catalog of double-lined spectroscopic binaries and triples in APOGEE by identifying targets whose cross-correlation function exhibited multiple peaks.¹³ Of the 610 targets in their catalog that were also in our initial sample, 574 were also classified as multiple systems by our pipeline, 5 were classified as SB1s, and 31 were classified as consistent with being single stars. Of the 574 stars classified as multiple systems by both pipelines, 514 are in our “potential close binary” subsample, which contains variable-velocity targets and binaries with large velocity offsets between the two components (Section 3.3.4). This $\sim 95\%$ agreement rate is encouraging, given the very different approaches of the two pipelines. A primary advantage of the method developed in this work is its increased sensitivity to long-period systems with negligible velocity offsets.

Recently, Badenes et al. (2017) studied the occurrence rate of short-period, velocity variable binaries in APOGEE. They found the multiplicity fraction for main-sequence stars to be a factor of ~ 2 higher in the lowest-metallicity tercile of their sample than in the highest-metallicity tercile. We find a similar result: for short-period systems (those with velocity shifts of at least 10 km s^{-1} between epochs), the multiplicity fraction is $\sim 60\%$ higher for the lowest-metallicity tercile of our sample ($[\text{Fe}/\text{H}] < -0.21$) than for the highest-metallicity tercile ($[\text{Fe}/\text{H}] > -0.02$). Badenes et al. (2017) studied systems with metallicities

¹³Their catalog is available at <http://astronomy.nmsu.edu/drewski/apogee-sb2/apSB2.html>

as low as $[\text{Fe}/\text{H}] = -2.5$; given the smaller range of metallicities in our sample ($[\text{Fe}/\text{H}] > -1$), these results are likely consistent. For long-period systems, we find the binary fraction to be consistent with being constant with metallicity. Some theoretical models (e.g. [Machida 2008](#)) predict that low-metallicity clouds should preferentially form short-period binaries, consistent with this result. However, we caution against over-interpreting this finding, as we have not attempted to quantify or correct for changes in the completeness of our method at lower metallicity.

Modeling approaches similar to the method developed in this work have previously been used on a case-by-case basis to fit “composite spectrum binaries”, a term that refers specifically to binaries containing a cool giant primary and a hot subgiant or main-sequence secondary (e.g. [González & Levato 2006](#); [Griffin & Griffin 2010](#)). Similar techniques have also been employed to spectroscopically detect and characterize unresolved binaries composed of very low-mass stars or brown dwarfs with different spectral types ([Burgasser 2007](#); [Burgasser et al. 2008](#)). Although this work focuses on modeling the spectra of binaries in which both components are main-sequence stars, the method we develop is flexible and can be straightforwardly extended to identify other flavors of binaries. The primary requirement is a robust training set spanning the range of single-star spectral types found in the dataset of interest.

For systems known to be double-lined binaries, a wide variety of techniques have been developed to disentangle the spectra of the two component stars in order to measure their individual velocities and stellar labels (e.g. [Bagnuolo & Gies 1991](#); [Simon & Sturm 1994](#); [Hadrava 1995](#); [Pavlovski & Hensberge 2010](#); [Czekala et al. 2017](#)). These techniques can reliably separate the spectra of the individual components of a binary even when lines are blended, but they generally require multi-epoch spectroscopy that captures the combined binary spectrum at several orbital configurations.

If only single-epoch spectroscopy is available or the binary is sufficiently wide that the orbital velocities of the two components do not change much between visits, the most common approach for measuring radial velocities is cross-correlation with a composite template spectrum ([Zucker & Mazeh 1994](#); [Halbwachs et al. 2017a](#)); this requires first estimating the labels of the individual stars. In such cases, most previous works have attempted to first model the primary star with a synthetic or empirical template, then subtract this template from the composite spectrum, and finally fit a model for the second star to the residual spectrum. However, it is difficult to ensure with this approach that the optimal binary model has been found, as the single-star model spectrum that best fits the combined binary spectrum does not in general correspond to the true best-fit parameters of the primary star (see [E18](#)). The method introduced in this work, which fits for the stellar parameters and velocities of both components simultaneously, avoids these complications.

Recently, a few works have shown that double-lined binaries can also be detected non-parametrically by identifying systems with peculiar spectra that are clustered outliers in a high-dimensional space of arbitrary summary statistics computed for all spectra collected by a survey ([Traven et al. 2017](#); [Reis et al. 2017](#)). While such methods thus far primarily identify binaries with large velocity offsets, we note that non-parametric methods can likely be further optimized for binary identification by searching for targets which are precisely




Figure 3.12: Schematic illustration of the range of binary periods and mass ratios that can be detected with different methods. Gray shading (identical in all panels) shows the distribution of periods and mass ratios for solar-type stars; hatches show the regions of parameter space that can be probed by radial velocity variability (left) and fitting a binary model to single- and multi-epoch spectra (middle, right). Conventional multi-epoch radial velocity surveys are sensitive to essentially all mass ratios, but only for short-period binaries, which represent roughly a third of the observed lognormal period distribution for solar-type stars. The binary spectral model introduced in this work is sensitive to all but the longest periods (as long as both stars fall within one spectroscopic fiber) with single-epoch observations, but only for intermediate mass ratio systems. When multi-epoch spectra are available, fitting a binary model can also detect all systems with variable radial velocities as SBIs.

the kind of outliers expected to result from binarity; for example, one could identify systems that cannot be well-described by a single combination of spectral PCA components but can be well-described by two sums of components with different velocities.

3.5.2 Future prospects

Improving the model

A straightforward way to make our model sensitive to a larger fraction of the binary population is to extend the single-star model to cooler temperatures. As discussed in Appendix 3.7.1, the lower limit of $T_{\text{eff}} = 4200$ K, which is due to shortcomings of ab-initio spectral models for main-sequence stars at lower temperatures, limits the model to only detecting binaries with mass ratios near $q = 1$ at low temperatures and prevents us from fitting spectra of the coolest stars altogether. Fitting cooler stars does not require any modification of our general approach, only a robust training set at lower temperatures, which currently does not readily exist. Due to the increased importance of molecular opacity from many species at lower T_{eff} , it may be helpful to include more abundances in the spectral

model for cooler stars.

Our model could also be improved by fitting for the projected rotation velocity $v \sin i$ explicitly instead of subsuming it under the Gaussian broadening of v_{macro} , since the single-star model currently performs worst for rapidly rotating stars. This is in principle simple to accomplish: rotation velocities for stars in the training set can be obtained straightforwardly in post-processing (e.g. Díaz et al. 2011), and the inferred $v \sin i$ can then be added as an additional label to the model. Rotation is currently not an important problem for most of the targets in our sample because stars with $T_{\text{eff}} \lesssim 6500$ K typically lose most of their angular momentum to magnetic braking and do not rotate rapidly on the main sequence (Glebocki et al. 2000; Schatzman 1962), and stars with $T_{\text{eff}} \geq 6500$ K represent less than 4% of our dataset. However, an improved treatment of rotation would make it possible to better model hot stars and would likely decrease the false-positive rate (see Appendix 3.8). This is particularly true for young stars, which can rotate significantly even at cooler temperatures (e.g. Terrien et al. 2014).

Hierarchical modeling

Beyond the Solar neighborhood, previous spectroscopic studies of the Galactic binary population have been limited to studying the short-period tail of the binary population. Because the model presented in this work does not depend on radial velocity variability or a line-of-sight velocity offset to detect binaries, it has the potential to substantially improve on existing constraints on the binary population of the Milky Way and/or its satellites when combined with a model for detection completeness and the survey selection function.

Existing radial velocity surveys of the Milky Way and nearby dwarf galaxies are sensitive to binaries with periods less than $\sim (1 - 10)$ years (e.g. Matijevič et al. 2011; Minor 2013; Hettlinger et al. 2015; Gao et al. 2017; Badenes et al. 2017). For the log-normal period distribution for solar-type stars found in the Solar neighborhood (Duchêne & Kraus 2013), $\sim 73\%$ percent of binaries have $P > 10$ years; most of these systems will be missed by such surveys. The most probable period for solar-type binaries is ~ 300 years; assuming random orbit orientations, the typical line-of-sight velocity separation for such systems is $\Delta v_{\text{los}} \sim 2 \text{ km s}^{-1}$, and the average radial velocity change over a one-year baseline is $\sim 0.02 \text{ km s}^{-1}$. This is an order of magnitude below the detectability thresholds of existing large spectroscopic surveys, though such weak radial velocity trends in SBIs may be marginally detectable with high-dispersion spectrographs typically used to study exoplanets (Konacki 2005; Katoh et al. 2013).

Irrespective of radial velocity variability, long-period binaries with favorable mass ratios can be detected with our model as long as both components fall within a single spectroscopic fiber. At a distance of 1 kpc, more than 80% of solar-type binaries will have projected separations of less than 1 arcsecond, so that both stars would fall with a single 2-arcsecond fiber; this fraction increases at larger distances. On the other hand, for long-period systems, the binary spectral model is sensitive only to intermediate mass ratio systems ($0.4 \lesssim q \lesssim 0.8$), in which the primary and secondary have qualitatively different spectral types, but the

secondary still contributes a non-negligible fraction of the total light (see Appendix 3.7.1 and E18). The distribution of mass ratios for solar-type binaries is approximately flat down to $q = 0.1$ (Duchêne & Kraus 2013), so the binary model will miss many high and low mass ratio systems with long periods.

We summarize the sensitivity of our method, as well as standard binary-detection methods based on velocity variability, to systems with different periods and mass ratios in Figure 3.12. Radial velocity variability can probe essentially all mass ratios, but only for the short-period tail of the binary population. On the other hand, fitting a binary spectral model to single-epoch observations can probe most of the period distribution, but only for a restricted subset of mass ratios. We thus expect that the method developed here can be fruitfully combined with existing multi-epoch radial velocity measurements from SB1s, such as the APOGEE constraints on the short-period binary fraction presented in Badenes et al. (2017) and measurements of the binary fractions of nearby dwarf galaxies presented by Minor (2013). This would enable a full hierarchical model for binary populations that is sensitive to an unprecedented range of periods and mass ratios.

An immediate advantage of our method is that it is sensitive to a large fraction of the binary population even when only single-epoch observations are available. With multi-epoch observations, our model can detect short-period systems as SB1s, with similar sensitivity to traditional methods. Our modeling approach can also be straightforwardly applied to spectra from other surveys. The precise range of mass-ratios to which it is sensitive will vary with wavelength coverage: surveys at optical wavelengths will be more sensitive to binaries with higher mass ratios ($0.8 \lesssim q \lesssim 0.9$; see E18) due to the increased spectral information content at shorter wavelengths, but they will be less sensitive at low q because a cooler secondary star contributes a greater fraction of a binary system’s total light in the near-infrared than at optical wavelengths.

In this work, we fit normalized spectra and only used the CMD to assess the reliability of our spectral model. A promising avenue for future work is to fit spectra and photometry simultaneously, or to place a photometric prior on q . This would make it possible to detect systems with $q \sim 1$ and negligible velocity offsets, which are twice as luminous as they would be if they were a single star. Particularly with improved parallaxes from future *Gaia* data releases, photometric constraints could substantially extend the fraction of the binary population to which our method is sensitive.

3.5.3 Summary

We have developed a flexible data-driven method for identifying and fitting the spectra of multiple-star systems and have applied it to $\sim 20,000$ main-sequence targets from the APOGEE survey. Unlike most previous work, our model performs well even for long-period systems in which the line-of-sight velocity offset between components is negligible, substantially expanding the fraction of the binary population that can be probed by observations. Our method is mostly automated and can be straightforwardly applied to other spectroscopic surveys with modest adjustments. Our main results are as follows:

1. *Spectral identification of long- and short-period binaries:* Unresolved binaries can be identified as systems whose spectrum can be better-fit by a sum of two single-star model spectra falling on a single isochrone than any single-star model (Figure 3.1). For systems with mass ratios $0.4 \lesssim q \lesssim 0.8$, in which the two stars have different spectral types, binaries can be identified spectroscopically even in the limit of no velocity offset and with only single-epoch observations. Spectral signatures of binarity are strengthened in the presence of a velocity offset of order one resolution element or greater (Figure 3.2); thus, close binaries can be detected even in the limit of $q \sim 1$.
2. *Photometric test of the model:* Nearly all spectroscopically identified binaries with accurate distance measurements fall above the main sequence on the CMD, as is predicted for true binaries, and triple systems fall above most binaries (Figure 3.9). Photometry does not enter our binary identification procedure, so this agreement with theoretical predictions provides independent validation of our spectral model.
3. *Dynamical mass ratios:* For short-period binaries in which the velocities of the two components change substantially between visits, it is possible to obtain a dynamical measurement of the mass ratio from the relative changes in the stars' radial velocities between visits (Figure 3.3). This provides a constraint on the mass ratio that is independent of the spectral mass ratio, which determines the contribution of the secondary star to the spectrum. We find good agreement between spectral and dynamical mass ratios, with a median difference of 0.048 and even better agreement for systems with high S/N spectra (Figure 3.4).
4. *Triple systems:* We identify 114 systems in which the contributions of three stars can be identified in the spectrum (Figure 3.6) and an additional 108 in which only two stars contribute significantly to the spectrum, but the presence of a third component can be inferred from its gravitational effects (Figure 3.8). Most identified triples are hierarchical, consisting of a close binary orbited by a third component with a much longer period; we have verified that these systems are all likely gravitationally bound (Figure 3.7).
5. *Orbital solutions:* For double-lined systems with a sufficient number of epochs and well-sampled radial velocity curves, we derive full Keplerian orbital solutions (Figure 3.10(b)); some of these systems are close binaries within hierarchical triples. We derive orbital solutions for 64 binaries with periods ranging from ~ 0.6 days to ~ 2 years and semimajor axes ranging from $\sim R_{\odot}$ to ~ 1 AU. Consistent with previous studies, we find that most binaries with $P \lesssim 10$ days have eccentricity consistent with 0 due to tidal circularization processes (Figure 3.11).

We make catalogs of best-fit labels for all identified multiple-star systems publicly available; these are described in Appendix 3.10.

Acknowledgements

We are grateful to the anonymous referee for a constructive report. We thank Gaspard Duchêne, Keith Hawkins, Jessica Lu, Hans-Günter Ludwig, Adrian Price-Whelan, and Silvia Toonen for helpful conversations. We are grateful to Jan Rybizki for assistance in creating mock catalogs with `Galaxia`, and to Anna Ho for help with *the Cannon*. This project was developed in part at the 2017 Heidelberg Gaia Sprint, hosted by the Max-Planck-Institut für Astronomie, Heidelberg. K.E. acknowledges support from the SFB 881 program (A3), a Berkeley Fellowship, a Hellman award for graduate study, and an NSF graduate research fellowship. H.W.R. received support from the European Research Council under the European Union’s Seventh Framework Programme (FP 7) ERC Grant Agreement n. [321035]. Y.S.T is supported by the Australian Research Council Discovery Program DP160103747, the Carnegie-Princeton Fellowship and the Martin A. and Helen Chooljian Membership from the Institute for Advanced Study at Princeton. E.Q. is supported in part by a Simons Investigator Award from the Simons Foundation. D.R.W. is supported by a fellowship from the Alfred P. Sloan Foundation. C.C. acknowledges support from NASA grant NNX15AK14G, NSF grant AST-1313280, and the Packard Foundation. The analysis in this paper relied on the python packages NumPy (Van Der Walt et al. 2011), Matplotlib (Hunter 2007), and AstroPy (Astropy Collaboration et al. 2013a).

3.6 Neural network spectral model

As mentioned in Section 3.3.1, we use a neural network to predict the normalized flux density at a given wavelength pixel as a function of stellar labels. As applied in this work, a neural network is essentially a flexible function produced through the composition of simple functions. It takes as its argument a vector of labels ($\vec{\ell}$; Equation 3.1) and returns the normalized flux density predicted at a particular wavelength pixel.

For the neural network used in this work, which contains a single hidden layer with 5 neurons, the normalized flux density at wavelength pixel λ is given by

$$\hat{f}_\lambda = \tilde{w}_\lambda^i \sigma \left(w_{\lambda i}^k \hat{\ell}_k + b_{\lambda i} \right) + \tilde{b}_\lambda \quad (3.10)$$

with implied summation over $k = 1 \dots N_{\text{labels}}$ and $i = 1 \dots N_{\text{neurons}}$. Here $\hat{\ell} = (\vec{\ell} - \vec{\ell}_{\text{min}}) / (\vec{\ell}_{\text{max}} - \vec{\ell}_{\text{min}}) - 0.5$ is a scaled label vector, $\vec{\ell}_{\text{max}}$ and $\vec{\ell}_{\text{min}}$ are vectors of the maximum and minimum values of each label in the training set, and $\sigma(z) = 1 / (1 + e^{-z})$ is the “sigmoid” activation function. The *weights*, w and \tilde{w} , and *biases*, b and \tilde{b} , parameterize the neural network; these are the free parameters that are adjusted during training.

In order to treat spectra with different line-of-sight velocities, all spectra are shifted to rest-frame and linearly interpolated onto a common wavelength grid. Training the model consists of minimizing a *loss function*, comparable to the χ^2 statistic, that quantifies how well the model can fit the training set. We use an L1 loss function, which minimizes the total

absolute difference between fluxes predicted by the neural network and those in the training set. We expect this to perform better than e.g. the χ^2 statistic during the iterative cleaning of the training set and re-training, since it is less sensitive to outliers. During training, we mask all pixels with S/N < 50, bad or missing pixels, and pixels with poor sky subtraction.

We implement and train the neural network using the python package PyTorch. We tested a wide range of network architectures, varying the network depth, width, and activation function, with both data-driven and synthetic spectral models. We find that using a small neural network and a large training set is the most straightforward way to prevent overfitting; using a substantially larger network with more neurons or hidden layers causes the model to reach lower losses (i.e., fit the training set better) but perform worse in cross-validation. We verified that our spectral model performs equally well on the training and cross-validation sets at fixed S/N, so it does not overfit the training set.

An advantage of using a neural network spectral model is that the neural network’s flexibility makes it possible to model a wide range of stellar parameters in a single model, rather than stitching together multiple models covering different regions of label space. However, our basic approach of constructing a binary spectral model does not depend critically on use of a neural network; a comparable binary model could likely be built from other forms of single-star model (e.g., *The Cannon*).

3.7 Model selection

Because the single-star model is a special case of the binary model, it is *always* possible to obtain a binary model that fits a data spectrum at least as well as does the best-fit single-star model. As the binary model is more complex than the single-star model, with three additional free parameters, one might expect to find a better fit, in a χ^2 sense, with the binary model even for targets which are true single stars. It is therefore necessary to formulate a heuristic to determine “how much” better a fit with a binary model is required to constitute reliable evidence in favor of the binary model.

The primary statistic used for model selection is the χ^2 difference, $\Delta\chi^2 = \chi_{\text{single}}^2 - \chi_{\text{binary}}^2$, which simply quantifies how much better a fit is obtained by the binary model. We also calculate a second statistic, the “improvement fraction” f_{imp} , to quantify how much better a fit the binary model achieves *relative to how different it is from the single-star model*. The basic idea here is that if a binary model spectrum is very different from the single-star model, but only achieves a slightly better fit to the data, this constitutes weaker evidence in favor of the binary model than a case with comparable $\Delta\chi^2$ in which most of the difference between the best-fit binary and single-star model goes toward improving the fit. The improvement fraction is defined as

$$f_{\text{imp}} = \frac{\sum \left\{ \left(\left| \hat{f}_{\lambda, \text{single}} - \hat{f}_{\lambda} \right| - \left| \hat{f}_{\lambda, \text{binary}} - \hat{f}_{\lambda} \right| \right) / \hat{\sigma}_{\lambda} \right\}}{\sum \left\{ \left| \hat{f}_{\lambda, \text{single}} - \hat{f}_{\lambda, \text{binary}} \right| / \hat{\sigma}_{\lambda} \right\}}, \quad (3.11)$$

where \hat{f}_λ and $\hat{\sigma}_\lambda$ are the normalized flux density and corresponding uncertainty, $\hat{f}_{\lambda,\text{single}}$ and $\hat{f}_{\lambda,\text{binary}}$ are the best-fit normalized single-star and binary model spectra, and the sum is over all wavelength pixels.

Our full acceptance criterion for preferring the binary model is given in Table 3.2. These thresholds were motivated in part by the $\Delta\chi^2$ and f_{imp} values calculated for semi-empirical binaries as described below, and in part by validation with the CMD (Figure 3.9). The adopted thresholds are conservative, and prevent us from identifying some binaries whose spectra can only be marginally better fit by a binary model; however, setting them to substantially lower values causes the model to begin categorizing more targets near the main sequence of the CMD as binaries, indicating a non-negligible false-positive rate. We have not tuned the minimum f_{imp} value for each $\Delta\chi^2$; we set the intuitive requirement that a higher f_{imp} should be required for systems with a lower $\Delta\chi^2$.

Different $\Delta\chi^2$ thresholds are required for systems identified as potential close binaries, because (a) for these systems, we fit multiple visit spectra simultaneously, and (b) we fit a total of 5 different models (see Section 3.3.4). For these systems, we begin with a fiducial threshold of $\Delta\chi^2 = 300 \times N_{\text{epochs}}$ for each increase in model complexity, where N_{epochs} is the number of visit spectra fit simultaneously; i.e., we require $\chi_{\text{single star}}^2 - \chi_{\text{SB1}}^2 > 300N_{\text{epochs}}$ for a system to be initially classified as an SB1; $\chi_{\text{SB1}}^2 - \chi_{\text{SB2}}^2 > 300N_{\text{epochs}}$ for a system to be initially classified as an SB2, etc. We then inspected the spectra of these targets individually and reclassified suspected false positives (see Appendix 3.8).

We experimented with generating single-star spectra directly from the spectral model, adding noise, and fitting them with both a binary and single-star model. This produces typical $\Delta\chi^2$ values of order unity, which is smaller than we find for the majority of APOGEE targets. This occurs because spectra generated in this way can necessarily be perfectly fit by the single-star spectral model, which is not necessarily the case for real spectra. We quantify the $\Delta\chi^2$ values expected for real binary spectra in the next section.

3.7.1 Tests with semi-empirical synthetic binary spectra

To assess the accuracy and potential systematics of our method, and to measure the expected $\Delta\chi^2$ values for true binaries with a particular mass ratio, we construct a library of $\sim 15,000$ “semi-empirical” synthetic binary spectra. These are created by combining randomly chosen pairs of APOGEE spectra in unnormalized space, following the method outlined in Section 3.3.2. We then normalize and fit these spectra following the same procedure used to fit real spectra. An advantage of constructing synthetic binary spectra by combining real spectra (as opposed to simply generating binary spectra from the data-driven model) is that this accounts for the possibility that the model does not capture all the variance in the real spectra; this is likely the case for our model, which only contains 5 labels.

We require that the two stars used to construct a semi-empirical binary spectrum have similar abundances (within 0.05 dex in [Fe/H] and $[\alpha/\text{Fe}]$) and fall within 0.03 dex in $\log g$ of a single isochrone. We only combine spectra consistent with being single stars; i.e., those which cannot be significantly better-fit by a binary model than a single-star model according




Figure 3.13: Results of fitting semi-empirical binary spectra with single-star and binary models. Semi-empirical binary spectra are created by adding together flux-calibrated APOGEE spectra of two stars with similar abundances. At $q \lesssim 0.8$, most semi-empirical binary spectra can be significantly better fit with a binary model, with $\Delta\chi^2 \gtrsim 1000$. The χ^2 difference is nearly always larger for systems with large velocity offsets; for systems with $q \sim 1$, only binaries with $\Delta v_{\text{los}} \gtrsim 10 \text{ km s}^{-1}$ have $\Delta\chi^2 \gtrsim 1000$. At fixed q and Δv_{los} , the typical $\Delta\chi^2$ is larger for systems with cooler primaries. Due to our single-star spectral model’s minimum T_{eff} of 4200 K, low mass ratio systems can only be modeled for hot primaries.

$\Delta\chi^2 = \chi_{\text{single}}^2 - \chi_{\text{binary}}^2$	minimum f_{imp}
$\Delta\chi^2 \geq 3000$	0
$2500 \leq \Delta\chi^2 < 3000$	0.05
$2000 \leq \Delta\chi^2 < 2500$	0.075
$1500 \leq \Delta\chi^2 < 2000$	0.1
$1000 \leq \Delta\chi^2 < 1500$	0.125
$750 \leq \Delta\chi^2 < 1000$	0.15
$600 \leq \Delta\chi^2 < 750$	0.175
$450 \leq \Delta\chi^2 < 600$	0.2
$300 \leq \Delta\chi^2 < 450$	0.225

Table 3.2: Minimum $\Delta\chi^2$ and improvement fraction f_{imp} (Equation 3.11) for a target to be classified as a binary. All systems with $\Delta\chi^2 < 300$, and all systems falling below the minimum f_{imp} for a given $\Delta\chi^2$, are classified as inconclusive; i.e., showing no strong evidence of binarity.

to the thresholds in Table 3.2. We assign realistic orbital parameters to each system following E18, drawing orbital periods from the log-normal period distribution for solar-type stars from Duchêne & Kraus (2013) and assuming random orbit orientations and phases. Results from fitting these semi-empirical binary spectra are shown in Figures 3.13, 3.14, 3.15, and 3.16.

In Figure 3.13, we show the χ^2 difference in favor of the binary model, $\Delta\chi^2 = \chi_{\text{single}}^2 - \chi_{\text{binary}}^2$, as a function of the mass ratio q , line-of-sight velocity offset Δv_{los} , and T_{eff} of the primary star. As expected, $\Delta\chi^2$ is a strong function of q : most binaries with $q \lesssim 0.75$ have $\Delta\chi^2 > 1000$, while most systems with $q \sim 1$ have much lower $\Delta\chi^2$. This is expected, because the two stars in binaries with $q \sim 1$ will have similar spectra, making the combined binary spectrum indistinguishable from that of either single star unless there is a sizable velocity offset between the two stars. The left panel of Figure 3.13 shows that $\Delta\chi^2$ is also a strong function of the velocity offset Δv_{los} : at fixed q , systems in which the velocity offset is larger nearly always have larger $\Delta\chi^2$. In particular, most binaries with $\Delta v_{\text{los}} \gtrsim 10 \text{ km s}^{-1}$ have $\Delta\chi^2 > 1000$, even at $q \sim 1$. Indeed, among binaries with large velocity offsets, the typical $\Delta\chi^2$ is largest for systems with $q \sim 1$; in such systems, both stars contribute significantly to the spectrum, and absorption lines are obviously split.

The right panel of Figure 3.13 shows the dependence of $\Delta\chi^2$, and the range of mass ratios to which our method is sensitive, on T_{eff} of the primary. At fixed mass ratio and Δv , the median $\Delta\chi^2$ is slightly lower for hot stars ($T_{\text{eff}} > 6000 \text{ K}$), particularly for systems with $q \sim 1$ and large Δv . This occurs because lines are on average weaker and more rotationally broadened in hot stars, reducing the information content of the spectrum. This panel also shows that, due to the minimum T_{eff} of 4200 K of the single-star spectral model, the minimum mass ratio that can be modeled varies with T_{eff} of the primary. This means that our completeness is higher for hot stars than for cool stars.

To determine whether our binary model fit has converged on the true globally optimal model rather than a local χ^2 minimum, we check whether the χ^2 value of the best-fit binary

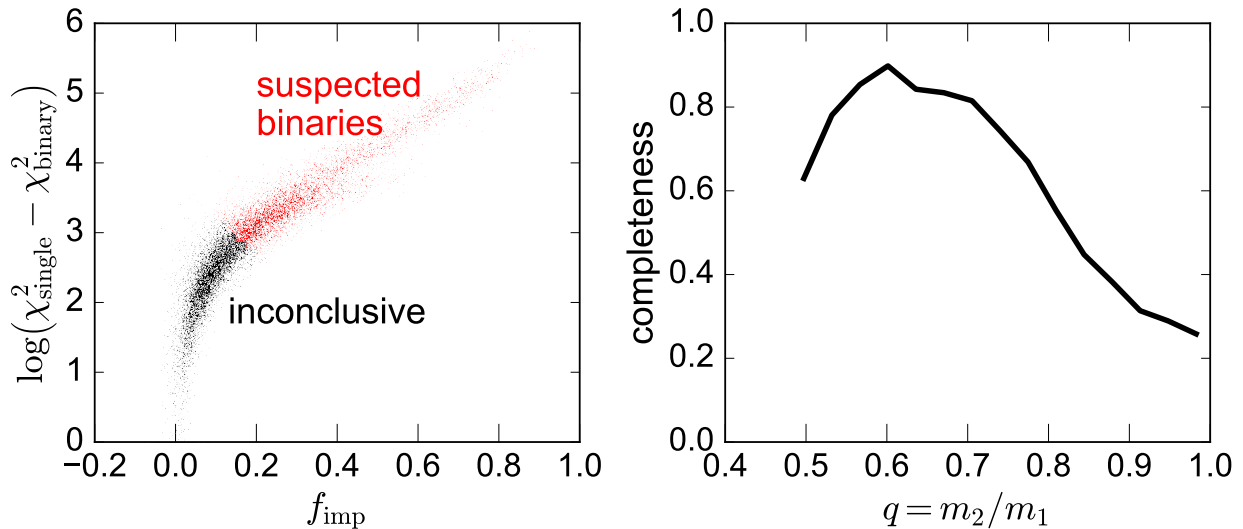


Figure 3.14: Results for fitting semi-empirical synthetic binary spectra. Left panel shows the χ^2 difference and improvement fraction (Equation 3.11); systems passing the adopted acceptance criterion for a binary candidate to be considered legitimate are plotted in red. Right panel shows the resulting completeness function; i.e., the fraction of semi-empirical binary systems at a given q that are successfully identified as binaries.

model is at least as low as that corresponding to the binary model with the true labels of the system. We find that our fit converges on the globally optimal solution for $\sim 99\%$ of all semi-empirical binaries. About half of the systems in which the fit converges on a local minimum are binaries with $q \sim 1$ in which the velocity assignments of the primary and secondary star are switched; in these cases, the derived stellar labels are still reasonably accurate.

In Figure 3.14, we show how our adopted model selection threshold translates to the range of mass ratios to which the model is sensitive. In the left panel, we plot the distribution of semi-empirical binaries in $\Delta\chi^2 - f_{\text{imp}}$ space. $\Delta\chi^2$ and f_{imp} are correlated: most systems whose spectrum can be significantly better fit by a binary model (high $\Delta\chi^2$) also have high f_{imp} . In the right panel, we plot the fraction of semi-empirical binaries at a given mass ratio that pass our adopted model selection criteria to be considered reliable binary candidates. As expected, this “completeness” function is a strong function of mass ratio. Most binaries with $0.55 \lesssim q \lesssim 0.75$ have spectra that are sufficiently different from any single-star model that they can be unambiguously identified as binaries; at higher mass ratios, the spectra of the two component stars become similar, so only the $\sim 20\%$ of binaries with $\Delta v_{\text{los}} \gtrsim 10 \text{ km s}^{-1}$ can be detected. At sufficiently low mass ratios, the secondary contributes a negligible fraction of the total light.

We emphasize that this completeness function does *not* represent our global completeness function for all binaries, for two reasons. First, the population of binaries in our semi-empirical library is not statistically representative of the Galactic binary population: because

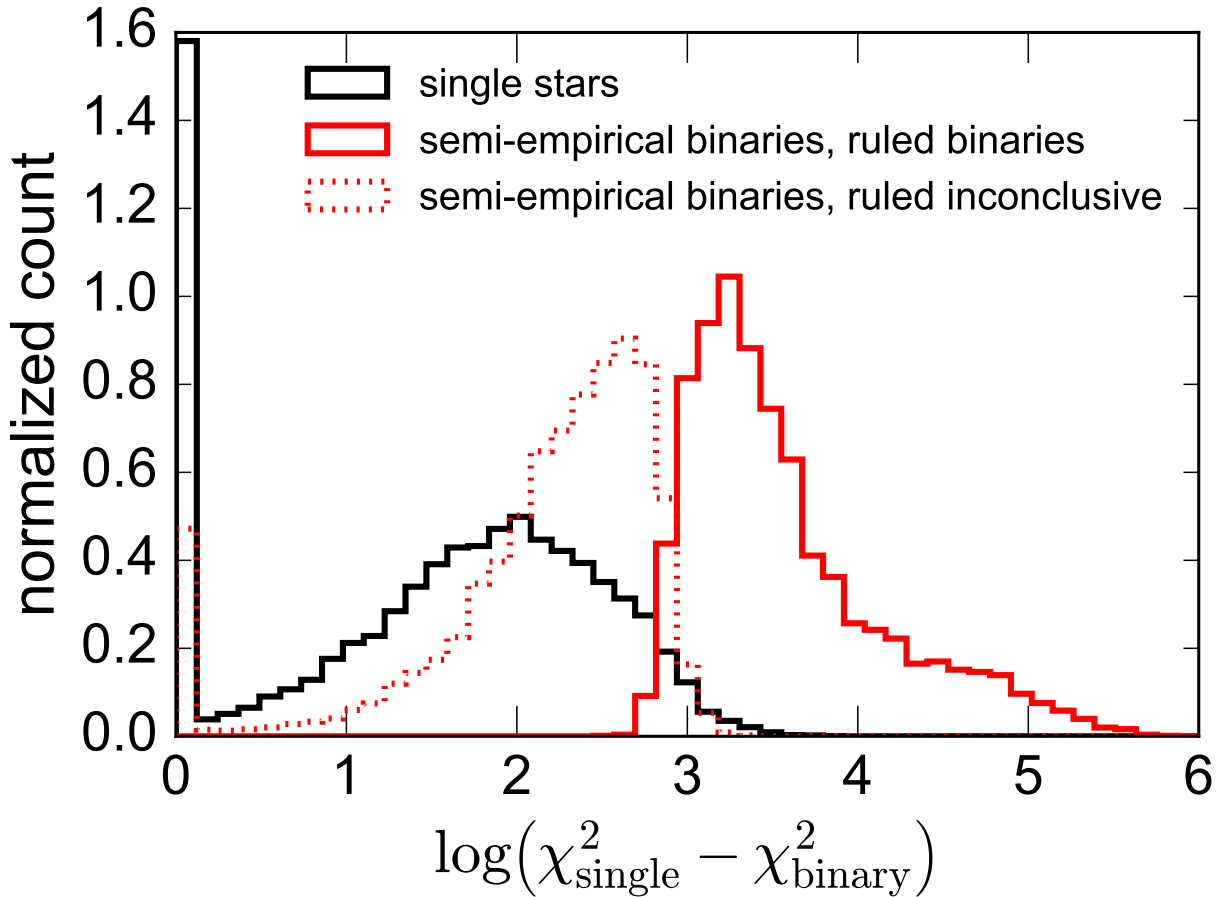


Figure 3.15: χ^2 differences between best-fit single and binary models. Black histogram shows suspected single stars from which semi-empirical binary spectra are constructed. Solid (dotted) red histogram shows semi-empirical binary spectra which pass (fail) the $\Delta\chi^2$ and f_{imp} binary acceptance criterion.

our single-star spectral model cannot model stars with $T_{\text{eff}} < 4200$ K, we cannot currently model low mass ratio binaries in which the primary is cool, as the secondary will have $T_{\text{eff}} < 4200$ K (see Section 3.3.2). Second, we have made no attempt to model the APOGEE selection function, which would complicate the distribution of T_{eff} at a given q . For the particular set of semi-empirical binaries analyzed here, $\sim 50\%$ of all binary systems pass the model selection threshold to be characterized as binaries. However, this is not representative of the global sensitivity of our model, since we make no attempt to use a realistic distribution of mass ratios in the semi-empirical library: the distribution of mass ratios in our semi-empirical binary library is skewed toward $q = 1$, which is precisely the regime in which the model performs poorly.

In Figure 3.15, we compare the distribution of $\Delta\chi^2$ values for the suspected single stars

used in constructing the semi-empirical binary library to those for semi-empirical synthetic binary spectra. Systems for which $0 < \Delta\chi^2 < 1$ are assigned $\log \Delta\chi^2 = 0$ on this plot; each histogram is normalized separately. The median $\Delta\chi^2$ is ~ 50 for single stars, ~ 200 for semi-empirical binaries that fail the detection criteria to be considered reliable binaries, and ~ 2600 for semi-empirical binaries that pass the threshold to be considered reliable. There is some overlap in the distribution of χ^2 values for single stars and semi-empirical binaries, since for systems with small Δv_{los} , the binary model in principle transitions smoothly to the single-star model both as $q \rightarrow 1$ and $q \rightarrow 0$. However, for binaries with favorable mass ratios, the typical $\Delta\chi^2$ is more than an order of magnitude greater than for single stars.

Cross validation

In Figure 3.16, we compare the best-fit labels for all semi-empirical binaries that pass the binary detection threshold to the true labels used in constructing the semi-empirical binary spectra. We set the “true” abundance for each binary as the average of the abundances of the two single-stars, which we required to be within 0.05 dex for $[\text{Fe}/\text{H}]$ and $[\alpha/\text{Fe}]$. T_{eff} and $\log g$ refer to the primary. These semi-empirical binaries were not used to train the spectral model, so this experiment constitutes cross validation of the binary spectral model. In each panel, we indicate the median signed error (bias) and absolute error (scatter) in the best-fit label. Overall, this experiment reveals that labels inferred from fitting our binary model are reasonably precise: the true and best-fit labels fall near the one-to-one line, with small scatter. The only label for which this is not obviously true is $v_{\text{macro},2}$; this occurs primarily because only hot stars have $v_{\text{macro}} \gtrsim 10 \text{ km s}^{-1}$, so only the few binaries in which both stars are hot have non-negligible $v_{\text{macro},2}$.

The median error in the best-fit q for the semi-empirical binaries is 0.021, which is smaller than the median difference of 0.048 between q_{dyn} and q_{spec} found for real binaries in Figure 3.4. This is not unexpected, because (a) q_{dyn} also has nonzero uncertainty, and (b) “ q_{true} ” for the semi-empirical binaries is calculated with the same isochrones used in the model from which q_{fit} is obtained. That is, the median difference of 0.021 does not account for uncertainties in the isochrones; the larger difference of 0.048 does, because q_{dyn} is independent of isochrones.

We emphasize that while Figure 3.16 shows our derived stellar labels to be *precise*, this does not guarantee that they are *accurate*. The reason for this is that the uncertainties in stellar parameters obtained from spectral fitting are often dominated by systematic uncertainties in the model, which enter primarily from errors in the “ground truth” labels of the training and validation sets, and are not accounted for in cross-validation. The cross-validation errors therefore are reasonable estimates of the precision of the model, but represent lower limits on the absolute uncertainties.

3.8 False positives

As discussed above, there are some targets for which our fitting and model selection formally prefers a binary model but visual inspection of the spectrum reveals that the evidence

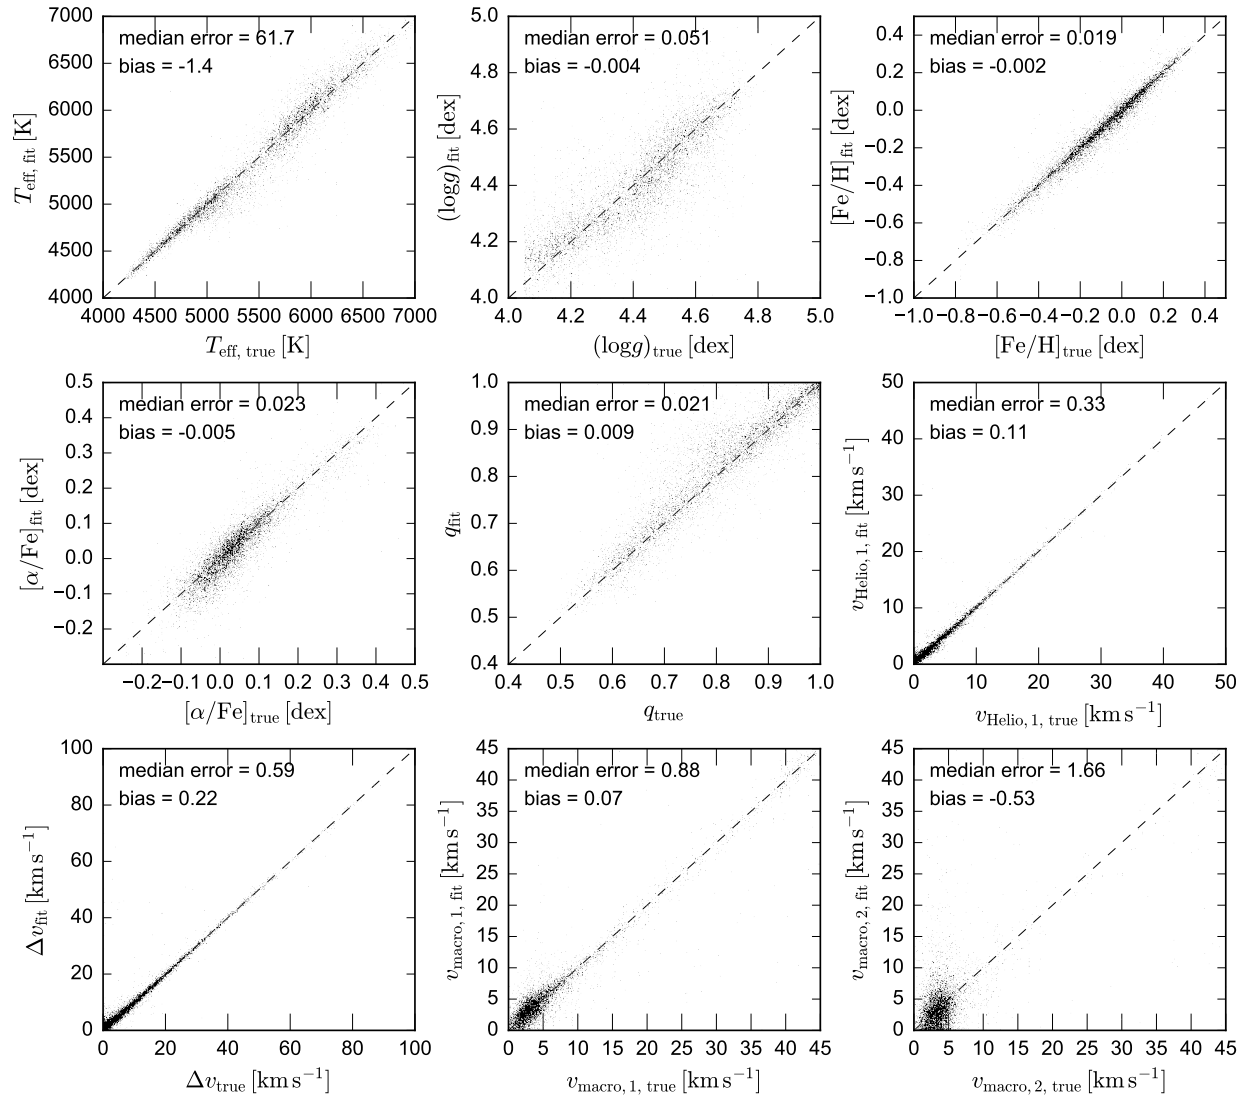


Figure 3.16: Label recovery diagnostic for semi-empirical binary spectra that pass our $\Delta\chi^2$ and f_{imp} threshold to be considered real detections. For each label ℓ_i , inset text indicates the median error, $\text{med}(|\ell_{i, \text{fit}} - \ell_{i, \text{true}}|)$, and bias, $\text{med}(\ell_{i, \text{fit}} - \ell_{i, \text{true}})$.

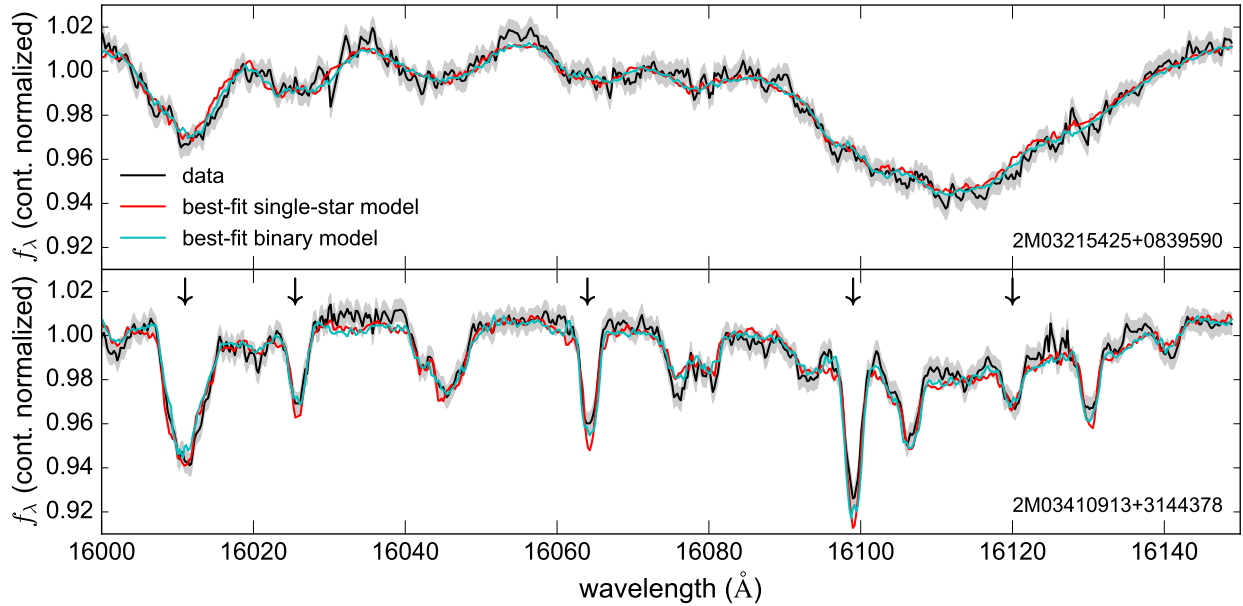


Figure 3.17: Spectra of two stars flagged as candidate binaries in our initial fitting that are likely false positives. **Top:** Nearly featureless spectrum of a hot ($T_{\text{eff}} \approx 7000$ K), rotationally broadened star. The main difference between the single-star and binary model (which has $q = 1$ and $\Delta v_{\text{los}} = 40 \text{ km s}^{-1}$) is that the binary model produces slightly broader lines. This target is an APOGEE telluric standard. **Bottom:** Although the binary model (with $q = 1$ and $\Delta v_{\text{los}} = 25 \text{ km s}^{-1}$) is formally preferred, it has many marginally split lines, marked with arrows, which are not present in the data spectrum. This target is a young star in an embedded cluster.

in favor of the binary model is weak and a single-star model should likely be preferred. After visually inspecting all ~ 3300 targets for which our model selection thresholds preferred a model other than a single star, we flagged ~ 300 targets as probable false-positives. The vast majority of such cases are hot stars ($T_{\text{eff}} \gtrsim 6500$ K) exhibiting significant rotational broadening.

In Figure 3.17, we show example spectra of two stars flagged as false positives. The top panel shows the spectrum of a hot, rotating star. Such spectra are less informative than spectra of cooler stars with less rotational broadening: most metallic lines are intrinsically weaker, since more species are ionized at higher T_{eff} , and the profiles of individual lines are smeared out due to rotation. Although the binary model is formally preferred, with $\Delta\chi^2 \sim 1000$, it does not obviously fit the profiles of individual lines better than the single-star model. Indeed, the main difference between the binary and single-star models is that the binary model produces slightly wider blended lines.

Most of the stars we flag as false positives have spectra similar to this target and large v_{macro} . There is in principle no reason to expect the method model to generically fail for such targets; the most likely explanation is that our 5-label single-star spectral model is

not sufficiently complex to fully characterize the results of rotation. More than half of these targets are APOGEE telluric standard stars, which are selected to be used for telluric correction because they have nearly featureless spectra. Because hot, rotationally-broadened stars constitute only a small fraction of our initial sample, we simply visually inspected and removed questionable targets by hand. A straightforward method for automating this procedure in future work would be to remove all telluric standards and/or targets for which a Fourier transform of the spectrum reveals little power at high frequencies.

The bottom panel of Figure 3.17 shows the spectrum of a false-positive candidate that is not featureless. The binary spectral model achieves a fit that is formally better than the single-star model ($\Delta\chi^2 \sim 1500$), but unlike the single-star model, the binary model spectrum has noticeably split lines as a result of the velocity offset between the two components. Because the data spectrum does not show such split line profiles, the binary model fit is likely erroneous and formally preferred only because it produces wider and shallower lines than the single-star model can accommodate. False-positives like this one are rare; only a few dozen spectra are identified in which the binary model produces a less realistic line profile than the single-star model despite achieving a formally better fit. A significant fraction of these are targets in young embedded clusters. Young stars often exhibit spectral features that are uncommon in older stars, such as chromospheric emission and increased rotation in cooler stars. Our method is more susceptible to incorrectly preferring a binary model if the spectrum cannot be well accommodated by the single-star model; this is likely the case for these targets. In future studies, such false-positives can potentially be eliminated by using a more complex spectral model for single stars and ensuring that different varieties of “unusual” spectra are represented in the training set.

Besides reclassifying stars flagged as false-positives upon visual inspection to be single stars, we also reclassified some of the potential close binary targets (Section 3.3.4) from one multiple-star class to another. For example, we reclassified SB2 systems for which the best-fit velocities of the primary and secondary fall on a one-to-one line with *positive* slope as SB1s, and SB3s with broad lines that were not obviously better fit by the SB3 model than the SB2 model as SB2s. We attempted to be conservative in classifying systems as triples; i.e., some triple systems are likely miscategorized as binaries, but all systems classified as triples are unambiguously better fit by the SB3 model than the SB2 model.

3.9 Orbit fitting convergence

Sparse radial velocity data can often be well-fit by several families of qualitatively different orbits, particularly when there are few radial velocity measurements and/or phase coverage is poor (e.g. Price-Whelan et al. 2017). Due to the complex multi-modal structure of the posterior in these cases, standard MCMC techniques fail to fully explore the orbital parameter space in finite time, meaning that there is no guarantee that our orbit-fitting procedure will converge on the true solution or that the orbital parameter uncertainties found by `emcee` are reliable. In the left panel of Figure 3.18, we show an example APOGEE binary

in which the radial velocity data are not sufficiently constraining to yield an unambiguous orbital solution: (at least) two qualitatively different orbital solutions can fit the measured radial velocities.

To assess the number of epochs and phase + velocity coverage required for reliable orbit constraints, we generated synthetic radial velocity measurements for a population of synthetic binaries (as described in E18) with $4 < N_{\text{epochs}} < 12$ and phase coverage $0.3 < U_N V_N < 1$ and fit them using the same procedure described in Section 3.4.4. To ensure realistic survey cadence, we drew observation times from observations of real APOGEE stars. We added Gaussian noise to the synthetic data with $\sigma_{RV} = 0.2 \text{ km s}^{-1}$, which is typical for our observations.

The results of this experiment are shown in Figure 3.18. We label each fit successful if the true orbital parameters fall within the marginalized 90% credibility regions returned by MCMC fitting. Although it is in principle possible to constrain the orbit of an SB2 with as few as 5 radial velocity epochs,¹⁴ reliable constraints for can only be obtained from realistic data with $\gtrsim 7$ epochs. At fixed number of epochs, constraints are more reliable for systems with larger $U_N V_N$, as expected. Based on the results of this experiment, we only attempt to fit orbits to systems with $N_{\text{epochs}} \geq 7$ and $U_N V_N \geq 0.5$. All but two of the targets for which we provide an orbital solution have $U_N V_N > 0.6$; we caution that orbital parameters for systems with lower $U_N V_N$ may be less reliable.

We stress that the orbit of an SB2 can in general be constrained with fewer radial velocity measurements than that of an SB1. The basic reason for this is that even with only a few epochs, having velocity measurements for both stars pins down the system velocity γ exactly. Therefore, many of the families of orbits with different combinations of P, e , and γ that would be permitted if radial velocity measurements were only available for the primary can be excluded when the secondary is detected.

Other works have used more conservative limits to determine whether the available velocity data were sufficient to constrain a binary orbit. For example, Halbwachs et al. (2017b) require $N_{\text{epochs}} \geq 11$. The true probability of convergence on a local minimum depends on the radial velocity uncertainties, number of epochs, and the uniformity of observational coverage. This experiment indicates that $N_{\text{epochs}} \geq 7$ is usually sufficient for radial velocity data similar to what we obtain in this work, but the probability of convergence on an erroneous orbital solution is, of course, lower for systems with a larger number of epochs.

3.10 Data products

Here we make available the best-fit labels for all targets identified as multiple-star systems. We also provide a list of targets consistent with being single stars in order to make it possible to reconstruct our initial sample of 20,142 targets.

¹⁴Although 7 parameters are required to parameterize a two-body orbit, the system velocity γ and mass ratio $q = K_1/K_2$ can be obtained “for free:” if the radial velocities of the primary and secondary fall on a line $v_{\text{Helio},2} = \alpha v_{\text{Helio},1} + \beta$, the system velocity is $\gamma = \beta / (1 - \alpha)$ and the mass ratio is $q = -1/\alpha$.




Figure 3.18: Left: Example APOGEE binary system with 8 epochs in which the available velocity data is not sufficient to fully constrain the orbit. Top and bottom panels show two qualitatively different orbital solutions which can both fit the measured velocity data (identical in the two panels) well. Because the measured velocities are sparse, the posterior is multimodal. We do not provide orbital solutions for such systems. **Right:** Results of fitting synthetic radial velocity data for an SB2 population with realistic orbital parameters and survey cadence. The “success rate” indicates the fraction of systems in each cell of $N_{\text{epoch}} - U_N V_N$ space for which our fitting procedure converged on the true orbit from which RV measurements were generated. Orbits can usually be reliably constrained for $N_{\text{epochs}} \gtrsim 7$, though systems with poor phase coverage can still converge on incorrect orbital solutions. We only attempt to fit orbits for systems with $N_{\text{epochs}} \geq 7$ and $U_N V_N > 0.5$.

APOGEE ID
2M00000233+1452324
2M00001701+7052395
2M00003475+5723259
2M00004578+5654428
...

Table 3.3: List of targets identified as single stars. This table is available in its entirety (with 16834 rows) in machine-readable form.

APOGEE ID	T_{eff} [K]	$\log g$ [dex]	[Fe/H] [dex]	[Mg/Fe] [dex]	v_{macro} [km s ⁻¹]	Δv_{max} [km s ⁻¹]
2M00010204+0049037	6008	4.26	-0.41	0.00	6.48	2.93
2M00031962-0017109	6157	4.23	-0.23	-0.06	42.95	65.17
2M00041803+1519505	5941	4.92	0.15	0.17	2.21	0.78
2M00041859+7104111	4921	4.36	0.13	-0.01	5.64	20.26
...

Table 3.4: Best-fit labels for targets identified as SB1s. Δv_{max} is the maximum change in radial velocity between visits. This table is available in its entirety (with parameters for 663 systems) in machine-readable form.

In Table 3.3, we list all targets consistent with being single stars; labels for these systems will be released by Ting et al. (in prep). Labels for targets identified as SB1s are listed in Table 3.4; these are obtained by simultaneously fitting visit spectra. Labels for targets classified as binaries are listed in Table 3.5. As described in Section 3.3.3, these are obtained by simultaneously fitting visit spectra for potential close binaries, which are primarily RV-variable systems, and by fitting the combined spectrum for all other systems. For close binaries in which the orbital configuration changes substantially between visits, we list dynamical mass ratios and center-of-mass velocities. Labels for targets classified as SB2s in which the gravitational effects of an with an unseen third component can be detected are listed in Table 3.6; these are obtained by simultaneously fitting visit spectra. Labels for targets identified as triples in which all three components contribute to the spectrum are listed in Table 3.7; these are obtained by simultaneously fitting visit spectra.

APOGEE ID	T_{eff} [K]	$\log g$ [dex]	[Fe/H] [dex]	[Mg/Fe] [dex]	q_{spec}	$v_{\text{macro},1}$ [km s ⁻¹]	$v_{\text{macro},2}$ [km s ⁻¹]	q_{dyn}	γ [km s ⁻¹]
2M00003968+5722329	4517	4.62	0.05	-0.04	0.88	3.09	8.19	—	—
2M00012717+0128193	5048	4.51	0.07	0.06	0.75	0.00	14.57	—	—
2M00023179+1521164	4589	4.37	-0.20	0.10	0.99	19.29	16.15	0.84	-0.92
2M00024073+6354560	5664	4.26	0.18	-0.01	0.67	1.42	3.34	—	—
...

Table 3.5: Best-fit labels for targets identified as binaries in which both components contribute to the spectrum. T_{eff} and $\log g$ refer to the primary. For the 623 targets in which the orbital configuration changes substantially between visits, we provide the dynamical mass ratios, q_{dyn} , and center-of-mass velocity, γ . This table is available in its entirety (with parameters for 2423 systems) in machine-readable form.

APOGEE ID	T_{eff} [K]	$\log g$ [dex]	[Fe/H] [dex]	[Mg/Fe] [dex]	q_{spec}	$v_{\text{macro},1}$ [km s ⁻¹]	$v_{\text{macro},2}$ [km s ⁻¹]
2M00103470+0043200	4200	4.50	-0.21	0.21	1.00	19.98	5.40
2M00265252+6359169	6416	4.56	-0.04	-0.29	0.85	10.34	14.38
2M00310678+8508494	5742	4.44	0.01	-0.07	0.62	4.91	3.11
2M01194897+8532293	5272	4.54	-0.08	0.03	0.90	2.68	4.80
...

Table 3.6: Best-fit labels for targets in which two components contribute to the spectrum but the gravitational effects a third components can be detected (e.g. Figure 3.8). T_{eff} and $\log g$ refer to the primary. This table is available in its entirety (with parameters for 108 systems) in machine-readable form.

APOGEE ID	T_{eff} [K]	$\log g$ [dex]	[Fe/H] [dex]	[Mg/Fe] [dex]	q_2	q_3	$v_{\text{macro},1}$ [km s ⁻¹]	$v_{\text{macro},2}$ [km s ⁻¹]	$v_{\text{macro},3}$ [km s ⁻¹]
2M00182859+6207248	5398	4.33	0.22	-0.20	1.00	1.00	8.69	4.33	0.19
2M00285967+5931138	4655	4.20	-0.22	0.08	0.94	0.94	1.39	30.07	36.15
2M00470197+1751448	6386	4.07	-0.47	0.03	0.89	0.88	19.03	3.50	4.46
2M01103850+6655525	4756	4.60	0.07	-0.00	0.92	0.81	3.71	1.34	16.35
...

Table 3.7: Best-fit labels for SB3s, targets in which three components contribute to the spectrum (e.g. Figure 3.6). T_{eff} and $\log g$ refer to the primary. This table is available in its entirety (with parameters for 114 systems) in machine-readable form.

Chapter 4

Imprints of white dwarf recoil in the separation distribution of Gaia wide binaries

An earlier version of this article was previously published in El-Badry, K. and Rix, H.-W., 2018, MNRAS, 480, 4884.

4.1 abstract

We construct from *Gaia* DR2 an extensive and very pure ($\lesssim 0.2\%$ contamination) catalog of wide binaries containing main-sequence (MS) and white dwarf (WD) components within 200 pc of the Sun. The public catalog contains, after removal of clusters and resolved higher-order multiples, $> 50,000$ MS/MS, $> 3,000$ WD/MS, and nearly 400 WD/WD binaries with projected separations of $50 \lesssim s/\text{AU} < 50,000$. Accounting for incompleteness and selection effects, we model the separation distribution of each class of binaries as a broken power-law, revealing marked differences between the three populations. The separation distribution of MS/MS systems is nearly consistent with a single power-law of slope -1.6 over at least $500 < s/\text{AU} < 50,000$, with marginal steepening at $s > 10,000$ AU. In contrast, the separation distributions of WD/MS and WD/WD binaries show distinct breaks at $\sim 3,000$ AU and $\sim 1,500$ AU, respectively: they are flatter than the MS/MS distribution at small separations and steeper at large separations. Using binary population synthesis models, we show that these breaks are unlikely to be caused by external factors but can be explained if the WDs incur a kick of $\sim 0.75 \text{ km s}^{-1}$ during their formation, presumably due to asymmetric mass loss. The data rule out typical kick velocities above 2 km s^{-1} . Our results imply that most wide binaries with separations exceeding a few thousand AU become unbound during post-MS evolution.

4.2 Introduction

Wide binaries (semimajor axes $a \gtrsim 100$ AU) are thought to form when initially unbound pairs of young stars become bound as their birth cluster dissolves (van Albada 1968; Moeckel & Bate 2010; Kouwenhoven et al. 2010; Moeckel & Clarke 2011). The binding energies of wide binaries are low, comparable to the kinetic energy of a star moving with velocity 1 km s^{-1} . Their orbits are therefore fragile, and can be easily disrupted. Such disruption can occur either due to external gravitational perturbations or due to internal kicks originating from one of the component stars.

Much of the interest in wide binaries to date has stemmed from the prospect of using their separation distribution as a probe of the Galactic gravitation potential (e.g. Bahcall et al. 1985; Chanamé & Gould 2004; Yoo et al. 2004; Quinn et al. 2009; Monroy-Rodríguez & Allen 2014; Correa-Otto et al. 2017; Coronado et al. 2018). Inhomogeneities in the potential due to stars, molecular clouds, black holes, or other dark objects accelerate the disruption of wide binaries (Retterer & King 1982; Bahcall et al. 1985; Carr & Sakellariadou 1999; Yoo et al. 2004). The wide binary separation distribution therefore places constraints on the number and mass distribution of gravitating objects in the disk and halo, and on the dynamical history of the Galaxy.

Because wide binaries can also potentially be disrupted by internal processes during post-main sequence evolution, the separation distribution of binaries containing one or more stellar remnants contains information about the late stages of stellar evolution. Wide binaries can be easily become unbound or “ionized” by $\sim \text{km s}^{-1}$ velocity kicks. As a consequence, the separation distributions of binaries containing a white dwarf places strong constraints on the occurrence of velocity kicks in dying stars. Such kicks have been theoretically predicted to occur as a result of anisotropic and/or non-adiabatic mass loss, but evidence for their existence is largely circumstantial (Fellhauer et al. 2003; Davis et al. 2008; Izzard et al. 2010). Constraining their occurrence is one of the primary goals of this work.

Historically, searches for wide binaries have identified pairs of stars with close angular separations and/or similar proper motions. Because the orbital timescales of wide binaries are too long for changes in positions of the component stars to be observed in real time, such studies rely on probabilistic arguments to identify pairs likely to be bound.

In the absence of astrometric data, the wide binary population can be studied statistically through the stellar angular two-point correlation function (Bahcall & Soneira 1981; Garnavich 1988; Gould et al. 1995; Sesar et al. 2008; Longhitano & Binggeli 2010; Dhital et al. 2015): although it is not possible to determine conclusively whether a given pair is a gravitationally bound binary or a chance alignment, the chance alignment rate can be estimated for stochastically distributed stars (e.g. Lépine & Bongiorno 2007a; Sesar et al. 2008; Dhital et al. 2015), allowing the number of true binaries to be estimated after the expected number of chance alignments at a given angular separation is subtracted.

Wide binaries can be identified more reliably with the aid of proper motion data, which can eliminate interloper pairs that appear close on the sky but have significantly different projected angular velocities. Astrometric searches for wide binaries have primarily targeted

stars with large proper motions (Luyten 1971, 1979b,a, 1988; Wasserman & Weinberg 1991; Poveda et al. 1994; Silvestri et al. 2002; Salim & Gould 2003; Chanamé & Gould 2004; Makarov et al. 2008; Dhital et al. 2010) in order to limit contamination from interlopers that most commonly are distant stars with correspondingly small proper motions.

The purity of a sample of wide binaries can be further improved with the inclusion of parallaxes (e.g. Close et al. 1990; Shaya & Olling 2011; Andrews et al. 2017; Hollands et al. 2018) or radial velocities (e.g. Latham et al. 1984; Close et al. 1990; Price-Whelan et al. 2017; Andrews et al. 2018a,b). Until recently, most wide binary searches that incorporated parallaxes were restricted to modest samples of stars in the immediate solar neighborhood ($d \lesssim 50$ pc). However, the recent *Gaia* data releases (Gaia Collaboration et al. 2016, 2018a) have dramatically expanded the sample of stars for which precise parallaxes and proper motions are available, making it possible to assemble large statistical samples of binaries with a $\ll 1\%$ contamination rate.

In this paper, we use *Gaia* DR2 data to produce a catalog of wide binaries within 200 pc with very high ($\sim 99.8\%$) purity. We divide binaries into main sequence–main sequence (MS/MS), white dwarf–main sequence (WD/MS), and white dwarf–white dwarf (WD/WD) pairs and study the separation distributions of each class of binaries separately. This sample is not volume complete. But the completeness as a function of angular separation can be well characterized (Appendix 4.7), allowing us to correct for observational biases and infer the intrinsic separation distributions.

Because white dwarf binaries are faint and rare, past studies of the WD/WD binary population have been limited to samples of a few dozen objects (e.g. Greenstein 1986; Sion et al. 1991; Andrews et al. 2012; Toonen et al. 2017). A few hundred candidate WD/MS binaries have been identified in common proper motion catalogs (Smith 1997; Silvestri et al. 2002; Chanamé & Gould 2004; Gould & Chanamé 2004), but a systematic study of their separation distribution has yet to be carried out. *Gaia* astrometry makes it possible to assemble a large, homogenous catalog of WD/MS and WD/WD binaries for the first time.

The remainder of this paper is organized as follows. In Section 4.3, we describe our strategy for selecting candidate binaries. Section 4.4 describes our method for extracting the intrinsic separation distributions, which are then presented in Section 4.4.3. These data reveal clear differences in the separation distributions of MS/MS, WD/MS and WD/WD binaries. In Section 4.5, we compare our results to the predictions of binary population synthesis models and use the separation distribution of binaries containing a white dwarf to constrain the occurrence of kicks during white dwarf birth. We summarize our findings in Section 4.6. Appendices 4.7–4.10 provide more information on the details of our model. The catalog of our high-confidence binaries is described in Appendix 4.11.

4.3 Methods

For each star in *Gaia* DR2 within 200 pc that passes our quality cuts (Section 4.3.1), we consider as potential binary companions all neighboring stars that pass similar quality cuts

and lie within an aperture of projected radius 5×10^4 AU. We further require the two stars in a pair to have parallaxes and proper motions consistent with a gravitationally bound binary, as described in Section 4.3.2. Finally, we remove binaries that are members of clusters, moving groups, and higher-order multiples (Section 4.3.3).

We limit our search to projected separations $s < 5 \times 10^4$ AU ≈ 0.25 pc, largely because the contamination rate from background sources (e.g. Lépine & Bongiorno 2007a; Andrews et al. 2017) and unbound associations (e.g. Caballero 2009; Oh et al. 2017a; Dhital et al. 2015) grows rapidly at larger separations. Disruption from external gravitational perturbations is expected to have significant effects only at the largest separations in our catalog ($s \gtrsim 0.1$ pc; Retterer & King 1982; Close et al. 1990; Weinberg et al. 1987; Jiang & Tremaine 2010), so works aiming to use wide binaries as dynamical probes typically focus on binaries with larger separations. As we show in Section 4.5.3, modest velocity kicks during white dwarf formation are expected to ionize binaries with substantially smaller separations.

4.3.1 General quality cuts

We require that both members of a candidate binary pair have a five-parameter astrometric solution and a successful measurement of the $G_{\text{BP}} - G_{\text{RP}}$ color. Following Lindegren et al. (2018) and Gaia Collaboration et al. (2018b), we require both stars to have low astrometric excess noise, satisfying $\sqrt{\chi^2/(\nu' - 5)} < 1.2 \times \max(1, \exp(-0.2(G - 19.5)))$, where χ^2 and ν' are respectively referred to as `astrometric_chi2_al` and `astrometric_n_good_obs_al` in the *Gaia* archive.

We further require the photometry for both stars to be relatively uncontaminated from nearby sources. Because G_{BP} and G_{RP} magnitudes are calculated by integrating over low-resolution spectra, which are more dispersed than the G -band point spread function, the degree of contamination from nearby sources can be assessed by comparing the total BP and RP flux to the G -band flux (Evans et al. 2018). We require both stars to satisfy $1.0 + 0.015(G_{\text{BP}} - G_{\text{RP}})^2 < \text{phot_bp_rp_excess_factor} < 1.3 + 0.06(G_{\text{BP}} - G_{\text{RP}})^2$.

Finally, we require both stars to have high-SNR photometry: $< 2\%$ flux uncertainties in the G -band and $< 5\%$ ($< 10\%$) uncertainties in BP and RP flux, for the primary and companion, respectively. These cuts remove a substantial number of spurious sources that fall in unphysical regions on the color-magnitude diagram, as well as some real sources (see Gaia Collaboration et al. 2018b).

In particular, for pairs with small parallaxes and proper motions – distant and slowly moving stars – it is essentially impossible to conclusively rule out the possibility that a given pair is a chance alignment of unassociated stars. Any search for wide binaries must therefore make compromises between completeness and purity. This work aims to construct a sample with high purity. This is essential for a reliable measurement of the separation distribution at large separations, where true binaries are rare and contaminants can easily dominate.




Figure 4.1: Joint distribution of projected physical separation and distance for all 55,128 binaries in our final sample. Marginalized histograms are on a logarithmic scale, with a dynamic range of 100 to 10,000 counts. Given our photometric and astrometric quality cuts (Section 4.3.1), the typical angular resolution limit of our sample is ~ 2 arcsec. This prevents the detection of close binaries at large distances.




Figure 4.2: On-the-sky physical velocity difference between members of candidate binaries, ΔV_{\perp} , versus their projected physical separation, $s = \theta \times d$. We select binary candidates as objects that fall below the dashed line within 3σ (Equation 4.7); i.e., the difference in their proper motions is small enough to be consistent with a bound Keplerian orbit. Bottom panel shows the subsample with $d < 100 \text{ pc}$.

4.3.2 Initial Search

We search for companions around stars that are nearby and have precise parallaxes, `parallax` > 5 and `parallax_over_error` > 20 . We then search a circle corresponding to a projected radius 5×10^4 AU around each such target, i.e., within an angular separation

$$\frac{\theta}{\text{arcsec}} \leq 50 \times \frac{\varpi}{\text{mas}}, \quad (4.1)$$

where ϖ is the parallax of the target. We also require potential companions to have a reasonably precise parallax, satisfying `parallax_over_error` > 5 . At large distances, one expects the parallaxes of two actual binary components to be essentially identical. However, for the nearest wide binaries, the orbital separation can be a non-negligible fraction of the distance to the binary. Given the high precision of *Gaia* parallaxes for nearby stars, this can cause individual stars in a true binary to have inconsistent parallaxes. For most orbits and orientations, the line-of-sight difference in distance between the two stars in a binary satisfies $\Delta d < 2s$. We reject pairs that do not satisfy this constraint within 3σ , i.e.,

$$\Delta d - 2s \leq 3\sigma_{\Delta d}, \quad (4.2)$$

where in suitable units, $\Delta d = |1/\varpi_1 - 1/\varpi_2|$, and $\sigma_{\Delta d} = (\sigma_{\varpi,1}^2/\varpi_1^4 + \sigma_{\varpi,2}^2/\varpi_2^4)^{1/2}$. Here ϖ_i and $\sigma_{\varpi,i}$ represent the parallax of a target and its standard error. The distribution of sources in our final binary catalog in distance–projected separation space is shown in Figure 4.1. The *Gaia* resolution limit leaves a clear imprint in the distribution, preventing close binaries from being detected at large distances.¹ At large separations, the catalog is dominated by the most distant binaries due to the larger volume at large distances.

We require the difference in proper motion of the two stars to be consistent with a bound Keplerian orbit. For a circular orbit of total mass $5 M_{\odot}$ and semi-major axis a , this translates to a maximum projected physical velocity difference of

$$\frac{\Delta V_{\text{orbit}}}{\text{km s}^{-1}} \leq 2.1 \left(\frac{a}{10^3 \text{AU}} \right)^{-1/2}. \quad (4.3)$$

In terms of observables, we translate this to

$$\frac{\Delta \mu_{\text{orbit}}}{\text{mas yr}^{-1}} \leq 0.44 \left(\frac{\varpi}{\text{mas}} \right)^{3/2} \left(\frac{\theta}{\text{arcsec}} \right)^{-1/2}, \quad (4.4)$$

where we have used the fact that $\Delta V_{\perp} [\text{km s}^{-1}] = 4.74 \times \Delta \mu [\text{mas yr}^{-1}] / \varpi [\text{mas}]$. We do not attempt to account for covariance between parallax and proper motion errors. Equation 4.4 implicitly assumes that the projected separation $s = \theta \times d$, where d is the distance to a

¹*Gaia* DR2 is reasonably complete for separations $\theta \gtrsim 0.7$ arcsec. However, the majority of pairs with $\theta < 2$ arcsec lack $G_{\text{BP}} - G_{\text{BP}}$ colors (see Figure 9 of Arenou et al. 2018a) and are therefore excluded from our catalog.

binary, is equal to the true semimajor axis a . This assumption does not hold in detail due to projection effects and non-circular orbits. But we show in Appendix 4.8 that it holds to within a factor of two in a large majority of cases. Equation 4.3 also does not hold in general for non-circular orbits, as the orbital velocity exceeds $\sqrt{GM/a}$ near periaipse. However, in this case the projected separation is also smaller than a , such that Equation 4.4 holds in the vast majority of cases for randomly oriented orbits (see also Andrews et al. 2017).

We require that all candidate binaries have proper motion differences within 3σ of the maximum velocity difference expected for a system of total mass $5M_{\odot}$ with circular orbits. For each pair, we determine the scalar proper motion difference,

$$\Delta\mu = [(\mu_{\alpha,1}^* - \mu_{\alpha,2}^*)^2 + (\mu_{\delta,1} - \mu_{\delta,2})^2]^{1/2}, \quad (4.5)$$

where $\mu_{\alpha,i}^* \equiv \mu_{\alpha,i} \cos \delta_i$. The corresponding uncertainty $\sigma_{\Delta\mu}$ is given by standard error propagation:

$$\sigma_{\Delta\mu} = \frac{1}{\Delta\mu} \left[\left(\sigma_{\mu_{\alpha,1}^*}^2 + \sigma_{\mu_{\alpha,2}^*}^2 \right) \Delta\mu_{\alpha}^2 + \left(\sigma_{\mu_{\delta,1}}^2 + \sigma_{\mu_{\delta,2}}^2 \right) \Delta\mu_{\delta}^2 \right]^{1/2}, \quad (4.6)$$

where $\Delta\mu_{\alpha}^2 = (\mu_{\alpha,1}^* - \mu_{\alpha,2}^*)^2$ and $\Delta\mu_{\delta}^2 = (\mu_{\delta,1} - \mu_{\delta,2})^2$. We reject pairs with $\sigma_{\Delta\mu} \geq 1.5 \text{ mas yr}^{-1}$. We then require

$$\Delta\mu \leq \Delta\mu_{\text{orbit}} + 3\sigma_{\Delta\mu}, \quad (4.7)$$

where $\Delta\mu_{\text{orbit}}$ is given by Equation 4.4.

The results of this selection are shown in Figure 4.2. The dashed line shows the fiducial cut from Equation 4.4. Pairs with significant proper motion uncertainties scatter above it; this scatter is smaller in the 100 pc sample, where typical proper motion and parallax uncertainties are smaller. To limit contamination from pairs with large uncertainties, we impose an absolute cut of $\Delta\mu \leq 2 \times \Delta\mu_{\text{orbit}}$.

Our selection method compares the on-sky projected proper motions of two stars, rather than the 3d space velocities. At large angular separations, this can potentially lead to a bias in the obtained binary population, because projection effects cause stars with the same 3d space velocity to have different proper motions (e.g. Oh et al. 2017a). We performed simulations to assess the magnitude of any possible biases introduced by this affect. Assuming a uniform spatial distribution of binaries and randomly oriented orbits, we found that within 200 pc, 0.7% of the widest binaries ($s > 10^4$ AU) are expect to have their projected physical velocity difference ΔV_{\perp} inflated by more than 0.25 km s^{-1} due to projection effects, such that they run a substantial risk of not passing the proper motion cut of Equation 4.7. Only binaries with large angular separations ($\theta \gg 10$ arcmin) are affected; so the effect is even smaller for binaries with smaller separations. Within 100 pc, the same fraction affected is 3%.

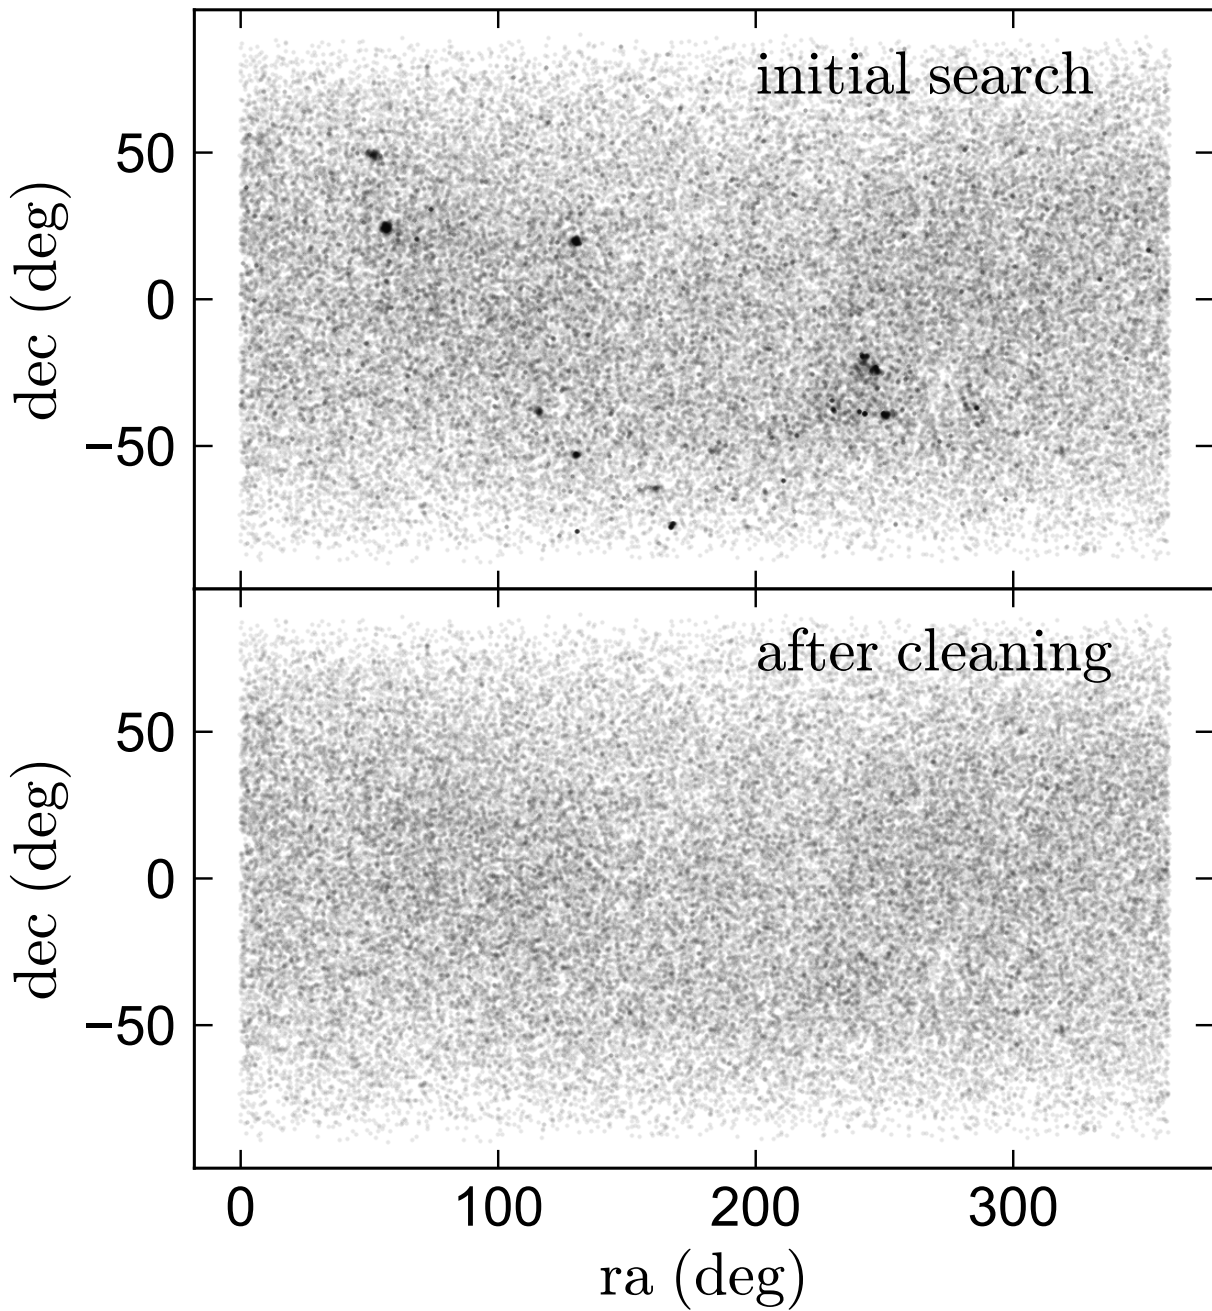


Figure 4.3: **Top:** On-sky distribution of $\sim 60,000$ wide, comoving pairs within 200 pc returned by our initial search. Overdensities correspond to stars in open clusters and moving groups, which meet our initial search criteria but are primarily not genuine bound binaries. **Bottom:** cleaned sample of binaries after $\sim 5,000$ pairs in clusters, moving groups, and resolved higher-order multiples have been removed.




Figure 4.4: Color-magnitude diagram for our final sample of wide binaries. We overplot a solar metallicity PARSEC isochrone and a $0.6 M_{\odot}$, hydrogen atmosphere white dwarf cooling sequence. The vast majority of objects in our catalog are main sequence stars and white dwarfs. In each binary, the “primary” is the object with the larger photometrically-inferred mass (Section 4.3.4).

4.3.3 Removing clusters, moving groups, and higher-order multiples

The on-sky distribution of objects returned by our initial search is shown in the upper panel of Figure 4.3. A number of obvious spatial overdensities are apparent. Most of these can be identified with known open clusters and moving groups within 200 pc. Members of these groups formally satisfy our selection criteria, but most are presumably not genuinely self-bound binaries. Indeed, once-bound associations can remain comoving for several Galactic dynamical times as they slowly dissolve (e.g. Jiang & Tremaine 2010; Shaya & Olling 2011; Oelkers et al. 2017; Oh et al. 2017a), so a substantial number of unbound moving groups are predicted to exist. We clean our sample of such objects through two cuts.

For each candidate binary pair, we count the number of nearby neighbor pairs in position–parallax–proper motion space. We define nearby neighbor pairs as those within 1 degree on the sky, $\pm 3 \text{ mas yr}^{-1}$ in both proper motion coordinates, and $\pm 5 \text{ pc}$ in $1/\varpi$. We remove from our sample pairs for which this search finds more than 5 neighbors. This cut, which was chosen through trial and error, removes ~ 2000 candidate binaries, all of which are in one of the overdensities visible in Figure 4.3.

We also repeat the search described in Section 4.3.2 but replace the proper motion cut from Equation 4.4 with an absolute cut on $\Delta\mu$ corresponding to a physical velocity difference of 2 km s^{-1} . In a majority of cases, this search returns either the same pair found in our initial search, or no results. However, in some cases, it returns many – in some cases, dozens – of additional companions to a given object within 50,000 AU. We discard from our sample all pairs for which this search returns more than one companion to either member of a pair. In addition to selecting the majority of cluster and moving group members selected by the first cut, this cut removes $\sim 2,000$ triples and resolved higher-order multiples in the field. We defer further analysis of these systems to future work.

The bottom panel of Figure 4.3 shows the effects of these cuts on our sample. Obvious clusters and associations are removed. Although not shown in the figure, we also find that obvious structure in proper motion space is removed. Our final sample contains 55,128 binaries. It does not contain any resolved higher-order multiples; i.e., if a star is a member of more than one pair, we remove both pairs.

4.3.4 Summary of the binary catalog

A color-magnitude diagram of all objects in our catalog is shown in Figure 4.4. We assign the parallax of the brighter star to the companion when calculating absolute magnitudes. We define objects with $M_G < 2.75(G_{BP} - G_{RP}) + 5.75$ as “main sequence” stars², where $M_G = G + 5 \log(\varpi/\text{mas}) - 10$. Objects with $M_G > 3.25(G_{BP} - G_{RP}) + 9.63$ are classified as white dwarfs. We do not attempt to correct for extinction, which is minor within 200 pc. 46 objects that passed our other quality cuts were not classified as either main sequence stars or white dwarfs (i.e., they fell between the main sequence and white dwarf cooling sequence);

²i.e., our “main sequence” category serves only to exclude white dwarfs; it includes a few giants.

we discarded the binary candidates containing these objects. Clear sequences of spatially unresolved binaries (e.g. [El-Badry et al. 2018g](#); [Widmark et al. 2018](#)) are apparent above the main sequence for both primaries and secondaries, indicating that quite a number of the wide binaries are really hierarchical systems containing unresolved closer binaries.

Given the color-magnitude diagram, we estimate the mass of each star photometrically by interpolating on a grid of isochrones. We use PARSEC isochrones ([Bressan et al. 2012](#)) for main sequence stars and cooling tracks for carbon-oxygen cores with hydrogen atmospheres ([Holberg & Bergeron 2006](#); [Kowalski & Saumon 2006](#); [Tremblay et al. 2011](#); [Bergeron et al. 2011](#)), in both cases using the revised DR2 filters. These mass estimates are crude and are not included in the public catalog, as we do not make any attempt to correct for unresolved binarity, extinction, or white dwarf spectral types, but in most cases they are expected to be accurate within $0.1 M_{\odot}$, which is sufficient for characterizing the catalog’s binary population in broad strokes.

In [Figure 4.5](#), we show distributions of several observables and inferred properties for each class of binaries. The top panels show the mass of the primary (which we always define to be the more massive star) and the mass ratio. For white dwarfs, the mass refers to the current mass (as opposed to the initial mass), so in some cases, the white dwarf in a WD/MS binary is the secondary. Most main sequence stars in the sample have masses $0.3 \lesssim M/M_{\odot} \lesssim 1.3$; most white dwarfs have $0.5 \lesssim M/M_{\odot} \lesssim 0.8$. The mass ratio distribution for MS/MS binaries is nearly uniform over $0.25 \lesssim q \lesssim 1$ but exhibits an excess of systems with $q \sim 1$, consistent with previous results (e.g. [Söderhjelm 2007](#)). The dearth of low-mass stars and low-mass ratio binaries can largely be attributed to incompleteness; the fall-off at low masses and mass ratios is less pronounced within 50 pc, where the catalog is more complete.

The middle left panel shows the distribution of apparent magnitude difference, $\Delta G = |G_1 - G_2|$, for binaries in each category. In most of the WD/WD binaries, the two stars have similar apparent magnitude (median $\Delta G = 0.66$ mag). The typical magnitude difference is significantly higher in WD/MS binaries (median $\Delta G = 3.50$ mag) and MS/MS binaries (median $\Delta G = 2.05$ mag). This is important for reliable determination of the intrinsic separation distribution, because the *Gaia* sensitivity to a companion at fixed angular separation varies with magnitude difference (see [Appendix 4.7](#)). Corrections for incompleteness at small separations thus affect the three populations differently.

[Figure 4.5](#) also shows the distributions of distance (middle right) and angular separation (bottom left). Because white dwarfs are intrinsically fainter than most main sequence stars, those in our catalog are on average nearer. The median distances for WD/WD, WD/MS, and MS/MS binaries in our catalog are 102 pc, 127 pc, and 147 pc, respectively. The typical angular separation for all binaries is of order 10 arcsec. Due to incompleteness at small angular separations ([Figure 4.1](#)), there are essentially no binaries with angular separations of less than 1.5 arcsec. A few of the nearest binaries have separations of tens of arcminutes, but these are rare; the 99th percentile separation for all binaries in our catalog is 4 arcmin.

Our requirement that stars in the catalog have precise parallaxes, proper motions, and photometry causes it to contain primarily relatively bright stars (lower right). The median apparent magnitude of MS stars in the catalog is $G = 14.5$; for white dwarfs, it is $G = 18.4$.

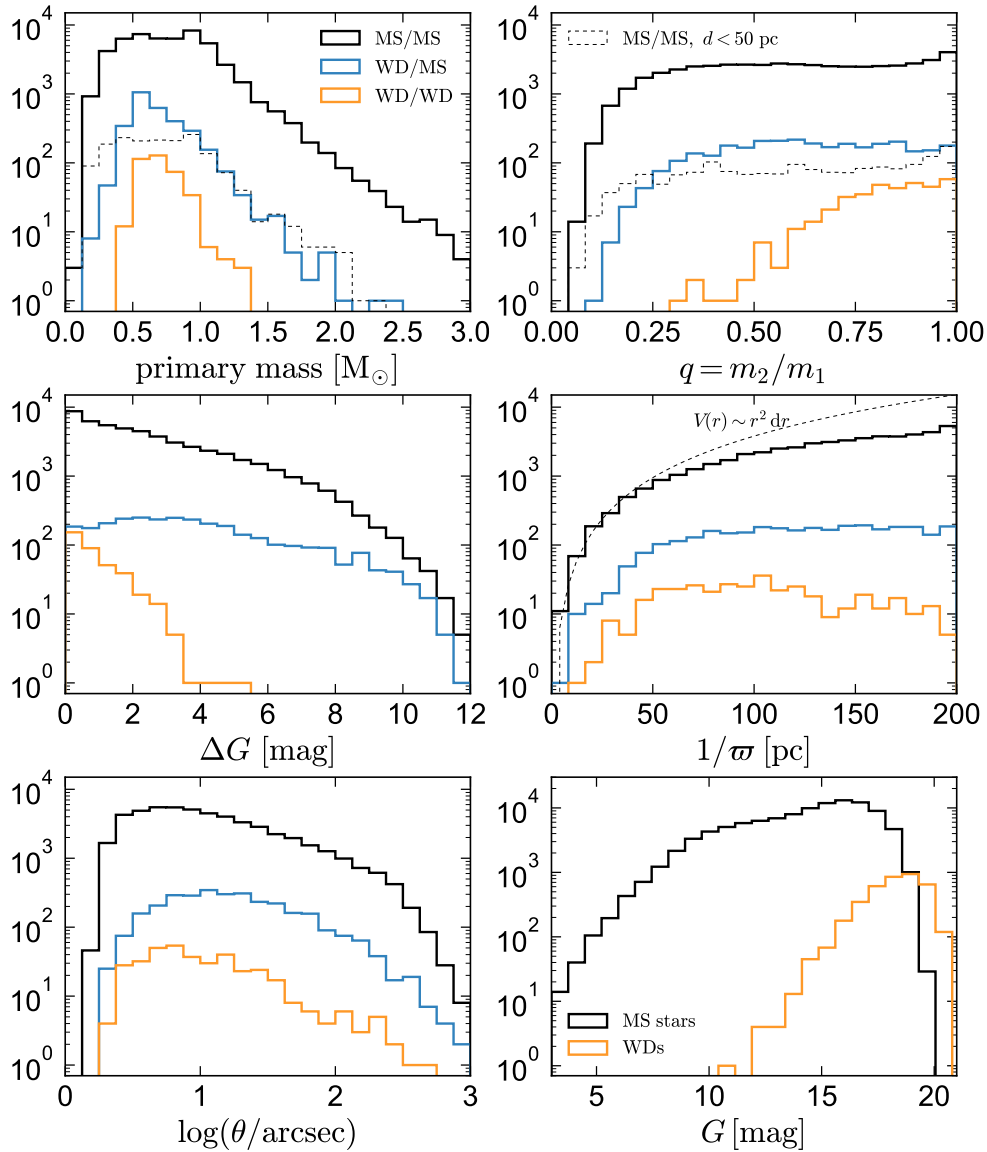


Figure 4.5: Distributions of primary mass, mass ratio, magnitude difference, distance, angular separation, and apparent magnitude for the 51,668 MS/MS, 3,085 WD/MS, and 375 WD/WD wide binaries in our catalog. Most of the main-sequence stars have masses $0.3 \lesssim M/M_\odot \lesssim 1.3$; most of the white dwarfs have $0.5 \lesssim M/M_\odot \lesssim 0.8$. The mass ratio distribution is roughly uniform for MS/MS binaries. The typical brightness contrast between the two stars in a binary is largest for MS/MS binaries and smallest for WD/WD binaries. Dashed line in the middle right panel shows the scaling expected for a uniform spatial distribution of binaries, illustrating the effects of incompleteness at large distances. The typical spatial angular separation is ~ 10 arcsec, with the widest binaries in the catalog having separations of tens of arcminutes.




Figure 4.6: Left: Heliocentric velocities for the 7.5% of MS/MS binaries in our catalog for which *Gaia* provides a precise radial velocity for both components. Radial velocities were not used in our selection of candidate binaries but provide a useful check on the contamination rate: in genuine binaries, the RVs of the two stars should agree within a few km s^{-1} (modulo radial velocity uncertainties of up to 3 km s^{-1}). **Right:** Most binaries in our catalog with substantially different radial velocities have larger-than-average radial velocity errors; the fraction of binaries with inconsistent radial velocities does not vary much with orbital separation. We estimate the contamination rate to be $\lesssim 0.2\%$.

Typical parallax and proper motion uncertainties are consequently also higher for white dwarfs.

Higher-order multiples

Although we have removed resolved higher-order multiples from the catalog, it is still expected to contain a substantial number of hierarchical systems in which one or both of the resolved components is really a spatially unresolved close binary. Unresolved binaries with mass ratios $q \gtrsim 0.5$ scatter above the main sequence in a color-magnitude diagram (e.g. Hurley & Tout 1998; Widmark et al. 2018). For solar-type stars, the “binary sequence” formed by such objects barely overlaps with the single-star main sequence, making it straightforward to estimate the binary fraction from the fraction of objects above the main sequence.

The unresolved binary sequence is most cleanly separated from the single-star main sequence at $1 \lesssim (G_{\text{BP}} - G_{\text{RP}}) \lesssim 2$ (Figure 4.4), so we use this region of the CMD to estimate the unresolved binary fraction for MS/MS binaries in our catalog. Main-sequence stars in this region of the CMD have masses $0.5 \lesssim m/M_{\odot} \lesssim 1$. For $1 < (G_{\text{BP}} - G_{\text{RP}}) < 2$, the line $M_G = 2.8(G_{\text{BP}} - G_{\text{RP}}) + 2.4$ divides the binary and single-star main sequences. We find that $\approx 10.5\%$ of primaries and secondaries with $1 < (G_{\text{BP}} - G_{\text{RP}}) < 2$ fall above this line and likely have an unresolved companion with $q \gtrsim 0.5$. Assuming a uniform mass-ratio distribution for unresolved binaries, this implies that $\sim 20\%$ of main-sequence primaries and secondaries have an unresolved main-sequence companion, such that at least one component has an unresolved main-sequence companion in $\sim 36\%$ of the wide binaries in our catalog. Although they cannot be easily identified in an optical CMD, we also estimate that a few percent of the main-sequence components in the catalog have an unresolved white dwarf companion (Willems & Kolb 2004).

4.3.5 Contamination Rate

Our search for wide binaries in *Gaia* DR2 uses only parallaxes, spatial coordinates, and proper motions to determine whether two stars are part of a binary system. We do not use *Gaia* radial velocities (RVs) because they are not available for most of our candidates. However, reliable RVs *are* available for both stars in 7.5% of the MS/MS binaries in our sample. We use these to validate our catalog and assess the contamination rate. RVs provide an independent check on whether candidate pairs are true binaries or chance alignments; in the latter case, one would expect the RVs of the two stars to be significantly different, with a typical RV difference comparable to the local velocity dispersion of the disk (e.g. Andrews et al. 2018a).

We consider all pairs in which a *Gaia* RV with uncertainty $\sigma_{\text{RV}} < 3 \text{ km s}^{-1}$ is available for both stars and a *Gaia*-RVS spectrum was obtained for at least 3 transits (see Sartoretti et al. 2018; Katz et al. 2018); σ_{RV} is calculated from the scatter in the RVs measured for individual transits. This cut eliminates most objects with bad RVs and yields 3813 pairs. All pairs with RVs are MS/MS, as DR2 does not include RVs for white dwarfs. In the left

panel of Figure 4.6, we compare the RVs of the two stars in each binary. As expected for true wide binaries, the vast majority of points lie within a few km s^{-1} of the one-to-one line. A few binaries do scatter far from the one-to-one line; these are potential contaminants.

In the right panel of Figure 4.6, we show the absolute RV difference, ΔRV , as a function of physical separation and the uncertainty in the RV difference. The majority of binaries with larger-than-average ΔRV also have larger-than-average $\sigma_{\Delta\text{RV}}$, suggesting that they are true binaries. There are 12 binaries with RV differences in excess of 10 km s^{-1} . Most of these also have relatively large $\sigma_{\Delta\text{RV}}$, such that their velocities are marginally consistent with a bound orbit. 6 binaries have $\Delta\text{RV} > 15 \text{ km s}^{-1}$. We regard these as likely contaminants, assuming no catastrophic failure in the RVs. This translates to a contamination rate of $\sim 0.15\%$. We note that one might expect the contamination rate for the full sample to be somewhat higher than that in the subsample of pairs with RVs, because the stars with measured RVs are brighter on average and thus have smaller-than-average parallax and proper motion uncertainties.

We obtain an independent estimate of the contamination rate by calculating the number of random alignment pairs expected to pass our selection criteria. We estimate this by applying the same selection criteria and quality cuts to the mock *Gaia* DR2 catalog presented in Rybizki et al. (2018), which implements the Besançon model of stellar population synthesis (Robin et al. 2003) and populates the Galactic distribution function using *Galaxia* (Sharma et al. 2011), assuming a similar selection function and uncertainty model to *Gaia* DR2. The mock catalog does not contain any true binaries, so *any* pairs that pass our selection criteria in the mock catalog are chance alignments. It also does not contain any white dwarfs, so we restrict our comparison to the MS/MS pairs.

Executing the same search on the mock catalog, we find 51 pairs. Given a total sample of $\sim 50,000$ MS/MS binaries, this implies a contamination rate of 0.1%. The majority (85%) of chance alignments found in the mock catalog have $s > 10^4 \text{ AU}$. This is expected for chance alignments simply because there is more area at large separations than at small separations. Intriguingly, the binary candidates in our catalog that we identify as likely contaminants because they have large ΔRV are not strongly concentrated at large s (Figure 4.6). This suggests that some of the pairs with large ΔRV may in fact be genuine binaries in which one of the objects has an unreliable radial velocity. One possibility is that some of these objects are in fact hierarchical triples in which one component is a close, unresolved binary. Most of these systems should be eliminated by the cut of $\sigma_{\text{RV}} < 3 \text{ km s}^{-1}$, but a few could escape detection if the different transits occurred at a similar orbital phase. In any case, these tests imply that our sample retains high purity at large separations; for example, the contaminant population estimated from the Rybizki et al. (2018) mock catalog implies a contamination rate of less than 1% even at $s > 10^4 \text{ AU}$.




Figure 4.7: Normalized distributions of projected physical separation for MS/MS binaries (black), WD/MS binaries (blue) and WD/WD binaries (gold). Top panel shows full sample within 200 pc; bottom panel shows objects within 100 pc. The dearth of systems with small separations reflects the fact that binaries with small angular separations are spatially unresolved, leading to incompleteness in our sample at $\log(s/\text{AU}) \lesssim 3$ in the top panel and $\log(s/\text{AU}) \lesssim 2.7$ in the bottom panel (dashed lines; see Appendix 4.7). The separation distribution of MS/MS binaries falls off as $dN/ds \sim s^{-1.6}$ (where Opik’s law is $dN/ds \sim s^{-1}$), at least for $s \gtrsim 500 \text{ AU}$. At large separations, the separation distributions of WD/WD and WD/MS binaries fall off more sharply than that of MS/MS binaries.

4.4 Separation distributions

Figure 4.7 shows raw distributions of the projected physical separation, $s = \theta \times (1/\varpi)$, for the three classes of binaries. The scarcity of binaries with small physical separations is due to incompleteness at small angular separations, where binaries are unresolved. At larger s , all binaries are resolved, so the slope of the raw separation distribution is representative of that of the intrinsic distribution. As we show in Appendix 4.7, within 200 pc the separation distribution is complete for $\Delta G < 5$ (84% of all binaries) at $s > 1000$ AU and is complete for all ΔG at $s > 2000$ AU. Considering only the 100 pc sample moves the completeness limit toward smaller separations by a factor of two.

Incompleteness effects become severe at $s \ll 1000$ AU (full sample) and $s \ll 500$ AU (100 pc sample), so we caution against overinterpretation of the raw distributions in this regime. Incompleteness affects the separation distribution of each class of binaries in different ways because they have different characteristic magnitude differences (see Figure 4.5 and Appendix 4.7). We describe how incompleteness at small orbital separations can be accounted for in Section 4.4.1.

The raw separation distribution of MS/MS binaries in Figure 4.7 clearly continues to rise toward smaller separations, to at least 500 AU. Although the distribution appears to flatten somewhat at smaller separations, it remains unambiguously steeper than Opik’s law (i.e., a flat distribution of $\log(s)$). This result is in conflict with the findings of some previous works (e.g. Poveda & Allen 2004; Chanamé & Gould 2004; Lépine & Bongiorno 2007a; Sesar et al. 2008), which have found the binary separation distribution to be consistent with Opik’s law out to a break at ~ 3000 AU, with steeper fall-off beginning only at larger separations. The existence of such a break has been interpreted theoretically as a consequence of dynamical effects within the clusters in which wide binaries are formed (Gould & Eastman 2006). Our distribution does not display any strong break. We deem it unlikely that this discrepancy is the result of unaccounted-for systematics in our catalog, as any incompleteness would be expected to introduce biases against *smaller* separations.

4.4.1 Inferring intrinsic separation distributions

Due to our quality cuts and the *Gaia* completeness limit, our catalog is not volume-complete: it is missing both spatially unresolved close binaries, and binaries in which either star is too faint to pass our quality cuts or be detected in the first place. Correcting for incompleteness due to unresolved binaries with small separations is important, as it leads to biases in the separation distribution. On the other hand, incompleteness arising from the fact that faint stars are not detected is only a problem if it preferentially affects stars with small or large separations.

Whether a companion is detected at fixed angular separation is a function primarily of its apparent magnitude, of the magnitude contrast between the two stars, and of the angular separation of the two stars. At sufficiently large angular separations, a binary will be detected as long as both stars are bright enough to be detected independently. However, at small

angular separations, the presence of a bright star makes it harder to detect a faint companion at fixed angular separation (see also Ziegler et al. 2018). This is especially true given the requirement that both stars be free from contamination from nearby sources (Section 4.3.1).

We characterize the probability of a companion’s detection at an angular separation θ with a function $f_{\Delta G}(\theta)$ which depends on ΔG , the G -band magnitude difference between the two stars. We empirically constrain $f_{\Delta G}(\theta)$ as described in Appendix 4.7. As one might expect, $f_{\Delta G}(\theta)$ drops rapidly from 1 at large angular separations³ to 0 at small separations; the characteristic angular scale at which the drop occurs increases with ΔG .

We consider an intrinsic separation distribution with functional form $\phi(s|\vec{m}) = dP/ds$, where \vec{m} is a set of free model parameters to be fit. Given a sample of binaries with projected separations s_i , the likelihood function is

$$L = p(\{s_i\}|\vec{m}) = \prod_i p(s_i|\vec{m}), \quad (4.8)$$

where $p(s_i|\vec{m})$ is the probability of detecting the i -th binary given model parameters \vec{m} . It can be calculated for each binary as

$$p(s_i|\vec{m}) = \frac{\phi(s_i|\vec{m})}{\int_{s_{\min}}^{s_{\max}} \phi(s|\vec{m}) f_{\Delta G}(s|d_i) ds}. \quad (4.9)$$

p_i is proportional to the probability that a binary with distance d_i , magnitude difference ΔG , and physical separation s_i is found in the catalog, given an intrinsic separation distribution with parameters \vec{m} . $f_{\Delta G}$ is explicitly a function of θ , the angular separation of the two stars (Equation 4.15); it is given by $f_{\Delta G}(\theta) = f_{\Delta G}(s/d_i)$. Here s_{\min} and s_{\max} are the minimum and maximum separations of the distribution from which the observed s_i are drawn, and $\phi(s|\vec{m})$ is normalized such that $\int_{s_{\min}}^{s_{\max}} \phi(s|\vec{m}) ds = 1$. We note that Equation 4.9 does not account for the observational uncertainties in s_i or d_i . These are expected to be small, as all the binaries in our sample have parallax errors smaller than 5%.

The denominator in Equation 4.9 represents the fraction of predicted binaries that could have been detected at a distance d_i and magnitude difference ΔG ; it accounts for the fact that at large distances and large ΔG , only binaries with large s can be detected. We set $s_{\max} = 5 \times 10^4$ AU. The choice of s_{\min} has no effect on our results, since the integrand in Equation 4.9 goes to 0 at small separations. We set $s_{\min} = 10^{-2}$ AU.

Whether a binary is detected also depends on the apparent magnitude of both stars, as systems in which either star is too faint will not be detected. This has no effect on the inferred separation distribution as long as the undetected binaries have the same intrinsic separation distribution – for a particular class of binaries – as those that are detected. This implies that the separation distribution for each class of binaries must be independent of both distance and the absolute magnitude of the two stars. That $\phi(s)$ does not vary significantly with distance is reasonable within our 200 pc sample, but the separation distribution *is* expected

³I.e., $f_{\Delta G}(\theta)$ is normalized relative to the fraction of companions that would be detected at asymptotically large separations, not the absolute number of companions that exist.

vary somewhat with the mass (and thus, absolute magnitude) of the two stars (e.g. [Duchêne & Kraus 2013](#); [Moe & Di Stefano 2017a](#)). Our inferred separation distribution should thus be viewed as a marginalization over the mass distributions in our sample, which consists primarily of solar and sub-solar mass stars ($0.3 \lesssim M/M_\odot \lesssim 1.3$; [Figure 4.5](#)).

4.4.2 Fitting the separation distribution

We model the separation distribution as a broken power-law:

$$\phi(s) = \phi_0 \begin{cases} s^{-\gamma_1}, & s \leq s_{\text{break}} \\ s_{\text{break}}^{\gamma_2 - \gamma_1} \times s^{-\gamma_2}, & s > s_{\text{break}} \end{cases}, \quad (4.10)$$

where ϕ_0 is a normalization constant. The break separation, s_{break} , where the distribution transitions from $\phi(s) \sim s^{-\gamma_1}$ to $\phi(s) \sim s^{-\gamma_2}$, is left free, so we fit for three parameters $\vec{m} = (\gamma_1, \gamma_2, \log(s_{\text{break}}/\text{AU}))$. We experimented with using a more flexible parameterization with more than one break in the power law but found that it did not improve the fit much. For a power-law parameterization of $\phi(s)$ and our adopted fitting function for $f_{\Delta G}(\theta)$ ([Equation 4.15](#)), the integral in [Equation 4.9](#) has an analytic solution in terms of hypergeometric functions.

We sample the posterior distribution for each class of binaries using `emcee` ([Foreman-Mackey et al. 2013b](#)). We use broad, flat priors on all model parameters.

4.4.3 Fitting results

We show the constraints obtained for the separation distribution of each class of binaries in [Figure 4.8](#). Marginalized parameter constraints are visualized using `corner` ([Foreman-Mackey 2016](#)). Remarkably, these distributions reveal for the first time distinct differences in the separation distributions of wide MS/MS, WD/MS, and WD/WD binaries.

These differences are already hinted at in the raw distributions, but become more obvious in the inferred intrinsic distributions, which account for the different angular completeness of the three classes of binaries. Consistent with the apparent trends in [Figure 4.7](#), the separation distribution of WD/WD binaries is almost flat out to $s \sim 1,500$ AU, after which it drops off sharply. The distribution for WD/MS binaries also exhibits a break, but at slightly larger separations. In contrast, the MS/MS separation distribution is best fit by a model with a weak break, formally at $\log(s/\text{AU}) = 3.8$. But the change in power law slope is modest (see also the raw distribution in [Figure 4.7](#)), and comparably good fits for MS/MS binaries can be obtained with a single power law model.

Although there is no strong evidence for a turnover in the separation distributions at small separations in our sample ([Figure 4.7](#)), we emphasize that the true uncertainty in the separation distribution at small s is larger than the shaded regions in the top panel of [Figure 4.8](#), particularly for MS/MS binaries. As [Figures 4.1](#) and [4.2](#) show, our catalog contains few binaries with small physical separations. The large number of binaries with

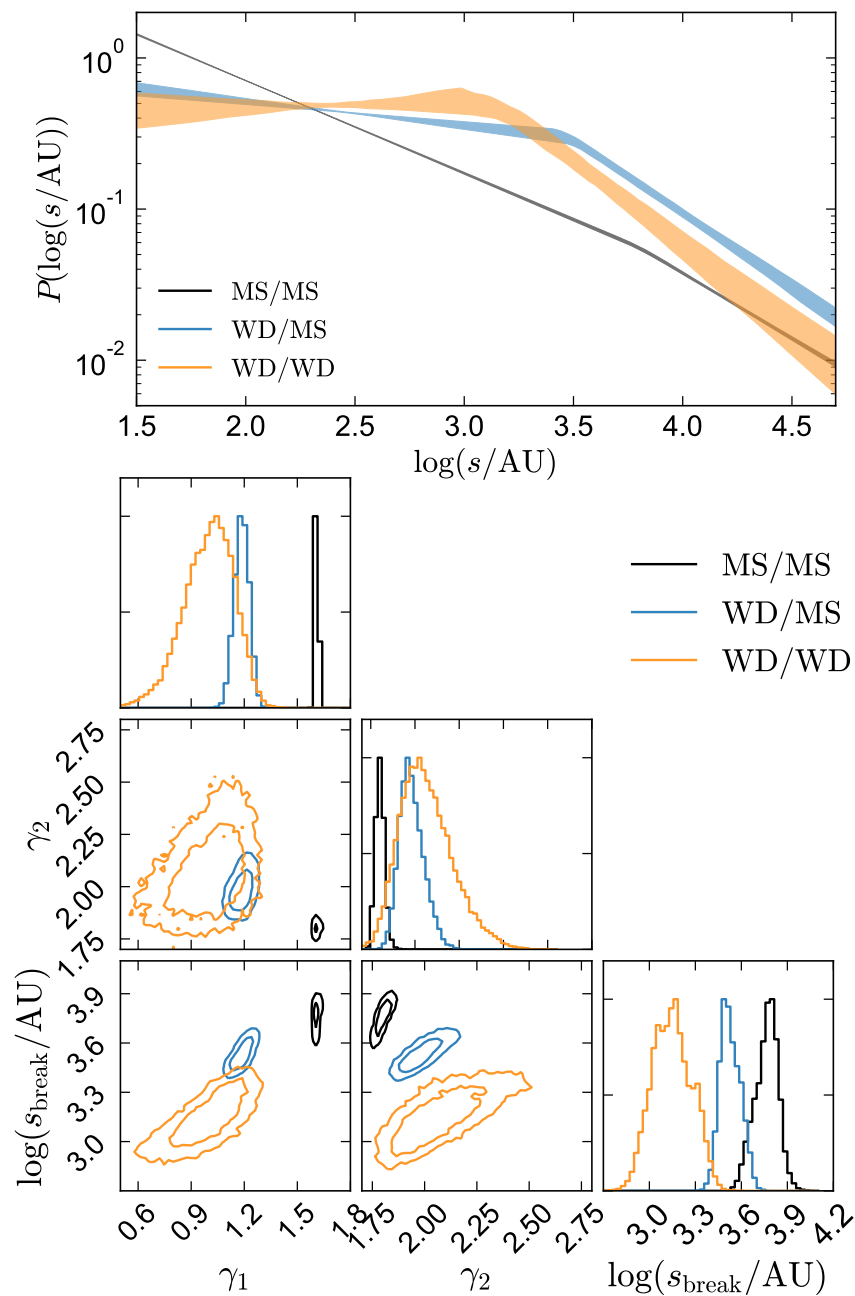


Figure 4.8: Top: Broken power-law fits to the intrinsic separation distributions of MS/MS, WD/MS, and WD/WD binaries, after accounting for incompleteness and selection effects. **Bottom:** 68 and 95% probability contours and marginalized probability distributions for parameters of broken power-law fits. The separation distributions of binaries containing a white dwarf are significantly different from the MS/MS binary separation distribution: they fall off more steeply at large separations and exhibit a stronger break at intermediate separations.

separations $2.5 \lesssim \log(s/\text{AU}) \lesssim 3.5$ set tight constraints on the power-law slope in that range; the slopes of all posterior samples remain tightly constrained at smaller separations because only a single break in the power law is allowed by the model. We find that we obtain comparably good fits to the separation distribution of MS/MS binaries when we fit a lognormal separation distribution that is allowed to flatten at $\log(s/\text{AU}) \lesssim 2.5$ (see Appendix 4.9). Moreover, when we fit a power law model with an extra break at $\log(s/\text{AU}) = 2.5$, constraints on the slope at smaller separations become weaker. On the other hand, a steepening in the slope of the separation distributions of WD/MS and WD/WD binaries at intermediate separations is a robust prediction; it is evident in the raw distributions and is recovered for all models that allow a change in the slope of the distribution.

In Figure 4.9, we divide the MS/MS sample by mass at a primary mass of $0.8 M_{\odot}$ and show separately the inferred separation distributions of MS/MS binaries with higher and lower masses. This cut results in roughly equal numbers of binaries in the two subsamples and is comparable to the minimum mass of a single star that can evolve to become a white dwarf within the age of the Universe.

Consistent with previous work (e.g. Duchêne & Kraus 2013; Moe & Di Stefano 2017a), we find that the separation distribution for MS/MS binaries varies somewhat with mass: the distribution for lower-mass systems falls off more strongly at large separations than that for higher-mass systems. Because the progenitors of white dwarfs in our catalog would fall in the higher-mass subsample, variation in the separation distribution with mass should be considered in interpreting the separation distributions of WD/MS and WD/WD binaries. We note that the trend of higher-mass binaries having wider separation distributions is the opposite of what would be expected if the different separation distributions for binaries containing a white dwarf were due only to the higher average masses of white dwarf progenitors.

In Appendix 4.9, we verify that our approach yields comparable constraints on the separation distribution when we fit the full catalog and when we fit only the 100 pc sample, which has a different raw separation distribution (Figure 4.7), completeness, and typical parallax and proper motion uncertainties. This suggests that our inference does a reasonable job of extracting the intrinsic separation distribution without large biases.

4.5 Theoretical interpretation

Our discovery that the separation distribution of WD/WD (WD/MS) binaries exhibits a break at ~ 1500 AU (~ 3000 AU) that is not seen in the distribution of MS/MS binaries calls for an explanation, which we develop now.

4.5.1 Semi-analytic binary evolution models

The wide binary separation distribution depends on the intrinsic initial separation distribution, on internal processes that can change the separation of a binary after it forms (e.g. mass transfer, tidal effects, mass loss, or velocity kicks), and on disruption due to

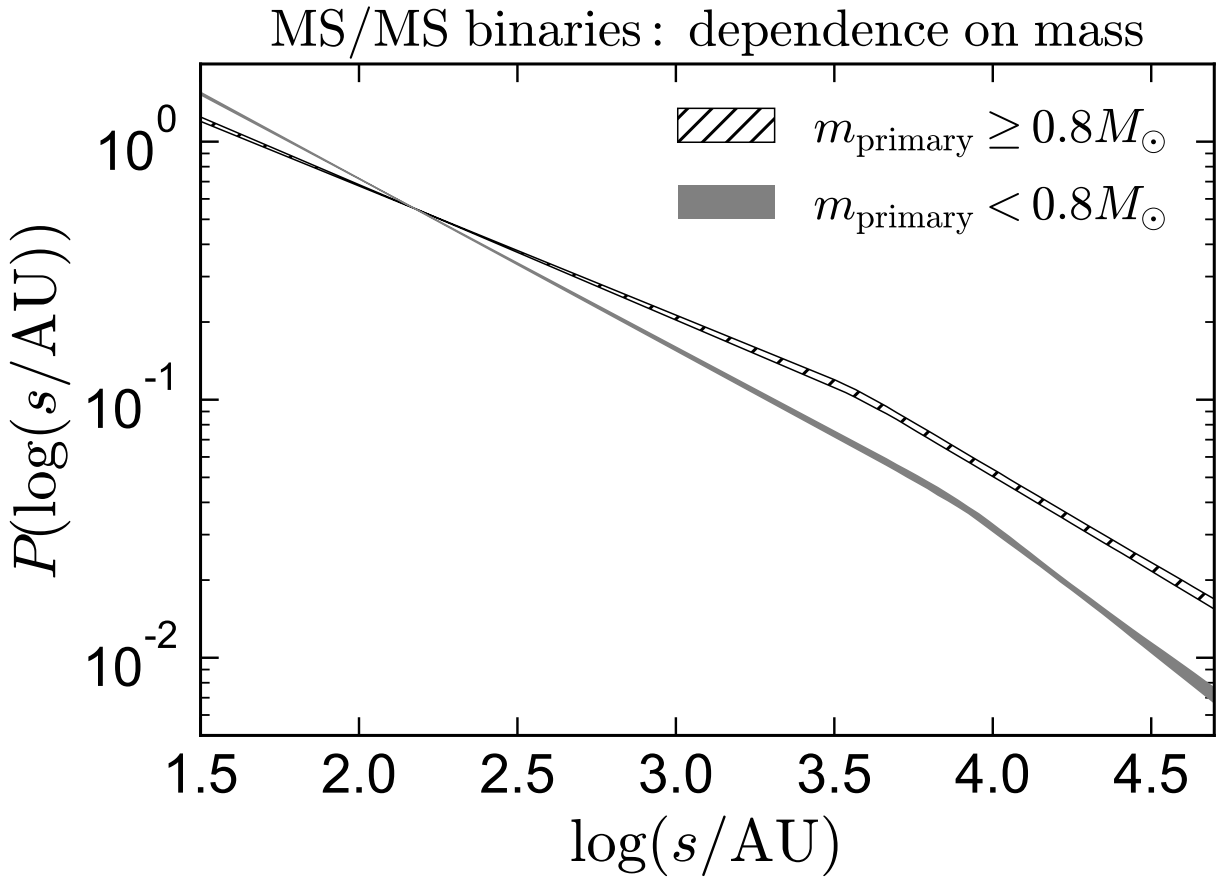


Figure 4.9: Constraints on the separation distribution of MS/MS binaries for systems with higher-mass (hatched) and lower-mass (shaded) primaries. The separation distribution of lower-mass systems falls off more strongly at large separations, consistent with results from previous work.

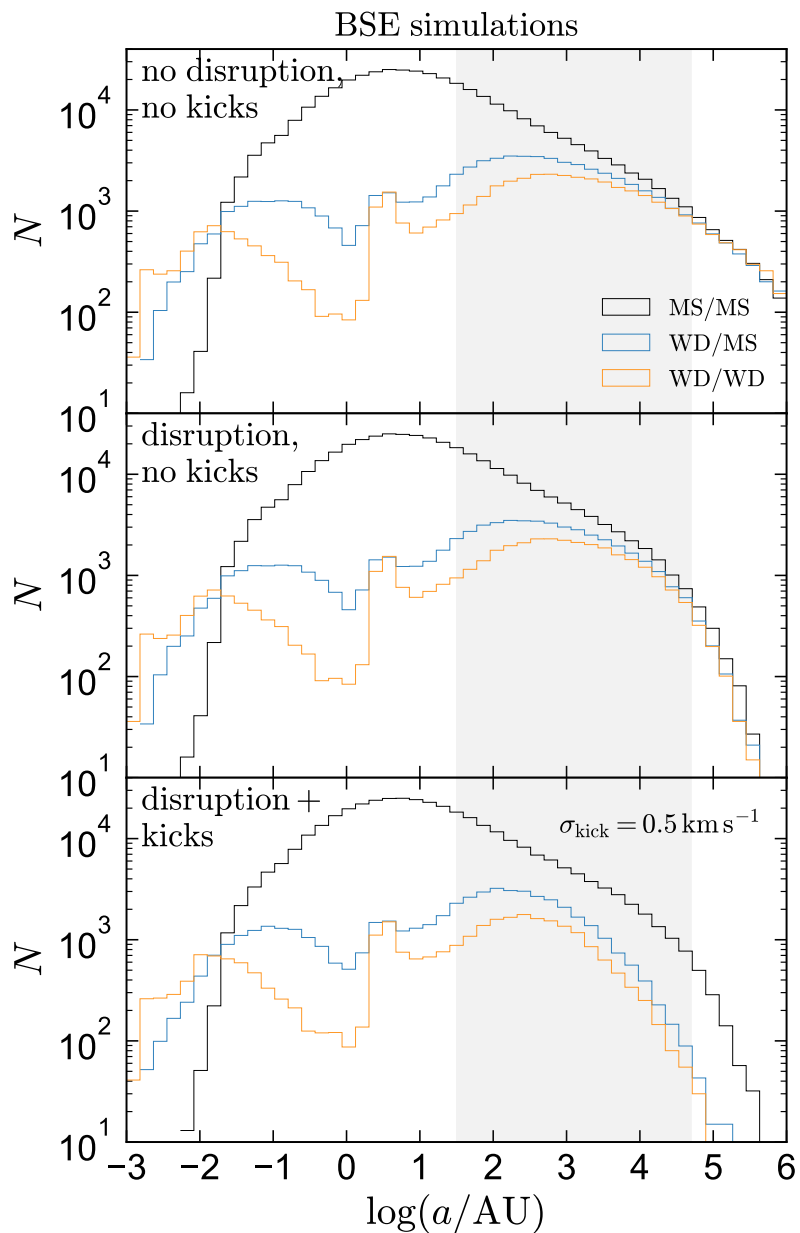


Figure 4.10: Full semi-major axis distributions predicted by our binary population synthesis simulations. Shaded region shows the range of separations probed by our data. Top panel shows the raw population produced by BSE; middle panel shows the results of accounting for the disruption of wide binaries due to external gravitational perturbations (Equation 4.11). In the bottom panel, a modest velocity kick is applied when white dwarfs form due to asymmetric mass loss during post-MS evolution (Section 4.5.3); this unbinds most binaries containing a proto-WD with $\log(a/\text{AU}) \gtrsim 3.5$ and also expands the orbits of binaries with smaller initial separations.

gravitational perturbations from other objects. To quantitatively assess the effects of these processes on the observable separation distributions, we carry out Monte Carlo simulations of a large population of synthetic binaries using the semi-analytic model BSE (“binary star evolution”; Hurley et al. 2002b).

We generate 5×10^5 binaries assuming a constant star formation history over 12 Gyr. We draw masses from a Kroupa (2001a) initial mass function between $0.3 M_{\odot}$ and $8 M_{\odot}$ and mass ratios from a uniform distribution over $[0, 1]$, roughly consistent with the photometrically-inferred mass distributions for our sample (Figure 4.5). For systems where the primary mass is greater than $0.75 M_{\odot}$, we draw orbital periods from a lognormal distribution with $\overline{\log(P/\text{day})} = 4.8$ and $\sigma_{\log(P/\text{day})} = 2.3$ (Duquennoy & Mayor 1991); for systems with lower mass primaries, we draw from a second lognormal with $\overline{\log(P/\text{day})} = 4.1$ and $\sigma_{\log(P/\text{day})} = 1.3$ (Fischer & Marcy 1992). Eccentricities are drawn from a uniform distribution over $[0, 0.95]$ (Duchêne & Kraus 2013). We use BSE to evolve all binaries from the zero-age main sequence to present day. We then divide the simulated binary population into MS/MS, WD/MS, WD/WD, and other classes of binaries according to the evolutionary stage reported by BSE. We do not include hierarchical systems containing unresolved binaries in these simulations.

The top panel of Figure 4.10 shows the semi-major axis distribution predicted for the three classes of binaries we consider, assuming default model parameters (see Table 3 of Hurley et al. 2002b). The distribution for MS/MS binaries is similar to the initial input distribution. The separation distributions of binaries containing a white dwarf are bimodal, with a shortage of systems with $a \sim 1$ AU. Binaries with initial separations $a \lesssim 10$ AU interact when their component stars leave the main sequence, and common envelope effects shrink many of their orbits to $a \ll 1$ AU (e.g. Paczynski 1976; Tout & Hall 1991; Ivanova et al. 2013).

At larger separations, the components of systems with modest eccentricities do not interact, and the orbital separation *grows* when each component leaves the main sequence. In this case – if mass loss is adiabatic and isotropic – the orbit gradually expands such that the quantity $(m_1 + m_2)a$ is conserved (Jeans 1924; Rahoma et al. 2009). This leads the orbits of typical MS/MS binaries to expand by a factor of ~ 2 on average when one star becomes a white dwarf, and by another factor of ~ 2 when the second star becomes a white dwarf. These processes make the separation distributions of binaries containing a white dwarf flatter than the MS/MS binary separation distribution in the range of separations probed by our data (top panel of Figure 4.10).

However, we have not yet accounted for processes that can disrupt the orbits of wide binaries.

4.5.2 Disruption from external perturbations

All wide binaries can be disrupted over time through gravitational interactions with inhomogeneities in the Galactic gravitational potential caused by stars, molecular clouds, or other dark objects (Heggie 1975; Bahcall et al. 1985; Weinberg et al. 1987; Mallada & Fernandez 2001). For the most weakly bound binaries, the large-scale Galactic tidal field

can also accelerate disruption (Jiang & Tremaine 2010).

For the range of separations probed by our data ($a \lesssim 0.25$ pc), the effects of the Galactic tidal field are negligible. Disruption due to interactions with stars and other objects can be modeled as a diffusive process described by the Fokker-Planck equation. Following Andrews et al. (2012), we estimate the timescale $t_{1/2}$ at which a binary of semi-major axis a has a 50% survival probability as

$$t_{1/2}(a) \approx 29 \text{ Gyr} \left(\frac{a}{10^4 \text{ AU}} \right)^{-1}. \quad (4.11)$$

This approximation is based on calculations by Weinberg et al. (1987). We apply this approximation to our Monte Carlo binary population statistically: the probability of a binary’s survival to the present day is set as

$$P_{\text{survive}}(a, \text{age}) = 2^{-\text{age}/t_{1/2}(a)}. \quad (4.12)$$

In practice, this prescription causes about half of all binaries with $a \sim 5 \times 10^4$ AU to be disrupted. Most binaries with $a > 10^5$ AU are disrupted, while most with $a < 10^4$ AU survive. The middle panel of Figure 4.10 shows the effects of disruption on the observable binary population. Disruption affects WD/MS and WD/WD binaries slightly more than MS/MS binaries at fixed semi-major axis because they are older on average, but the difference is minor. This means that the separation distribution of MS/MS binaries can straightforwardly constrain the efficiency of disruption from external perturbations, without potential complications due to the effects of mass loss on binary orbits.

4.5.3 White dwarf recoil

We now assess how *internal* perturbations arising from kicks to the component stars during post-main sequence evolution are expected to change the observable separation distribution.

In the calculations described above, BSE assumes that mass loss on the AGB occurs isotropically and quasisteadily over long timescales, such that mass loss leads to gradual orbital expansion. This result will break down if mass loss occurs asymmetrically or too rapidly. If mass loss is asymmetric, the proto-white dwarf will recoil to conserve momentum, resulting in a velocity “kick” near the end of the TP-AGB phase or as it forms a planetary nebula. Given typical wind velocities of $(10 - 20) \text{ km s}^{-1}$, even 1% deviations from spherical symmetry in the integrated mass-loss history can produce a recoil velocity of order a km s^{-1} (Fellhauer et al. 2003). Resolved observations of AGB stars find tentative evidence for anisotropic mass loss (Zijlstra 2006; Blasius et al. 2012; Höfner & Olofsson 2018), and some asymmetric mass loss is likely required to explain the observed rotation rates of single white dwarfs (Spruit 1998).

Although the proposed kicks are too weak to detect directly in observations of proto-

white dwarfs in planetary nebulae (PNe),⁴ several circumstantial lines of evidence hint at their existence. First, the number of observed white dwarfs in nearby open clusters is lower than is predicted by stellar evolution models for a standard initial mass function, initial mass-age relation, and initial-final mass relation (e.g. Weidemann 1977; Kalirai et al. 2001); this could be explained if the typical kick speed is comparable to the clusters' few-km s⁻¹ escape velocity (Fellhauer et al. 2003). Second, the velocity dispersion of young white dwarfs in globular clusters has been found to be higher than predicted in the absence of natal kicks (Davis et al. 2006, 2008; Calamida et al. 2008). A typical kick velocity of order 1-2 km s⁻¹ could explain these observations (Heyl 2007, 2008b,a), and could also play a role in staving off core collapse in globular clusters (Heyl & Penrice 2009; Fregeau et al. 2009). Few-km s⁻¹ velocity kicks during white dwarf formation could also explain the nonzero observed eccentricities of binaries containing a white dwarf and a barium star (Izzard et al. 2010), as well as the high occurrence rate of dusty disks around young white dwarfs (Stone et al. 2015).

The separation distributions of WD/MS and WD/WD binaries provide a sensitive probe of white dwarf kicks, as even a small kick is sufficient to unbind a weakly-bound wide binary. For example, for two stars of equal mass M in a circular orbit, the kick velocity required to unbind the system is

$$v_{\text{kick}} = \left(\frac{GM}{a}\right)^{1/2} = 0.5 \text{ km s}^{-1} \left(\frac{M}{M_{\odot}}\right)^{1/2} \left(\frac{a}{10^{3.5} \text{ AU}}\right)^{-1/2}. \quad (4.13)$$

To test explicitly how kicks during white dwarf formation change the separation distributions of WD/MS and WD/WD binaries, we modify BSE to include a velocity kick when a white dwarf forms. We use the same algorithm BSE uses for (stronger) kicks following SNe explosions. When a white dwarf is formed, we draw a kick velocity magnitude, v_{kick} , from a Maxwellian distribution characterized by a free parameter, σ_{kick} :

$$P(v_{\text{kick}}) = \sqrt{\frac{2}{\pi}} \frac{v_{\text{kick}}^2}{\sigma_{\text{kick}}^3} \exp\left[-\frac{v_{\text{kick}}^2}{2\sigma_{\text{kick}}^2}\right]; \quad (4.14)$$

this distribution peaks at $v_{\text{kick}} = \sqrt{2}\sigma_{\text{kick}}$. The kick is applied instantaneously in a random direction and occurs at a random time in the orbit. Following a kick, BSE self-consistently solves for the new eccentricity and semi-major axis, as described in Appendix A1 of Hurley et al. (2002b). If a kick unbinds the system, we consider the binary disrupted and discard it from our subsequent analysis.

In the bottom panel of Figure 4.10, we show the effects of applying kicks during WD formation with $\sigma_{\text{kick}} = 0.5 \text{ km s}^{-1}$ to the same Monte Carlo synthetic binary population. The main observable effect of including kicks is that the number of very wide WD/MS and WD/WD binaries is reduced, such that their distributions fall off more steeply at

⁴Within the typical lifetimes of PNe are $\sim 10^4$ yrs (Badenes et al. 2015), a 1 km s^{-1} kick will transport the WD a distance of order 0.01 pc, which is small compared to the pc-scale size of PNe.

$\log(a/\text{AU}) \gtrsim 3$ than the distribution of MS/MS binaries. The disrupted WD/MS and WD/WD binaries with large a are partially replaced by widened binaries with smaller initial separations. But because the orbits of the widest binaries are so weakly bound, only a small number of binaries have initial separations such that their orbits are significantly expanded without being fully disrupted. As a result, even weak kicks with $\sigma_{\text{kick}} = 0.5 \text{ km s}^{-1}$ cause the separation distributions of WD/WD and WD/MS binaries to fall off significantly more steeply at large a than the MS/MS distribution. For $\sigma_{\text{kick}} = 0.5 \text{ km s}^{-1}$, 21% of MS/WD binaries and 33% of WD/WD binaries are disrupted. Kicks unbind a majority of WD/MS binaries with initial separations $\log(a_{\text{init}}/\text{AU}) > 3.1$ and a majority of WD/WD binaries with $\log(a_{\text{init}}/\text{AU}) > 2.6$.

Comparison with observational constraints

We now compare the separation distributions predicted by BSE to our constraints from the *Gaia* 200 pc sample. The observational data constrain the distribution of *projected* separations, $P(s)$, while BSE returns the distribution of intrinsic orbital semi-major axes, $P(a)$. To carry out a fair comparison between these distributions, we compute in Appendix 4.8 the probability distribution, $P(a/s)$, of the conversion factor between s and a by mock-observing an ensemble of synthetic binary orbits along random lines of sight. For each binary evolved by BSE, we divide the intrinsic a by a random draw from $P(a/s)$.

The projected separation distributions predicted by BSE for different kick velocities are shown in Figure 4.11. Of course, kicks during WD formation have no effect on the separation distribution of MS/MS binaries. The observed constraints on the separation distribution of MS/MS binaries are reasonably similar, though not identical, to the distribution produced by BSE.

The separation distributions of WD/MS and WD/WD binaries appear most consistent with the simulations with $\sigma_{\text{kick}} = 0.5 \text{ km s}^{-1}$, corresponding to a typical kick speed of $\sim 0.75 \text{ km s}^{-1}$. Simulations with no kicks produce separation distributions that drop off significantly less steeply than the observational constraints. Conversely, stronger kicks cause the separation distribution of the simulated binary population to drop off more steeply than is observed. This does not rule out the possibility that such kicks occur for *some* white dwarfs, but it suggests that they are not widespread.

4.6 Summary and Discussion

We have constructed a catalog of high-confidence wide binaries using astrometric data from *Gaia* DR2. The catalog is publicly available and is described in Appendix 4.11. On its basis we determined with high fidelity the physical separation distributions of MS/MS, WD/MS, and WD/WD binaries over the separation range $1.5 \lesssim \log(a/\text{AU}) \lesssim 4.7$. We discovered that these distributions differ among these types of binaries, and use these differences to constrain the occurrence of kicks during white dwarf formation. We summarize our study below.




Figure 4.11: Comparison of our separation distribution constraints (shaded regions) to binary population synthesis simulations (“BSE” ; lines). We compare the results predicted for simulations in which mass loss in post-MS stars is isotropic and adiabatic (“no kicks”; solid lines) to simulations in which a stochastic velocity kick is applied during white dwarf formation (see Section 4.5.3). The separation distributions of WD/MS and WD/WD binaries are best matched by simulations with $\sigma_{\text{kick}} = 0.5 \text{ km s}^{-1}$ (corresponding to a most common kick speed of $\sim 0.75 \text{ km s}^{-1}$; see Equation 4.14); simulations with no kicks or larger kick velocities produce substantially worse agreement with the data. The separation distribution of MS/MS binaries (bottom) is independent of kicks but is also reasonably well matched by the simulations.

1. *Sample selection:* we require both members of a binary pair to have $d < 200$ pc and have positions, parallaxes, and proper motions consistent with a bound orbit (Figures 4.1 and 4.2). Members of open clusters, moving groups, and wide higher-order multiples are excised from the catalog (Figure 4.3). We classify each component as a main sequence star or a white dwarf using its location on a color-magnitude diagram (Figure 4.4). The catalog contains $\sim 50,000$ MS/MS binaries, ~ 3000 WD/MS binaries, and ~ 400 WD/WD binaries; the main sequence stars are a mix of solar-type stars and M-dwarfs ($(0.3 - 1.3)M_{\odot}$; Figure 4.5).
2. *Purity and completeness:* we estimate the contamination rate of the catalog by comparing the radial velocities of the components of candidate binaries (Figure 4.6) and by measuring the chance alignment rate in a mock catalog (Section 4.3.5). Precise parallaxes for both components make it possible to achieve unprecedentedly high purity, with an estimated contamination rate of order $\sim 0.2\%$.

The catalog is not volume complete. Incompleteness at small angular separations due to blending is straightforward to characterize empirically (Appendix 4.7) and is accounted for in our inference (Section 4.4.1). We do not attempt to correct for incompleteness due to missing faint sources, as it does not lead to any obvious biases in the inferred separation distribution.

3. *Separation distributions:* We infer the intrinsic separation distributions for each class of binaries (Figure 4.8) from the raw distributions (Figure 4.7). We discovered that the separation distributions of WD/MS and WD/WD binaries fall off more sharply at large separations than the corresponding distribution for MS/MS binaries. The break in the separation distribution occurs at smaller separations for WD/MS binaries than for MS/MS binaries, and at even smaller separations for WD/WD binaries. The log-separation distribution of MS/MS binaries continues to rise toward smaller separations (i.e., it is steeper than Opik’s law) to at least $s = 500$ AU (Figure 4.7). We infer consistent separation distributions for the full sample and the 100 pc sample (Appendix 4.9).
4. *Constraints on white dwarf kicks:* We use binary population synthesis simulations to predict the separation distributions of each class of binaries given a reasonable initial separation distribution and dynamical disruption rate (Figure 4.10). These simulations fail to reproduce the observed separation distributions of WD/MS and WD/WD binaries, producing too many binaries with wide separations ($s \gtrsim 5000$ AU). When modest velocity kicks during white dwarf formation due to asymmetric mass loss are included in the simulations (with typical velocities of ~ 0.75 km s $^{-1}$), they reproduce the observed separation distributions much more closely (Figure 4.11). Ubiquitous kicks with substantially larger kick velocities ($v_{\text{kick}} \gtrsim 2$ km s $^{-1}$) can be ruled out, though we cannot exclude the possibility that they occur in a fraction of systems.

4.6.1 Discussion

Introducing a velocity kick during white dwarf formation is of course not the only way to change the predicted separation distribution for binaries containing a white dwarf: one can also change the input initial separation distribution or the prescription for disruption due to external gravitational perturbations, both of which are uncertain. However, the nontrivial differences between the separation distributions of MS/MS binaries and binaries containing a white dwarf (Figure 4.8) calls for an additional disruption mechanism that affects only binaries containing a white dwarf.

Toonen et al. (2017) also considered the possibility that wide binaries are disrupted during white dwarf formation. Rather than modeling explicit natal kicks resulting from anisotropic mass loss, they assumed that mass loss is isotropic but can occur non-adiabatically, on a timescale that is short relative to the orbital timescale. In this case, wide binaries can be unbound when a WD forms because there is not enough time for the orbit to adjust to the reduced binding energy after mass loss (e.g. Savedoff 1966; Veras et al. 2011). Given a typical AGB lifetime of 1 Myr, one naively expects such effects to become important for $t_{\text{orbit}} \gg 1$ Myr, corresponding to $a \gg 10^4$ AU for solar-type stars.

Although Toonen et al. (2017) did not attempt to model the observed separation distribution, their calculations showed that rapid mass loss could unbind a fraction of white dwarfs in binaries with $a_{\text{init}} \gtrsim 5000$ AU. The primary goal of their simulations was to test whether sudden mass loss could explain the observed deficit of wide WD/WD binaries in the immediate solar neighborhood ($d < 20$ pc, where the observed white dwarf population is nearly complete) relative to the predictions of populations synthesis models. They found that rapid mass loss would not unbind enough systems to explain the shortage of observed WD/WD binaries, as only a small fraction of all binaries are wide enough to be disrupted.

We do not attempt to model the normalization of the white dwarf binary population in our sample, because its absolute completeness is not well characterized. Dissolution arising from rapid mass loss like that modeled by Toonen et al. (2017) could also be partially responsible for the shortage of very wide WD/MS and WD/WD binaries in our sample. However, we note that even if mass loss is instantaneous, it will not necessarily lead to disruption if it is isotropic. For example, a circular orbit will only be disrupted by instantaneous isotropic mass loss if the total mass of both stars decreases by more than 50% (Huang 1963), irrespective of the initial separation. This means that equal-mass binaries with circular orbits will *always* survive if only one star loses mass. For non-circular orbits, the probability of disruption also depends on orbital phase, but the majority of binaries are expected to survive rapid, isotropic mass loss if the total mass of the binary decreases by less than $\sim 70\%$ (Savedoff 1966).

Our population synthesis models assume kicks occur instantaneously; i.e., on timescales that are short compared to the orbital time. For the orbital separations we consider, this is always a good approximation if asymmetric mass loss occurs during PNe formation, but it can become problematic if deviations from spherically symmetric mass loss build gradually, over the full extent of the AGB phase. Using direct N-body simulations of a range binaries

orbits, we find that a majority of wide binaries that are disrupted by instantaneous kicks would also be disrupted if the same total velocity change were delivered through a gradual acceleration over many orbital timescales. We thus remain agnostic of the timescale over which anisotropic mass loss is likely to occur.

Irrespective of the details of the process responsible for ionizing binaries during post-MS evolution, one might also expect the magnitude of kicks to vary with the mass of the white dwarf, as the progenitors of massive white dwarfs lose a larger fraction of their mass during white dwarf birth (e.g. [El-Badry et al. 2018f](#)). Splitting the WD/MS binaries in our catalog by photometric white dwarf mass, we do not detect a significant difference between the separation distributions of binaries with lower- and higher-mass white dwarfs. However, the range of white dwarf masses in our catalog is not very large, and the photometric mass estimates are subject to nontrivial uncertainties. Further investigation of any dependence on the mass of the white dwarf and/or main sequence star represents a useful avenue for future work.

Finally, we note that kicks during white dwarf formation could accelerate the rate of WD/WD mergers and collisions in hierarchical triples, with potentially important implications for the SNe Ia rate. Kozai oscillations can drive inner WD/WD binaries in triples to extremely eccentric orbits, leading to a reduced gravitation wave inspiral timescale and/or head-on collisions of the WDs (e.g. [Thompson 2011](#); [Katz & Dong 2012](#)). The efficiency of this mechanism depends on the mutual inclination of the inner and outer binaries. Weak kicks during white dwarf formation could rearrange stable hierarchical triples containing an inner WD/WD binary, driving them to higher mutual inclinations where they are susceptible to the Kozai mechanism (e.g. [Toonen et al. 2018](#)). Simulations are needed to assess the efficiency of such kick-driven orbital rearrangement.

Using the catalog

Our sample of binaries serves as a potential resource for studies of the Galactic binary population, including the physics of star formation, ISM chemistry, cluster dissolution, and dynamical disruption. We make the catalog of high-confidence binaries within 200 pc publicly available in the hope that others find it useful.

Although the current catalog is truncated at separations of $50,000 \text{ AU} \approx 0.25 \text{ pc}$, extending the search to larger separations is straightforward (see [Appendix 4.10](#)); doing so may be useful for studies focused on dynamical disruption. At significantly larger separations, a more careful contamination model may be needed to account for unavoidable chance alignments of gravitationally unassociated pairs.

As discussed in [Section 4.3.2](#), the fact that our binary selection uses on-sky projected proper motions can lead to a slight bias against binaries with large angular separations, as stars with identical 3d space velocities can have different proper motions. The effects of this bias are expected to be negligible for our full sample and for the 100 pc sample, but it can lead to a non-trivial selection effect against the widest physical separations if only the nearest binaries are considered. The simulations described in [Section 4.3.2](#) show that if

only binaries within 50 pc (20 pc) are considered, 12% (60%) of binaries with $s > 10^4$ AU are likely to be missed. We therefore recommend a less stringent cut on proper motion for studies focused on the nearest binaries.

For more than 4000 of the binaries in our catalog, at least one component has been observed by a Galactic spectroscopic survey such as LAMOST, RAVE, APOGEE, and GALAH; for several hundred, a spectrum has been obtained for both stars. The metallicities and detailed abundances of binaries in the catalog can be used to investigate the chemical homogeneity of star forming regions, the creation rate through dynamical encounters of wide binaries that did not form together, and the dependence, or lack thereof, of the wide binary fraction on metallicity.

Wide binaries containing a white dwarf are particularly useful for the determination of stellar ages, as the total age of the white dwarf can be constrained from its cooling age and mass, given an initial-final mass relation and an initial mass-age relation. Assuming that the two components in a binary are co-eval, this age constraint can be transferred to a main-sequence companion, whose age could otherwise not be measured precisely (e.g. Foesneau et al. 2018). Such analysis is also useful for calibrating empirical age constraints for main sequence stars based on stellar activity or gyrochronology (e.g. Chanamé & Ramírez 2012; Godoy-Rivera & Chanamé 2018). The strongest constraints can be obtained for systems in which spectroscopic abundance constraints are available. In addition to the systems which have already been observed by Galactic spectroscopic surveys, we anticipate that *Gaia* DR3 will deliver abundances for a large fraction of binaries in our catalog.

Acknowledgements

We are grateful to the referee, Andrei Tokovinin, for a constructive report. We also thank Carles Badenes, Matteo Cantiello, Andy Casey, Trent Dupuy, Morgan Foesneau, Andy Gould, Ted von Hippel, David W. Hogg, Sergey Koposov, Chao Liu, Adrian Price-Whelan, Eliot Quataert, Jan Rybizki, David Spergel, and Daniel R. Weisz for helpful discussions. KE was supported by the SFB 881 program (A3) and an NSF graduate research fellowship. We are grateful to Livia DeMarchis for her hospitality during the writing of this manuscript. This project was developed in part at the 2018 NYC *Gaia* Sprint, hosted by the Center for Computational Astrophysics of the Flatiron Institute in New York City. This work has made use of data from the European Space Agency (ESA) mission *Gaia* (<https://www.cosmos.esa.int/gaia>), processed by the *Gaia* Data Processing and Analysis Consortium (DPAC, <https://www.cosmos.esa.int/web/gaia/dpac/consortium>). Funding for the DPAC has been provided by national institutions, in particular the institutions participating in the *Gaia* Multilateral Agreement.

4.7 Gaia spatial resolution and dependence on flux ratio

Our method of identifying wide binaries requires that the two component stars be spatially resolved and that both stars pass the quality cuts we impose (Section 4.3.1). Whether a wide binary is detected thus depends both on the angular separation of the two stars and on their flux ratio. The secondary is more likely to be outshone or contaminated by light from the primary at fixed angular separation if the flux ratio is large than if it is small. This must be accounted for in modeling the intrinsic separation distribution.

We assess the sensitivity of *Gaia* photometry to a companion at a given angular separation using a method similar to that described in Arenou et al. (2018a). We query all sources in a field of radius 1.5 degrees and measure the number of pairs passing the quality cuts listed in Section 4.3.1 as a function of angular separation and ΔG , the G -band magnitude difference between the two sources. The resulting two-point correlation function is shown in the top panel Figure 4.12. For each bin in ΔG , we plot the number of pairs in 0.66 arcsec-wide bins. With no restriction on parallax or proper motion, the vast majority of pairs are chance alignments. At any flux ratio and angular separation θ , the expected number of such random alignments scales as $N_{\text{pairs}} \sim 2\pi\theta d\theta$.

The top panel of Figure 4.12 shows that such a linear scaling is indeed found at large separations for all ΔG bins. To estimate the total number of pairs at a given θ and ΔG , including those that are missed due to incompleteness at small separations, we extrapolate a linear fit to the number of pairs found at large separations. We then estimate the fraction of all pairs that are detected by dividing the empirically measured number of pairs by this extrapolation (bottom panel of Figure 4.12). For all ΔG , the sensitivity drops rapidly from 1 to 0 at some angular separation θ_0 , where θ_0 increases with ΔG . We fit the sensitivity curves for each ΔG with an empirical fitting function:

$$f_{\Delta G}(\theta) = \frac{1}{1 + (\theta/\theta_0)^{-\beta}}. \quad (4.15)$$

Here θ_0 characterizes the angular separation below which the sensitivity drops to 0, and β determines how rapidly the sensitivity falls off at $\theta \ll \theta_0$. We estimate the value of θ_0 and β appropriate for a given ΔG by interpolating on the fits for the discrete values of ΔG shown in Figure 4.12. We find that $\beta = 10$ works reasonably well for all ΔG . For θ_0 , we find $\theta_0 \approx 2.25$ arcsec at $\Delta G < 1.5$ mag, and $\theta_0 \approx 0.9(\Delta G + 1)$ at $\Delta G > 1.5$ mag.

We find that this parameterization of the sensitivity works well over a wide range of colors and source surface densities. However, we caution that it applies only for the particular choice of quality cuts listed in Section 4.3.1. Ziegler et al. (2018) found a comparable sensitivity scaling in the efficiency with which *Gaia* recovers known binaries, but with a different normalization θ_0 due to less stringent quality cuts on the *Gaia* photometry.




Figure 4.12: **Top**: total number of pairs with separation θ (i.e., two-point autocorrelation function) in a dense region ($l, b = -30^\circ, -4^\circ$). Different lines correspond to pairs in bins of different magnitude difference, $\Delta G = |G_1 - G_2|$. Dashed lines show linear fits to data at $\theta > 10$ arcsec. **Bottom**: ratio of data in top panel to linear extrapolations, representing the fraction of all pairs with a given ΔG that are detected at angular separation θ . Dashed lines show fits using Equation 4.15.




Figure 4.13: Distribution of the conversion factor, $P(a/s)$, from projected separation s to semi-major axis a for randomly oriented orbits. We compare distributions for a binary population with a uniform eccentricity distribution (as is assumed in our population synthesis simulations; black) and a thermal eccentricity distribution ($f(e) de = 2e de$, red).

4.8 Relation between projected separation and semimajor axis

When comparing BSE simulations to our observational constraints, it is necessary to translate between distributions of the intrinsic semi-major axis, a , and the projected separation, s . To determine an appropriate conversion factor distribution, we mock-observe an ensemble of 3×10^6 simulated binaries.

Each binary is observed along a random line-of-sight (resulting in an $a p(i) di = \sin(i) di$ inclination distribution) at a random time in its orbit. We consider two eccentricity distributions: a uniform distribution, which is roughly consistent with observations (e.g. Duchêne & Kraus 2013; Tokovinin & Kiyaveva 2016) and is also assumed in our BSE simulations, and a thermal distribution, $p(e) de = 2e de$, which is predicted if the phase-space distribution of binaries is a function of energy alone (e.g. Jeans 1919). The distributions of stellar mass, mass ratio, and orbital period have no effect on the conversion factor between s and a .

The results of these simulations are shown in Figure 4.13; they are identical to the findings of Dupuy & Liu (2011), who performed a similar experiment. For each binary simulated in our population synthesis experiments in Section 4.5, we multiply the semimajor axis by a random draw from the $P(a/s)$ distribution for a uniform eccentricity distribution in order to construct the $P(s)$ distribution. For both eccentricity distributions, s and a differ by less than a factor of two for a large majority of binaries. Because we consider separation distributions over separations varying by many orders of magnitude, $P(s)$ and $P(a)$ are nearly indistinguishable.

4.9 Distance dependence

In Figure 4.14, we compare the separation distributions obtained for our full binary sample to those obtained when we apply our inference only to the binaries within 100 pc (i.e., $\varpi > 10$ mas). We expect the intrinsic separation distribution of the 100 pc sample to be similar to that of the full sample. Because the typical parallax and proper motion uncertainties, angular separation, and expected contamination rate of the 100 pc sample are different from those of the full catalog, this experiment provides a useful test of the inference described in Section 4.4.1.

The 100 pc sample contains 10179, 977, and 177 MS/MS, WD/MS, and WD/WD binaries, respectively. Figure 4.14 shows that the constraints obtained for each class of binary are quite similar in the 100 pc sample to in the full sample. Although there are slight differences in the constraints, the separation distributions are consistent within 2σ at all separations, and none of our conclusions would change if we were to fit only the 100 pc sample.

In the top panel of Figure 4.14, we also show the best-fit distribution obtained for MS/MS binaries when we replace the default broken power law parameterization of $\phi(s)$ with a lognormal; i.e., a Gaussian in $\log(s)$. The formal best-fit parameters are $\mu_{\log s/\text{AU}} = -0.93$ and $\sigma_{\log s/\text{AU}} = 1.69$. However, we caution that because there is no strong evidence of a turnover in $\phi(s)$ over the range of separations probed by our data, $\mu_{\log s}$ and $\sigma_{\log s}$ are strongly degenerate, and the best-fit value of $\mu_{\log s}$ is meaningless (see El-Badry et al. 2018a, their Figure 1). For example, a lognormal with $\mu_{\log s/\text{AU}} = 1.0$ and $\sigma_{\log s/\text{AU}} = 1.26$ provides a comparably good fit.

4.10 ADQL query

The ADQL query used to obtain our initial sample of candidate binaries (before removing clusters, triples, etc., and before applying all the quality cuts listed in Section 4.3.1) is reproduced below. The search approach is modeled after that used by Fouesneau et al. (2018). We follow the variable naming convention from the *Gaia* Archive.

```
SELECT g2.source_id as source_id2,
       g2.ra as ra2, g2.dec as dec2,
```





Figure 4.14: 1σ constraints on the separation distributions of different classes of wide binaries. We compare constraints obtained from fitting the full sample within 200 pc (solid; identical to Figure 4.8) to those obtained when fitting only the nearest objects within 100 pc (hatched); they are essentially consistent. Distance distributions are shown in Figure 4.5. In the top panel, we also show the best-fit lognormal distribution for the full MS/MS sample.


```

g2.astrometric_chi2_al as
astrometric_chi2_al_2,
g2.astrometric_n_good_obs_al
as astrometric_n_good_obs_al2,
g2.phot_g_mean_flux_over_error
as phot_g_mean_flux_over_error2,
g2.phot_rp_mean_flux_over_error
as phot_rp_mean_flux_over_error2,
g2.phot_bp_mean_flux_over_error
as phot_bp_mean_flux_over_error2,
g2.phot_bp_rp_excess_factor
as phot_bp_rp_excess_factor2,
g2.pmra as pmra2, g2.pmra_error
as pmra_error2, g2.pmdec as pmdec2,
g2.pmdec_error as pmdec_error2,
g2.phot_g_mean_mag as
phot_g_mean_mag2, g2.phot_bp_mean_mag
as phot_bp_mean_mag2,
g2.phot_rp_mean_mag as
phot_rp_mean_mag2, g2.parallax as
parallax2, g2.parallax_over_error as
parallax_over_error2, g2.radial_velocity
as radial_velocity2,
g2.radial_velocity_error as
radial_velocity_error2, g2.rv_nb_transits
as rv_nb_transits2, t1.source_id, t1.ra,
t1.dec, t1.pmra, t1.pmra_error, t1.pmdec,
t1.pmdec_error, t1.phot_g_mean_mag,
t1.phot_bp_mean_mag, t1.phot_rp_mean_mag,
t1.parallax, t1.parallax_over_error,
t1.astrometric_chi2_al,
t1.astrometric_n_good_obs_al,
t1.phot_g_mean_flux_over_error,
t1.phot_rp_mean_flux_over_error,
t1.phot_bp_mean_flux_over_error,
t1.phot_bp_rp_excess_factor,
t1.radial_velocity, t1.radial_velocity_error,
t1.rv_nb_transits,
distance(POINT('ICRS', t1.ra, t1.dec),
POINT('ICRS', g2.ra, g2.dec)) AS pairedistance
FROM (select * from gaiadr2.gaia_source where
parallax between 5 and 1000 and

```

```

parallax_over_error > 20 and
phot_g_mean_flux_over_error > 50 and
phot_rp_mean_flux_over_error > 20 and
phot_bp_mean_flux_over_error > 20) as t1,
(select * from gaiadr2.gaia_source
where bp_rp is not null and
parallax_over_error > 5 and
phot_g_mean_flux_over_error > 50 and
phot_rp_mean_flux_over_error > 10 and
phot_bp_mean_flux_over_error > 10) as g2
WHERE 1 = contains(POINT('ICRS', g2.ra, g2.dec),
CIRCLE('ICRS', t1.ra, t1.dec,
1.4e-2*t1.parallax))
AND t1.source_id != g2.source_id
AND abs(1/t1.parallax - 1/g2.parallax) -
2*0.01745*distance(POINT('ICRS', t1.ra,
t1.dec), POINT('ICRS', g2.ra, g2.dec))/
t1.parallax < 3*sqrt(power(t1.parallax_error,
2)/power(t1.parallax, 4) + power(
g2.parallax_error, 2)/power(g2.parallax, 4))
AND sqrt(power((t1.pmra - g2.pmra), 2) +
power((t1.pmdec - g2.pmdec), 2)) -
(7.42e-3 * power(t1.parallax, 1.5) *
power(distance(POINT('ICRS', t1.ra, t1.dec),
POINT('ICRS', g2.ra, g2.dec)), -0.5)) <
3*sqrt(((power((t1.pmra_error), 2) + power(
(g2.pmra_error), 2)) * power((t1.pmra -
g2.pmra), 2) + (power((t1.pmdec_error), 2) +
power((g2.pmdec_error), 2)) * power(
(t1.pmdec - g2.pmdec), 2))/(power(
(t1.pmra - g2.pmra), 2) + power(
(t1.pmdec - g2.pmdec), 2))))
AND sqrt(((power((t1.pmra_error), 2) +
power((g2.pmra_error), 2)) * power(
(t1.pmra - g2.pmra), 2) + (power(
(t1.pmdec_error), 2) + power(
(g2.pmdec_error), 2)) * power((t1.pmdec -
g2.pmdec), 2))/(power((t1.pmra - g2.pmra),
2) + power((t1.pmdec - g2.pmdec), 2)))) < 1.5

```

In practice, this query will not finish within the default time limits on the *Gaia* Archive. We therefore split it on source id into many smaller queries. That is, we add the condition

where `source_id` between A and B, where A and B are integers, to both clauses in the FROM block. We find splitting on healpix level 2 (i.e., 192 regions of area $\approx 215 \text{ degree}^2$) to work well. Splitting up the sky during the search could in principle introduce a bias against pairs with large separations, which are more likely to be missed because the two component stars fall in different regions of the sky. We have verified through simulations that any resulting bias is at the $< 3\%$ level, because the typical angular separations of even the widest binaries in our catalog are small compared to the size of each search region.

4.11 Catalog

Table 4.1 provides a summary of our catalog of wide binaries within 200 pc. A full version of the catalog is available online.

Table 4.1: Catalog description

Column	units	Description
astrometric_chi2_al		astrometric goodness-of-fit (χ^2) in the along-scan direction; star 1
astrometric_chi2_al_2		astrometric goodness-of-fit (χ^2) in the along-scan direction; star 2
astrometric_n_good_obs_al		number of good CCD transits; star 1
astrometric_n_good_obs_al2		number of good CCD transits; star 2
binary_class		type of binary; MSMS, WDMS, or WDWD
dec	deg	declination; star 1
dec2	deg	declination; star 2
pairedistance	deg	angular separation between star 1 and star 2
parallax	mas	parallax; star 1
parallax2	mas	parallax; star 2
parallax_over_error		parallax divided by its error; star 1
parallax_over_error2		parallax divided by its error; star 2
phot_bp_mean_flux_over_error		integrated BP mean flux divided by its error; star 1
phot_bp_mean_flux_over_error2		integrated BP mean flux divided by its error; star 2
phot_bp_mean_mag	mag	integrated BP mean magnitude; star 1
phot_bp_mean_mag2	mag	integrated BP mean magnitude; star 2
phot_bp_rp_excess_factor		ratio of total integrated BP and RP flux to G-band flux; star 1
phot_bp_rp_excess_factor2		ratio of total integrated BP and RP flux to G-band flux; star 2
phot_g_mean_flux_over_error		integrated G-band mean flux divide by its error; star 1
phot_g_mean_flux_over_error2		integrated G-band mean flux divide by its error; star 2
phot_g_mean_mag	mag	G-band mean magnitude (Vega scale); star 1
phot_g_mean_mag2	mag	G-band mean magnitude (Vega scale); star 2
phot_rp_mean_flux_over_error		integrated RP mean flux divided by its error; star 1
phot_rp_mean_flux_over_error2		integrated RP mean flux divided by its error; star 2
phot_rp_mean_mag	mag	integrated RP mean magnitude; star 1
phot_rp_mean_mag2	mag	integrated RP mean magnitude; star 2
pmdec	mas yr ⁻¹	proper motion in the declination direction; star 1
pmdec2	mas yr ⁻¹	proper motion in the declination direction; star 2
pmdec_error	mas yr ⁻¹	standard error of proper motion in the declination direction; star 1
pmdec_error2	mas yr ⁻¹	standard error of proper motion in the declination direction; star 2
pmra	mas yr ⁻¹	proper motion in right ascension direction; i.e., $\mu_\alpha^* = \mu_\alpha \cos \delta$; star 1
pmra2	mas yr ⁻¹	proper motion in right ascension direction; i.e., $\mu_\alpha^* = \mu_\alpha \cos \delta$; star 2
pmra_error	mas yr ⁻¹	standard error of proper motion in right ascension direction; star 1
pmra_error2	mas yr ⁻¹	standard error of proper motion in right ascension direction; star 2
ra	deg	right ascension; star 1
ra2	deg	right ascension; star 2
radial_velocity	km s ⁻¹	spectroscopic barycentric radial velocity; star 1
radial_velocity2	km s ⁻¹	spectroscopic barycentric radial velocity; star 1
radial_velocity_error	km s ⁻¹	standard error of spectroscopic barycentric radial velocity; star 1
radial_velocity_error2	km s ⁻¹	standard error of spectroscopic barycentric radial velocity; star 2
rv_nb_transits		number of transits used to compute radial velocity; star 1
rv_nb_transits2		number of transits used to compute radial velocity; star 2
s_AU	AU	projected physical separation between two stars
source_id		<i>Gaia</i> source id (int64); star 1
source_id2		<i>Gaia</i> source id (int64); star 2

Note: Each row in the catalog corresponds to a single binary; “star 1” and “star 2” designations in each binary are arbitrary. Full descriptions of *Gaia* variables can be found at <https://gea.esac.esa.int/archive/documentation/GDR2/>.

Chapter 5

The wide binary fraction of solar-type stars: emergence of metallicity dependence at $a < 200$ AU

An earlier version of this article was previously published as El-Badry, K. and Rix, H.-W., 2019, MNRASL, 482, 139.

5.1 Abstract

We combine a catalog of wide binaries constructed from *Gaia* DR2 with $[\text{Fe}/\text{H}]$ abundances from wide-field spectroscopic surveys to quantify how the binary fraction varies with metallicity over separations $50 \lesssim s/\text{AU} \lesssim 50,000$. At a given distance, the completeness of the catalog is independent of metallicity, making it straightforward to constrain *intrinsic* variation with $[\text{Fe}/\text{H}]$. The wide binary fraction is basically constant with $[\text{Fe}/\text{H}]$ at large separations ($s \gtrsim 250$ AU) but becomes quite rapidly anti-correlated with $[\text{Fe}/\text{H}]$ at smaller separations: for $50 < s/\text{AU} < 100$, the binary fraction at $[\text{Fe}/\text{H}] = -1$ exceeds that at $[\text{Fe}/\text{H}] = 0.5$ by a factor of 3, an anti-correlation almost as strong as that found for close binaries with $a < 10$ AU. Interpreted in terms of models where disk fragmentation is more efficient at low $[\text{Fe}/\text{H}]$, our results suggest that $100 < a/\text{AU} < 200$ is the separation below which a significant fraction of binaries formed via fragmentation of individual gravitationally unstable disks rather than through turbulent core fragmentation. We provide a public catalog of 8,407 binaries within 200 pc with spectroscopically-determined $[\text{Fe}/\text{H}]$ for at least one component.

5.2 Introduction

The statistical properties of the binary star population and their potential variation with metallicity provide powerful diagnostics of the star formation process ([White & Ghez 2001](#);

Machida et al. 2009; Duchêne & Kraus 2013; Bate 2014; Badenes et al. 2018; Moe et al. 2018). Different binary formation mechanisms predict different metallicity dependences of the resulting binary fraction (Moe et al. 2018). Measurements of metallicity dependence as a function of orbital separation can thus constrain the separation regimes in which different formation mechanisms operate.

The *close* binary fraction is strongly anti-correlated with metallicity: the fraction of primaries that have a companion with $a < 10$ AU decreases by a factor of 4 over $-1 < [\text{Fe}/\text{H}] < 0.5$, from $(40 \pm 6)\%$ at $[\text{Fe}/\text{H}] = -1$ to $(10 \pm 3)\%$ at $[\text{Fe}/\text{H}] = 0.5$ (Raghavan et al. 2010; Badenes et al. 2018; Moe et al. 2018). The existence of such anti-correlation was long controversial (Jaschek & Jaschek 1959; Carney 1983; Latham et al. 2002; Hettinger et al. 2015), in large part because the sensitivity of most binary detection methods varies with metallicity. But in a recent re-analysis of the binary populations probed by 5 surveys, Moe et al. (2018) found that after observational biases are corrected for, all 5 show consistent evidence of a metallicity-dependent close binary fraction.

On the other hand, studies of the *wide* binary fraction have found it to be approximately metallicity-invariant (e.g. Zapatero Osorio & Martín 2004; Chanamé & Gould 2004; El-Badry et al. 2018g). Most of the binaries studied in these works have separations of order 1000 AU, but the variation of metallicity dependence with orbital separation has yet to be studied in detail. The recent *Gaia* data releases (Gaia Collaboration et al. 2016, 2018a) have substantially simplified the process of reliably identifying spatially resolved binaries, making it possible to study the intermediate-to-wide binary population with unprecedented precision (e.g. Oh et al. 2017a; Oelkers et al. 2017; Andrews et al. 2017). In this Letter, we seek to constrain the transition between the close and wide binary regimes, pinpointing the separation at which metallicity dependence emerges.

5.3 Methods

5.3.1 Wide binary catalog and extension

We extend the wide binary catalog described in El-Badry & Rix (2018a, hereafter ER18), which was constructed by searching *Gaia* DR2 for pairs of stars whose positions, proper motions, and parallaxes are consistent with being gravitationally bound. The ER18 catalog contains $\sim 50,000$ binaries with separations $50 \lesssim s/\text{AU} < 50,000$, with an estimated contamination rate of $\sim 0.2\%$. To maintain high purity, it only contains binaries that are within 200 pc of the Sun and have high-quality astrometry and photometry.

To separate main sequence stars and white dwarfs, ER18 required both stars to have a measured `bp_rp` color, and to have well-resolved photometry as quantified by the `phot_bp_rp_excess_factor` (see Section 2.1 of ER18). As a consequence of these requirements, the ER18 catalog has an effective resolution limit of ~ 2 arcsec and contains few binaries with separations $s < 200$ AU. To find more binaries with small separations, we now extend the catalog by removing the restrictions on `bp_rp` and `phot_bp_rp_excess_factor` *only for pairs with projected separa-*

Table 5.1: Sources of the spectroscopic abundances

Sample	$N_{\text{in binary}}$	$N_{\text{tot}, d < 200 \text{ pc}}$
RAVE	3,261	33,792
LAMOST	3,729	46,729
APOGEE	660	5,796
GALAH	537	6,759
Hypatia	660	3,954
Total stars	8,847	97,030
Total number of binaries	8,407	

tions $s < 500$ AU. We still require both components to pass the other quality cuts in ER18, including having reliable astrometry and precisely measured, mutually consistent parallaxes, and we apply the same procedure for removing members of clusters, moving groups, and resolved higher-order multiples.

This extended search yields 23,079 new binaries not included in the initial catalog. Many of the new additions have angular separations between 0.5 and 2 arcsec. Combining them with the sample from ER18 results in a total of 78,207 binaries, including 8,284 with $s < 200$ AU. Although the photometry of the objects that did not pass the ER18 cuts is less clean, we expect essentially all of them to be bona fide binaries, as the contamination rate from chance alignments is negligible at close separations.

5.3.2 Spectroscopic metallicities

We cross-matched the expanded wide binary catalog with several wide-field spectroscopic surveys: LAMOST (DR5; Zhao et al. 2012), RAVE (DR5; Kunder et al. 2017), APOGEE (DR14; Majewski et al. 2017, using the abundances derived by Ting et al. 2018), and GALAH (DR2; Buder et al. 2018). We also cross-matched with the Hypatia catalog (Hinkel et al. 2014), which is a compilation of high-resolution spectroscopic abundances for stars within 150 pc of the Sun. We limit our sample to main-sequence binaries in which the primary has absolute magnitude $2.5 < M_G < 9.5$, corresponding to $0.45 \lesssim M/M_\odot \lesssim 1.5$.

The resulting catalog is summarized in Table 5.1. Cross-matching yields a spectroscopic [Fe/H] for at least one component of 8,407 binaries; in 440, a spectroscopic [Fe/H] is available for both components. We assign binaries in which both components have a spectroscopic [Fe/H] the mean of the two components; when only one component has a measured [Fe/H], we adopt that value. For stars that were observed by more than one survey, we prioritize abundances from surveys in the reverse order listed in Table 5.1. The catalog is available online.

We also construct a spectroscopic “control sample” that consists of the 97,030 stars within 200 pc that were observed by the spectroscopic surveys listed in Table 5.1 and pass the *Gaia* quality and magnitude cuts applied to the wide binary sample. Our method for identifying wide binaries is metallicity-blind, and the spectroscopic surveys did not preferentially target

or avoid wide binaries. Therefore, any metallicity bias in the spectroscopic binary sample will, at fixed distance, affect the control sample in the same way it affects the binary sample.

Binaries with angular separations of less than a few arcsec may not be spatially resolved by ground-based spectroscopic surveys. The resulting errors in $[\text{Fe}/\text{H}]$ are expected to be less than 0.1 dex on average, with negligible systematic biases for the mid-resolution optical spectra that constitute the majority of our sample (Schlesinger et al. 2010; El-Badry et al. 2018b).

5.4 Results

5.4.1 Metallicity Distribution

Figure 5.1 compares the metallicity distribution functions (MDFs) of binaries in different separation bins to the MDF of the control sample. Because close binaries are unresolved at large distances, their distributions of heliocentric distance are different from those of the full 200 pc control sample. To avoid biases arising from the distance-dependence of the MDF, for each bin in s , we compare to a random subset of the control sample with the same distance distribution as the binaries in that s bin.

At small separations, the MDFs of binaries are biased toward low $[\text{Fe}/\text{H}]$ relative to the control sample. This bias is strongest in the $50 < s/\text{AU} < 100$ bin but is present in all separation bins up to $s = 250$ AU. No strong bias toward higher or lower $[\text{Fe}/\text{H}]$ is evident at large separations, though there are hints of a slight excess of binaries with $[\text{Fe}/\text{H}] \sim 0$ in the largest separation bin. The latter is likely attributable to age effects: at the widest separations, there is a non-negligible probability for binaries to be dynamically disrupted by gravitational perturbations from other stars and molecular clouds. Lower-metallicity binaries are on average older, allowing more time for dynamical disruption.

5.4.2 Inferring the dependence of binarity on $[\text{Fe}/\text{H}]$

To quantify the metallicity dependence of the binary fraction implied by the binary MDFs in Figure 5.1, we fit a model that describes the ratio of the MDF of binaries to that of all stars as a function of metallicity and separation. Given a normalized, empirically measured MDF for all stars, $\psi([\text{Fe}/\text{H}]) = dP/d[\text{Fe}/\text{H}]$, we define the MDF of binaries with separation s as

$$\psi_b([\text{Fe}/\text{H}]|s, \vec{m}) = \psi([\text{Fe}/\text{H}]) w([\text{Fe}/\text{H}]|s, \vec{m}) \times c. \quad (5.1)$$

Here w is an arbitrary, positive-definite “weight function” parameterized by a vector of model parameters \vec{m} , and c is a normalization constant,

$$c(s, \vec{m}) = \left[\int \psi([\text{Fe}/\text{H}]) w([\text{Fe}/\text{H}]|s, \vec{m}) d[\text{Fe}/\text{H}] \right]^{-1}. \quad (5.2)$$




Figure 5.1: Metallicity distribution functions (MDFs) of binaries (red) and all stars with spectroscopic metallicities (black). We separate binaries in bins of projected physical separation, s ; in each panel, the black histogram shows the MDF of a random subset of the control sample with the same distance distribution as the binaries in that panel. At small separations ($s < 250$ AU), the binary MDF shows a shortage of metal-rich binaries and an excess of metal-poor binaries relative to the control sample. At larger separations, the MDF of binaries is nearly identical to that of all stars, with a slight excess of metal-rich systems at the widest separations.

If the MDF of all binaries is identical to that of single stars, then $w = \text{const}$. If the binary fraction varies with metallicity in way that affects binaries of all separations equally, w is a function of $[\text{Fe}/\text{H}]$ alone. In the general case where the binary fraction varies with $[\text{Fe}/\text{H}]$ in a separation-dependent way, $w = w([\text{Fe}/\text{H}], s)$. The absolute normalization of w is arbitrary due to the normalization condition, but its dependence on $[\text{Fe}/\text{H}]$ and s determines the binary fraction relative to its value at a fixed metallicity; e.g.

$$\frac{f_{\text{bin}}(s, [\text{Fe}/\text{H}])}{f_{\text{bin}}(s, [\text{Fe}/\text{H}] = 0)} = \frac{w(s, [\text{Fe}/\text{H}])}{w(s, [\text{Fe}/\text{H}] = 0)}. \quad (5.3)$$

We parameterize w with a flexible function that allows the metallicity dependence to vary with separation:

$$w = \frac{1}{2} \left[(w_0 + 1) - (w_0 - 1) \tanh \left(\frac{[\text{Fe}/\text{H}] - [\text{Fe}/\text{H}]_0}{\gamma} \right) \right]. \quad (5.4)$$

This parameterization causes w to asymptote to 1 at $[\text{Fe}/\text{H}] \gg [\text{Fe}/\text{H}]_0$ and asymptote to w_0 at $[\text{Fe}/\text{H}] \ll [\text{Fe}/\text{H}]_0$. The abruptness of the transition between the two regimes is determined by γ . To allow the asymptotic behavior to vary with s , we parameterize w_0 as

$$w_0 = 1 + \frac{A - 1}{1 + (s/s_0)^\beta}, \quad (5.5)$$

which asymptotes to $w_0 = 1$ at $s \gg s_0$ and to $w_0 = A$ at $s \ll s_0$. The abruptness of the transition is determined by β , which is required to be positive.

Given a sample of binaries with metallicities $[\text{Fe}/\text{H}]_i$ and separations s_i , the likelihood function is

$$L = p(\{[\text{Fe}/\text{H}]_i, s_i\} | \vec{m}) = \prod_i \psi_b([\text{Fe}/\text{H}]_i | s_i, \vec{m}), \quad (5.6)$$

with ψ_b given by Equation 5.1. We tabulate the ψ appropriate for the distance distribution of binaries as a function of s , using a subset of stars from the control sample with the same distance distribution as binaries of separation s . We sample the posterior distribution using `emcee` (Foreman-Mackey et al. 2013b), with broad, flat priors on all model parameters $\vec{m} = (A, \log(s_0/\text{AU}), \beta, [\text{Fe}/\text{H}]_0, \gamma)$.

We plot the resulting constraints on the binary fraction in Figure 5.2, normalizing relative to $[\text{Fe}/\text{H}] = 0$. Consistent with the qualitative picture from Figure 5.1, the binary fraction becomes metallicity-dependent at $s \lesssim 250$ AU. Over $-0.5 < [\text{Fe}/\text{H}] < 0.5$, the dependence on metallicity is nearly linear, with hints of flattening at lower $[\text{Fe}/\text{H}]$. The emergence of metallicity dependence towards smaller separations is quite rapid: there is essentially no dependence at $s = 300$ AU, while at $s = 100$ AU, the binary fraction at $[\text{Fe}/\text{H}] = -1$ is a factor of 3 higher than at $[\text{Fe}/\text{H}] = 0.5$. The metallicity dependence at the smallest separations probed by our catalog is fully consistent with that found at $a < 10$ AU by Moe et al. (2018), though our median constraints lean toward somewhat weaker metallicity




Figure 5.2: Metallicity-dependence of the binary fraction. Top panels show constraints at fixed metallicity as a function of separation; bottom panels show constraints at fixed separation as a function of metallicity. In all panels, we normalize by the separation-dependent binary fraction at $[\text{Fe}/\text{H}] = 0$. Hints of metallicity-dependence first appear at $s \sim 250$ AU and become stronger with decreasing separation. For comparison, we show constraints on the metallicity-dependence of the close binary fraction ($a < 10$ AU) from [Moe et al. \(2018\)](#).

dependence at $[\text{Fe}/\text{H}] < -0.5$. The uncertainties in our constraints are substantial at small separations due to the small number of binaries in our catalog with $s < 100$ AU. Nevertheless, the data rule out a metallicity-independent binary fraction with greater than 2σ significance for all $s < 150$ AU.

5.4.3 Separation distributions

In Figure 5.3, we show inferred intrinsic separation distributions over $50 < s/\text{AU} < 500$ (where the binary selection function is consistent and well-characterized; see Section 5.3.1) of binaries in 4 metallicity bins. The relative detection efficiency drops below 50% at an angular separation of $\theta_0 \approx 1$ arcsec, so the maximum distance at which a binary of separation s can be detected is $d_{\text{max}}/\text{pc} \approx \min \{(s/\text{AU}) / (\theta_0/\text{arcsec}), 200\}$. In inferring the intrinsic separation distribution, we weight each observed binary by the fraction of objects in the control sample at a particular $[\text{Fe}/\text{H}]$ that are at distances greater than d_{max} .

Over $50 < s/\text{AU} < 500$, the separation distributions of lower- $[\text{Fe}/\text{H}]$ systems consistently show an excess of smaller-separation binaries relative to those of higher- $[\text{Fe}/\text{H}]$ systems. Integrating over all metallicities, the total separation distribution peaks at $a \approx 200$ AU and is relatively flat over the separations shown in Figure 5.3 (Duchêne & Kraus 2013). However, Figure 5.3 shows that the peak in the separation distribution occurs at smaller (larger) separations for binaries with low (high) metallicity. Using a weighted two-sample KS test, we verified that the separation distributions for both of the two highest- $[\text{Fe}/\text{H}]$ bins are inconsistent with those of each of the two lowest- $[\text{Fe}/\text{H}]$ bins at at least the 4σ level ($p_{\text{KS}} < 3 \times 10^{-5}$). Indeed, such variation is required to explain a binary fraction that is metallicity-dependent at small separations but not at asymptotically large separations (see Moe et al. 2018, their Figure 19).

5.5 Discussion and Conclusions

We have shown that the metallicity distribution functions (MDFs) of binaries with separations $50 \lesssim s/\text{AU} \lesssim 250$ exhibit a shortage of high- $[\text{Fe}/\text{H}]$ binaries and an excess of low- $[\text{Fe}/\text{H}]$ binaries relative to a control sample subject to the same selection function (Figure 5.1). We fit a flexible, parameterized model for the modification of the binary MDF relative to the MDF of all stars, thus constraining the $[\text{Fe}/\text{H}]$ -dependence of the binary fraction as a function of separation (Figure 5.2). Metallicity dependence is weak at $s = 250$ AU but ramps up rapidly with decreasing separation: at $50 < s/\text{AU} < 100$, the binary fraction increases by a factor of 3 over $-1 < [\text{Fe}/\text{H}] < 0.5$. For these separations, the metallicity-dependence is roughly linear over $-0.5 < [\text{Fe}/\text{H}] < 0.5$. It begins to flatten at $[\text{Fe}/\text{H}] < -0.5$, albeit with substantial uncertainty. The separation distribution is similarly metallicity-dependent, with low-metallicity binaries concentrated at smaller separations (Figure 5.3).

The projected separation s of a wide binary can exceed the semi-major axis a by at most a factor of 2. Projection effects will tend to smooth out the transition between the




Figure 5.3: Normalized separation distributions over $50 < s/\text{AU} < 500$ for binaries of different metallicities, after correcting for incompleteness at small separations. Low- (high-) $[\text{Fe}/\text{H}]$ binaries are weighted toward smaller (larger) separations.

metallicity-dependent and independent regimes, so the weak metallicity dependence detected at $s \gtrsim 200$ AU could be a consequence of stronger metallicity dependence at $a \sim 150$. However, a and s are usually similar for realistic orbits (see ER18; their Figure B1). For a uniform eccentricity distribution, 16% of randomly observed orbits satisfy $s > 1.33a$, and 2.3% satisfy $s > 1.75a$. Metallicity dependence out to $s \sim 250$ AU thus implies dependence to at least $a = 150$ AU.

In the model proposed by Moe et al. (2018), the binary fraction is metallicity-dependent only at separations where binaries formed primarily via the fragmentation of gravitationally unstable disks: the turbulent core fragmentation process that produces wider binaries is independent of metallicity, at least for the range of $[\text{Fe}/\text{H}]$ considered here. Interpreted in terms of this model, our results suggest that $100 < a/\text{AU} < 200$ is the separation below which a substantial fraction of solar-type binaries were formed via disk fragmentation.

In simulations, disk fragmentation typically occurs at separations of $a/\text{AU} \sim 50 - 100$ (Burkert et al. 1997; Machida et al. 2009; Stamatellos & Whitworth 2009; Kratter et al. 2010), in agreement with observations that find Class 0 binary protostars formed by disk fragmentation to have typical separations of $s \sim 75$ AU (Tobin et al. 2013, 2016). After fragmentation, the orbital separation can decrease due to viscous dissipation or three-body dynamics (e.g. Moe & Kratter 2018) or increase, e.g. following accretion of high-angular momentum gas. Our results imply that systems formed by disk fragmentation constitute a significant fraction of the binary population out to $a \sim 200$ AU, with wider systems forming primarily through turbulent core fragmentation. Other lines of evidence also point toward a change in the binary formation mechanism at $a \sim 200$ AU. The masses of the components of solar-type binaries are correlated and inconsistent with random pairings from the IMF out to separations of 200 AU (Moe & Di Stefano 2017a). Similar trends with separation are found for the correlation in accretion rates of binary protostars (White & Ghez 2001) and the mutual inclination of orbits in hierarchical triples (Tokovinin 2017a).

It remains important to systematically measure the metallicity dependence of the binary fraction at intermediate separations ($10 \lesssim a/\text{AU} \lesssim 50$), between the regimes probed by close binaries (e.g. Badenes et al. 2018; Moe et al. 2018) and our study. Some previous works have found metallicity-dependence in the binary fraction in this regime to be weak (Raghavan et al. 2010; Rastegaev 2010), and others have found evidence of a *positive* correlation with metallicity, at least for low-mass binaries (Riaz et al. 2008; Jao et al. 2009; Lodieu et al. 2009; Ziegler et al. 2015). However, the common strategy of photometrically selecting metal-poor subdwarfs from below the main sequence for follow-up imaging can lead to a bias against binaries (Moe et al. 2018). While speckle interferometry and AO + *HST* imaging likely represent the most promising route to probing the binary fraction at $10 \lesssim s/\text{AU} \lesssim 50$, studies based on spectroscopically selected metal-poor samples would be less prone to biases against binaries. Constraints on the wobble of astrometric binaries in future *Gaia* data releases will also probe the intermediate-separation regime.

Acknowledgements

We thank Max Moe and Carles Badenes for helpful discussions. KE was supported by the NSF GRFP. This project was developed in part at the 2018 NYC Gaia Sprint, hosted by the Center for Computational Astrophysics of the Flatiron Institute in NYC, and in part at the workshop “Dynamics of the Milky Way System in the Era of Gaia,” hosted at the Aspen Center for Physics, which is supported by NSF grant PHY-1607611. This work has made use of data from the ESA *Gaia* mission, processed by the *Gaia* Data Processing and Analysis Consortium (DPAC).

The GALAH survey is based on observations made at the Australian Astronomical Observatory, under programmes A/2013B/13, A/2014A/25, A/2015A/19, A/2017A/18. We acknowledge the traditional owners of the land on which the AAT stands, the Gamilaraay people, and pay our respects to elders past and present. Funding for the Sloan Digital Sky Survey IV has been provided by the Alfred P. Sloan Foundation, the U.S. Department of Energy Office of Science, and the Participating Institutions. SDSS-IV acknowledges support and resources from the Center for High-Performance Computing at the University of Utah. The SDSS web site is www.sdss.org. Funding for RAVE (www.rave-survey.org) has been provided by institutions of the RAVE participants and by their national funding agencies. Guoshoujing Telescope (the Large Sky Area Multi-Object Fiber Spectroscopic Telescope LAMOST) is a National Major Scientific Project built by the Chinese Academy of Sciences. Funding for the project has been provided by the National Development and Reform Commission. LAMOST is operated and managed by the National Astronomical Observatories, Chinese Academy of Sciences. The research shown here acknowledges use of the Hypatia Catalog Database, an online compilation of stellar abundance data as described in Hinkel et al. (2014, *AJ*, 148, 54), which was supported by NASA’s Nexus for Exoplanet System Science (NExSS) research coordination network and the Vanderbilt Initiative in Data-Intensive Astrophysics (VIDA).

Chapter 6

The geometric challenge of testing gravity with wide binaries

An earlier version of this article was previously published in El-Badry, K., 2019, MNRAS, 482, 5018.

6.1 Abstract

Wide binaries provide promising laboratories for testing general relativity (GR) in the low-acceleration regime. Recent observational studies have found that the difference in the proper motions and/or radial velocities of the components of nearby wide binaries appear larger than predicted by Kepler's law's, indicating a potential breakdown of GR at low accelerations. These studies have not accounted for projection effects owing to the different position of the two stars on the celestial sphere. I show that two stars in a wide binary with *identical* 3D space velocities often have significantly different proper motions and radial velocities purely due to projection effects. I construct a sample of simulated binaries that follow Kepler's laws and have similar phase-space distributions to the observed samples of nearby binaries. Beyond separations of ~ 0.1 pc, direct comparison of the components' proper motions would suggest strong tensions with GR, even though the simulated binaries follow Kepler's laws by construction. The magnitude of the apparent disagreement is similar to that found observationally, suggesting that the apparent tension between observations and GR may largely be due to projection effects. I discuss prospects for constraining gravity at low accelerations with wide binaries. Robust tests of GR are possible with current data but require measurements of 3D velocities. Further work is also needed to model contamination from unbound moving groups and unrecognized hierarchical triples.

6.2 Introduction

Wide binaries provide an intriguing testbed for modified theories of gravity that predict deviations from general relativity (GR) in the Newtonian limit at very low accelerations. Many such theories have been proposed to explain observations on galactic scales, potentially alleviating the need for dark matter to explain observations of galactic dynamics (e.g. [Famaey & McGaugh 2012](#)). Testing modified gravity on galactic scales is challenging because many aspects of the galaxy formation process remain imperfectly understood. In the idealized case, the orbits of wide binaries provide a less complicated test, as the two-body problem has a simple solution both in GR and in many modified gravity theories ([Zhao et al. 2010](#)).

Modified gravity theories designed to explain galactic dynamics typically deviate from the Newtonian limit of GR at accelerations below the scale $a_0 \approx 10^{-8} \text{ cm s}^{-2}$. This is comparable to the acceleration in a solar-type binary with separation 10,000 AU (0.05 pc). At wider separations, GR and modified gravity theories predict different relations between the physical separation and orbital velocity of two stars in a gravitationally bound binary. The orbital timescales of such wide binaries are long ($\gtrsim 1$ Myr), so their full 3D separation and orbital velocity are generally not directly observable. However, for a statistical sample of binaries observed at random snapshots in their orbits, GR and modified gravity theories predict different relations between the projected physical separation, s , and the instantaneous one-dimensional velocity difference between the two stars (typically measured with proper motions and/or radial velocities), ΔV . The magnitude of the difference compared to GR varies significantly between modified gravity theories and is much larger in theories that do not include an external field effect than in theories that do ([Banik & Zhao 2018](#); [Pittordis & Sutherland 2018](#)). GR (in the Newtonian limit) predicts $\Delta V \sim s^{-1/2}$; at $s \sim 1$ pc, it predicts bound solar-type binaries to have ΔV of only a few tens of m s^{-1} . Measurements of binaries with wide separations and large ΔV can in principle rule out both GR and modified gravity theories that include an external field effect ([Famaey & McGaugh 2012](#); [Pittordis & Sutherland 2018](#)).

Recently, [Hernandez et al. \(2018\)](#) used astrometry from *Gaia* for a sample of wide binary candidates with projected separations $0.01 \lesssim s/\text{pc} \lesssim 10$ to measure the relation between ΔV and s , with the aim of testing classical gravity at low accelerations. They found values of ΔV that are substantially larger than the Newtonian prediction at large separations ($s \gtrsim 0.1$ pc) and interpreted them as evidence for a possible breakdown of GR at low accelerations. Their binary catalog contains 83 wide binary candidates within ~ 120 pc of the Sun that were originally identified by [Shaya & Olling \(2011\)](#) using astrometry from the *Hipparcos* and *Tycho-2* catalogs and were classified as unlikely to be chance alignments of stars that are not physically associated. The second *Gaia* data release significantly improved the precision of proper motions and parallaxes for most stars observed by *Hipparcos*, making it possible to measure sky-projected velocities for nearby stars with $\ll 0.1 \text{ km s}^{-1}$ precision ([Gaia Collaboration et al. 2018b](#)). As a result, the formal constraints on ΔV obtained by [Hernandez et al. \(2018\)](#) are quite strong.

Prior to the *Gaia* mission, two other studies considered the relation between ΔV and s




Figure 6.1: Distributions of heliocentric distance (left), total proper motion ($\mu_{\text{tot}} = \sqrt{\mu_{\alpha}^2 \cos^2 \delta + \mu_{\delta}^2}$; middle) and radial velocity (right). The observed binary sample from [Hernandez et al. \(2018\)](#) is shown in black; the simulated sample described in Section 6.3 is shown in red. The simulated sample is constructed to have a similar distribution of distance and proper motion to the observed sample.

as a potential constraint on the force law at low accelerations. [Hernandez et al. \(2012\)](#) used proper motions of wide binary candidates from the SDSS “SLoWPoKES” catalog ([Dhital et al. 2010](#)), as well as those of some of the wide binary candidates in the [Shaya & Olling \(2011\)](#) catalog, and reached qualitatively similar conclusions to [Hernandez et al. \(2018\)](#), but with larger uncertainties due to less precise astrometry. [Scarpa et al. \(2017\)](#) compared the radial velocities (RVs) of the components of some of the nearby binaries studied by [Hernandez et al. \(2012\)](#). They found that a large fraction had RV differences larger than the naive prediction for Keplerian orbits, indicating either a breakdown in classical gravity or that the binary candidates are not actually gravitationally bound.

In this paper, I draw attention to a potential complication in measuring the true velocity differences of the components of wide binaries. Direct comparison of RVs and proper motions of binary components entails the projection of velocity vectors in spherical coordinates onto a local Cartesian frame centered on each component. This will yield correct velocity differences in the limit where the two components are at the same position on the sky. However, as the angular separation of binary components grows, Cartesian planes that are normal to the unit sphere at the position of each component become rotated with respect to each other, leading to projection effects (e.g. [Shaya & Olling 2011](#)). These effects can cause two stars with very similar or identical 3D space velocities to have significantly different proper motions and radial velocities. Left unaccounted for, projection effects can lead to apparent tension with GR at wide separations, even in the absence of any true deviation from Keplerian orbits. In Section 6.3 below, I describe simulations to quantify the magnitude of this effect.

6.3 Simulations

I construct a sample of simulated binaries with the goal of comparing to the sample studied by [Hernandez et al. \(2018\)](#). I first sample positions for the center of mass of each binary assuming a uniform spatial distribution. I then reject a random subset of the simulated binaries such that their distribution of heliocentric distance is similar to that of the observed sample. Center-of-mass velocities for each simulated binary are drawn from a 3D Gaussian with $\sigma_{1D} = 25 \text{ km s}^{-1}$. This value of σ_{1D} is comparable to that measured in the solar neighborhood ([Sharma et al. 2014](#)) and is also equal to the dispersion in RV for the observed sample of binaries that have measured RVs. In [Figure 6.1](#), I compare the distributions of distance, proper motion, and RV of the simulated binary sample to those of the observed sample.

For each simulated binary, I draw an eccentricity from a uniform distribution $e \in [0, 1]$, a primary mass from a uniform distribution of $m_1 \in [0.5M_{\odot}, 0.6M_{\odot}]$, and a mass ratio from a uniform distribution of $q = m_2/m_1 \in [0.8, 1]$. These choices are designed to yield a binary population similar to the observed sample. I assume random orbital orientations, corresponding to a $p(i) di = \sin(i) di$ distribution of inclinations. Periods are drawn from a log-uniform distribution. This is not realistic ([Duchêne & Kraus 2013](#)), but the period distribution is not important for this study, which aims only to measure the enhancement of ΔV at a particular separation. Finally, each binary is mock observed at a random time $t_{\text{obs}} \in [0, P]$.

At time t_{obs} , I calculate the celestial coordinates of each component star, and θ , the angular separation between them. For a binary with center-of-mass distance d , the projected separation is $s = d \times \theta$, and the projected physical velocity difference in each proper motion component is

$$\frac{\Delta V_i}{\text{km s}^{-1}} = 4.74 \times 10^{-3} \frac{\Delta \mu_i}{\text{mas yr}^{-1}} \times \frac{d}{\text{pc}}. \quad (6.1)$$

Here $\Delta \mu_i$ represents the difference in each proper motion component ($\Delta \mu_{\alpha} \cos(\delta)$ for right ascension; $\Delta \mu_{\delta}$ for declination) between the two components of a binary. I also calculate the RV of each component, and the true velocity difference between the components of each binary along the three Cartesian axes. Transformations between coordinate frames are carried out using the `astropy.coordinates` package ([Astropy Collaboration et al. 2018a](#)).

6.3.1 Projection effects

[Figure 6.2](#) compares the true one-dimensional ΔV to the value obtained by directly comparing the proper motions of the two components (Equation 6.1). Each binary contributes two points, as proper motions in the RA and Dec components are considered independently. At small separations, projection effects are negligible, and the median ΔV computed from Equation 6.1 is identical to that of the true 1D velocity difference computed from the Cartesian velocity components. However, at $s \gtrsim 0.1 \text{ pc}$ (corresponding to an angular separation of




Figure 6.2: One-dimensional velocity difference versus projected separation for simulated binaries. In the top panel, red points show the true velocity difference, while black points show the result of directly comparing proper motion components; i.e. applying Equation 6.1. Solid and dashed lines show the median ΔV in bins of projected separation. At $s \gtrsim 0.1$ pc, projection effects become important, and the value of ΔV obtained by direct comparison of proper motion components significantly exceeds the true velocity difference. In the bottom panel, I remove binaries in which $\Delta V > 4$ km s $^{-1}$ in any component, as such pairs are unlikely to be classified as binaries in the first place. Blue points and error bars show the binned median and middle 68% of ΔV for the observed binary sample from Hernandez et al. (2012). The increase in ΔV at large separations is quite similar to that predicted to result from projection effects in the simulations.

$\theta \gtrsim 10$ arcminutes for the distance distribution of our sample), the true and projected values of ΔV begin to diverge widely. At a separation of 1 pc (10 pc), the typical enhancement in ΔV due to projection effects is 0.5 km s^{-1} (2 km s^{-1}).

The bottom panel of Figure 6.2 compares the values of ΔV predicted by the simulation to the observed binary sample from Hernandez et al. (2018). Here I have removed binaries in which the RVs or either proper motion component of the two stars differ by more than 4 km s^{-1} ; such pairs are unlikely to be classified as genuine binaries.¹ I plot the median and 1σ (middle 68%) range of ΔV for observed binaries in each bin of projected separation (note that this is not identical to the rms velocity in each bin plotted by Hernandez et al. 2018). Values of ΔV for the observed binaries are generally in good agreement with the simulation, even at large separations where they disagree substantially from the prediction for bound Keplerian orbits. This agreement suggests that the tension between the observed sample and the classical gravity prediction at large separations may in large part be a consequence of projection effects.

Although it is not shown in Figure 6.2, I find that projection effects enhance ΔV in the RV component in a manner essentially identical to the enhancement in proper motion difference.

6.3.2 Correcting for projection effects

With measurements of both radial velocity and proper motion, projection effects can be straightforwardly corrected for. The most direct approach is to transform the velocities of both components into a Cartesian frame and compare each of the velocity components in that frame. With *Gaia* data for nearby binaries, it is typically the case that proper motion uncertainties are much smaller than RV uncertainties. This fact can be exploited by transforming velocities for both stars into a Cartesian frame centered on one of the components and normal to the celestial sphere (i.e., with the \hat{z} axis pointing toward the Sun). In this case, a high-precision comparison of velocities in the \hat{x} and \hat{y} components is possible even with relatively low-precision RVs.

This is demonstrated in Figure 6.3, which shows the ΔV predicted for simulated binaries after transforming to a Cartesian frame centered on one component. I assume a proper motion uncertainty of 5 mas yr^{-1} in both components in all cases (adding Gaussian noise to the proper motions during mock observations) and vary the available RV information in different panels. To correct for projection effects, I chose one component of each binary at

¹This is similar but not identical to the cuts used in constructing the observed binary sample studied by Hernandez et al. (2018). Their sample was constructed from the Shaya & Olling (2011) binary catalog, after removal of pairs in which both components have RVs measured by *Gaia* that differ by more than 4 km s^{-1} , as well as pairs in which Shaya & Olling (2011) estimated the probability that the system is a chance alignment to be greater than 10%.

Shaya & Olling (2011) did attempt to account for geometrically-induced proper motion differences in constructing their catalog. Because Hernandez et al. (2012) used measurements of $\Delta\mu$ from Shaya & Olling (2011) for part of their sample, these measurements may be less affected by projection effects than those in Hernandez et al. (2018).




Figure 6.3: Predicted one-dimensional ΔV recovered from precise proper motions when projection effects are corrected for (see Section 6.3.2). In the top two panels, I assume RVs are available for both components. The third panel assumes RVs available for only one component, and the fourth panel assumes that no RVs are available for either component. Correcting for projection effects ameliorates the enhancement in ΔV at large separations compared to the direct comparison of proper motions in all cases, but some RV information is required to recover accurate velocities in the separation regime where alternative gravity theories predict strong deviations from GR.

random and transform the 3D velocities of both stars to a Cartesian frame that is centered on that component with \hat{z} axis pointed toward the Sun. I calculate ΔV separately for the \hat{x} and \hat{y} components in this frame. If no RVs are available, I assume $\text{RV} = 0 \text{ km s}^{-1}$ for both components. If an RV is measured for one component only, I assume both component have that same RV.

When no RVs are available (bottom panel), transforming to a Cartesian frame reduces the enhancement of ΔV due to projection effects by a factor of ~ 2 on average, but projection effects still dominate at $s \gtrsim 0.2 \text{ pc}$. Inclusion of RVs for both components makes the geometric correction effective to typical separations of 1-3 pc, depending on the precision of the RVs (top two panels). Obtaining a precise RV for one component and assuming that the 2nd component has the same RV is also effective at $s \lesssim 1 \text{ pc}$.

6.4 Discussion

I have shown that projection effects can lead to apparent disagreement in the proper motions and RVs of the components of wide binaries, even when the 3D space velocities follow Keplerian orbits. The magnitude of the disagreement caused by projection effects is comparable to the disagreement found observationally (Scarpa et al. 2017; Hernandez et al. 2018). It is of course possible that the observed wide binary candidates really do have discrepant 3D velocities, even after geometric distortions are corrected for. Here I only argue that consistency with Keplerian orbits is not ruled out.

I have considered only geometric complications to attempts to measure the velocity difference between the two components of a binary. Below, I briefly highlight two additional challenges to tests of gravity with wide binaries, and their potential resolutions.

1. *Unbound comoving pairs*: There is no guarantee that wide pairs of comoving stars with consistent astrometry are gravitationally bound. Recent studies (e.g. Oh et al. 2017a; Simpson et al. 2018; Faherty et al. 2018) have identified large numbers of unbound “moving groups”; i.e., associations of stars that are likely not gravitationally bound but follow very similar orbits through the Galaxy. Moving groups originate from dissolving star clusters that drift apart on $\sim 100 \text{ Myr}$ timescales. Their components have typical separations of a few pc and are essentially indistinguishable from genuinely bound binaries at large separations.

Failure to remove unbound pairs will result in an artificial enhancement of ΔV at large separations. Unbound pairs are more common at large separations, where the number of genuine binaries decreases, the phase-space volume for potential companions increases, and the binding energy of genuine binaries decreases. At separations exceeding the Jacobi limit, $r_J \sim 1.7 \text{ pc}$, the Galactic tidal field is stronger than the gravitational attraction between the components of solar-type binaries (Binney & Tremaine 2008a), so essentially no binaries at larger separations are expected to be genuinely bound. Several previous studies have attempted to account for the presence of unbound pairs

by comparing to the predictions of [Jiang & Tremaine \(2010\)](#), who include in their calculations disrupted, unbound binaries in the process of drifting gradually apart. But because [Jiang & Tremaine \(2010\)](#) model only co-moving pairs that were initially gravitationally bound binaries – i.e., their calculations do not attempt to account for the fact that most stars are born in clusters that gradually dissolve – their prediction at large separations likely underestimates the number of unbound comoving pairs with velocities that differ by a few km s^{-1} .

At the widest separations where previous studies have attempted to use wide binaries as probes of gravity, it is highly improbable that the identified pairs are binaries in any useful dynamical sense. At a separation of 8 pc, typical orbital periods are ~ 1.5 Gyr, many times larger than the Galactic dynamical time! It is therefore advisable to use binaries with separations $s < 1$ pc. The rate of contamination from dissolving clusters can also be substantially decreased by targeting binaries on halo-like orbits, which are primarily old and free from their birth associations.

2. *Hierarchical triples*: nearly half all wide binaries are really hierarchical triples and higher-order multiples ([Tokovinin et al. 2006](#); [El-Badry & Rix 2018b](#)) in which the additionally components are either faint or unresolved. These systems are generically not expected to follow the Keplerian prediction of $\Delta V \sim s^{-1/2}$, because the dynamics of one component are dominated by its closer companion. Of order half of such systems can be identified as containing photometric or spectroscopic binaries ([El-Badry et al. 2018g](#)), but it is unavoidable that a substantial fraction of hierarchical triples will go undetected; namely, those in which the unseen companion is faint and the inner orbit is wide enough that little RV variation is expected on observable timescales. Robust tests of gravity with samples of wide binaries should therefore include a population model for unrecognized higher-order multiples.

Acknowledgements

I thank Xavier Hernandez, Andrian Price-Whelan, Eliot Quataert, and Daniel Weisz for useful discussions and am grateful to the anonymous referee for a constructive report. I acknowledge support from an NSF graduate research fellowship.

Chapter 7

Discovery of an equal-mass “twin” binary population reaching 1000+ AU separations

An earlier version of this article was previously published in El-Badry, K., Rix, H.-W., Tian, H., Duchêne, G., Moe, M., 2019, MNRAS, 489, 5822.

7.1 Abstract

We use a homogeneous catalog of 42,000 main-sequence wide binaries identified by *Gaia* to measure the mass ratio distribution, $p(q)$, of binaries with primary masses $0.1 < M_1/M_\odot < 2.5$, mass ratios $0.1 \lesssim q < 1$, and separations $50 < s/\text{AU} < 50,000$. A well-understood selection function allows us to constrain $p(q)$ in 35 independent bins of primary mass and separation, with hundreds to thousands of binaries in each bin. Our investigation reveals a sharp excess of equal-mass “twin” binaries that is statistically significant out to separations of 1,000 to 10,000 AU, depending on primary mass. The excess is narrow: a steep increase in $p(q)$ at $0.95 \lesssim q < 1$, with no significant excess at $q \lesssim 0.95$. A range of tests confirm the signal is real, not a data artifact or selection effect. Combining the *Gaia* constraints with those from close binaries, we show that the twin excess decreases with increasing separation, but its width ($q \gtrsim 0.95$) is constant over $0.01 < a/\text{AU} < 10,000$. The wide twin population would be difficult to explain if the components of all wide binaries formed via core fragmentation, which is not expected to produce strongly correlated component masses. We conjecture that wide twins formed at closer separations ($a \lesssim 100$ AU), likely via accretion from circumbinary disks, and were subsequently widened by dynamical interactions in their birth environments. The separation-dependence of the twin excess then constrains the efficiency of dynamical widening and disruption of binaries in young clusters. We also constrain $p(q)$ across $0.1 \lesssim q < 1$. Besides changes in the twin fraction, $p(q)$ is independent of separation at fixed primary mass over $100 \lesssim s/\text{AU} < 50,000$. It is flatter than expected for random pairings from the IMF

but more bottom-heavy for wide binaries than for binaries with $a \lesssim 100$ AU.

7.2 Introduction

Binary stars are ubiquitous: roughly half of all field stars have binary companions (Duquennoy & Mayor 1991; Duchêne & Kraus 2013; Moe & Di Stefano 2017a), and the binary fraction is even higher in star forming regions (e.g. Ghez et al. 1993; Leinert et al. 1993; Mathieu 1994; Connelley et al. 2008; Sadavoy & Stahler 2017; Duchêne et al. 2018). Significant progress has been made in developing theoretical models to explain the population statistics of observed binaries (e.g. Kroupa 1995; Marks et al. 2011; Bate 2012; Lomax et al. 2015), but fundamental aspects of the binary formation process remain imperfectly understood.

The distribution of binary mass ratios has been a subject of interest for at least a century (e.g. van Biesbroeck 1916; Öpik 1924; Kuiper 1935). As a final outcome of the binary formation process, the mass ratio distribution provides useful constraints on theoretical models of star formation. Unlike the distributions of orbital separation and eccentricity, the mass ratio distribution has been suggested to be insensitive to dynamical evolution after formation (such that binaries of different mass ratios are disrupted at similar rates; e.g. Parker & Reggiani 2013). Mapping the mass ratio distribution over a range of binary masses and separations has thus been the focus of many studies (e.g. Trimble 1974, 1987, 1990; Eggleton et al. 1989; Mazeh et al. 1992; Hogeveen 1992; Shatsky & Tokovinin 2002; Burgasser 2007; Söderhjelm 2007; Raghavan et al. 2010; Tokovinin 2014a; Gullikson et al. 2016).

Observational studies of the mass ratio distribution are complicated by incompleteness. All binary detection methods are biased against low-mass ratio companions, which produce weaker radial-velocity shifts at fixed separation (e.g. Shahaf & Mazeh 2019), contribute less light to the observed spectra of unresolved binaries (e.g. El-Badry et al. 2018c,g), cause weaker eclipses (e.g. Moe & Di Stefano 2013), and are less likely to be detected as visual companions (e.g. Tokovinin 2011; El-Badry & Rix 2018c). The detection efficiency also varies with primary mass and orbital separation. This complicates measurement of the mass ratio distribution, because the distributions of separation, primary mass, and mass ratio are not independent (e.g. Moe & Di Stefano 2017a). It is thus important for demographic studies of binaries that the selection function of observed samples is well understood. If possible, binaries of different masses and physical separations should be considered independently.

A puzzling feature of the mass ratio distribution identified by previous works is the so-called “twin” phenomenon, which refers to a purported statistical excess of nearly equal-mass binaries with mass ratios $0.95 \lesssim q < 1$. Most studies that find an excess of equal-mass twins have focused on spectroscopic binaries with close separations ($a \ll 1$ AU, e.g. Lucy & Ricco 1979; Hogeveen 1992; Tokovinin 2000; Halbwachs et al. 2003; Lucy 2006; Pinsonneault & Stanek 2006; Simon & Obbie 2009; Kounkel et al. 2019). Indeed, several studies have reported a sharp drop-off in the fractional excess of twin binaries beyond periods of 40 days ($a \lesssim 0.2$ AU; e.g. Lucy & Ricco 1979; Tokovinin 2000; Simon & Obbie 2009). Recent studies

have not confirmed such a sharp drop-off in the twin excess for solar-type stars (Tokovinin 2014a; Moe & Di Stefano 2017a), but have still found it to decline monotonically with increasing separation. At least for $M_1 \gtrsim 1M_\odot$ (the mass range on which most previous studies have focused), the twin excess has been found to decrease with increasing primary mass at fixed separation and to extend to wider separations for lower-mass primaries (Moe & Di Stefano 2017a).

Some works have also argued that the twin phenomenon may be a selection effect (Mazeh et al. 2003; Cantrell & Dougan 2014), since equal-mass binaries are brighter, can be detected at larger distances, and are preferentially selected by several binary detection methods (e.g. Branch 1976). Such biases are minimized for binary samples that are nearly volume-complete (e.g. Raghavan et al. 2010; Tokovinin 2014a) and/or have well-understood completeness.

If it is a real effect, the physical origin of the twin phenomenon is not fully understood. Several mechanisms have been proposed that could lead to preferential formation of equal-mass binaries, including fragmentation during the late stages of protostellar collapse, mass transfer between pre-main sequence stars, and competitive accretion (see Tokovinin 2000 for discussion of different formation mechanisms). Several simulations have predicted that accretion of high-angular momentum gas, particularly from a circumbinary disk, tends to drive binary mass ratios toward unity (e.g. Bate 2000; Farris et al. 2014; Young & Clarke 2015). However, it is not obvious why, when averaged over a large population of binaries, accretion from a circumbinary disk would produce a sharp peak in the mass ratio distribution at $q \gtrsim 0.95$ as opposed to a gradual increase.

There have also been hints of an excess of twins among spatially resolved wide binaries with separations ranging from tens to thousands of AU (Trimble 1987; Söderhjelm 2007). The selection functions of the wide binary samples studied in these works were poorly understood, causing investigators to remain agnostic of whether the excess of equal-brightness pairs in their catalogs was of astrophysical origin or rooted in selection biases. Recently, Moe & Di Stefano (2017a) measured the twin excess at different separations in a small but relatively complete sample of solar-type binaries within 25 pc of the Sun. They found the twin excess to decline with separation, but found it inconsistent with 0 out to separations of 200 AU. At even wider separations, they set an upper limit of $\sim 5\%$ on the excess twin fraction.

High-quality parallaxes and proper motions from the recent *Gaia* data releases (Gaia Collaboration et al. 2016, 2018b) have simplified the process of constructing samples of wide binaries with (a) little contamination from chance alignments and (b) a well-understood selection function. Using data from *Gaia* DR2, El-Badry & Rix (2018c, hereafter ER18) constructed a high-purity catalog of wide binaries within 200 pc of the Sun consisting mainly of AFGKM dwarfs. In this paper, we use a subset of that catalog to constrain the mass ratio distribution over a wide range of primary masses ($0.1 \lesssim M/M_\odot \lesssim 2.5$), mass ratios ($0.1 \lesssim q < 1$) and separations ($50 \lesssim s/\text{AU} < 50,000$). The large size of the catalog allows us to constrain $p(q)$ in narrow bins of primary mass and separation independently. This approach makes it possible to measure variation in $p(q)$ with mass and separation, and it minimizes the sensitivity of our results to imperfectly known inputs such as the initial mass function (IMF) and separation distribution.

A striking result of our investigation is the unambiguous evidence that twins are not purely a close-binary phenomenon: a significant excess of equal-mass ($q \gtrsim 0.95$) binaries persists out to separations as wide as 10,000 AU. We derive constraints on the excess fraction of twins and the width of the twin excess as a function of mass and separation. We also provide constraints on the full mass ratio distribution over $0.1 \lesssim q \leq 1$ in all bins.

The remainder of this paper is organized as follows. Section 7.3 describes the binary catalog and tests we have done to verify that the twin excess is real. In Section 7.4, we describe how we forward-model synthetic binary populations to fit for the intrinsic mass ratio distribution. Results of this fitting are presented in Section 7.5. In Section 7.6, we compare to previous work and discuss implications of our results for models of binary star formation and dynamical evolution. The appendices provide additional details about several aspects of the data and model. There we discuss sensitivity to the adopted parametric form of $p(q)$, (Appendix 7.8), systematic uncertainties in our model (Appendix 7.9), evidence for a twin excess in archival wide binary catalogs (Appendix 7.10), empirical determination of the selection function (Appendix 7.11), and validation of the adopted Galactic model and selection function (Appendix 7.12). Constraints on fitting parameters are tabulated in Appendix 7.14.

7.3 Data

Our primary analysis uses the binary catalog assembled in ER18. This catalog contains $\sim 55,000$ spatially resolved wide binaries with main sequence (MS), giant, and white dwarf components, and projected physical separations of $50 \lesssim s/\text{AU} < 50,000$. We refer to ER18 for a full description of the catalog’s contents. In brief, it was constructed by searching *Gaia* DR2 for nearby ($d < 200$ pc) pairs of stars whose positions, proper motions, and parallaxes are consistent with being gravitationally bound. Resolved higher-order multiples and suspected members of bound and dissolving clusters were removed. The catalog is designed to be pure but not complete: cuts on photometric and astrometric precision ensure that the contamination rate from chance alignments is low ($\ll 1\%$), but they also reduce the number of faint binaries.

We do not use the full catalog from ER18, but impose the following additional cuts:

- We only consider MS/MS binaries, removing binaries in which either component is suspected to be a white dwarf, subgiant, giant, or pre-main sequence star. We identify non-MS components from the color-magnitude diagram (CMD; see below).
- We require *both* components to have `parallax > 5`, `parallax_over_error > 20`, `phot_bp_mean_flux_over_error > 20`, and `phot_rp_mean_flux_over_error > 20`. ER18 used these same cuts for the primary, but used less stringent cuts for the secondary. Here we apply the same cuts to both components in order to symmetrize the selection function.

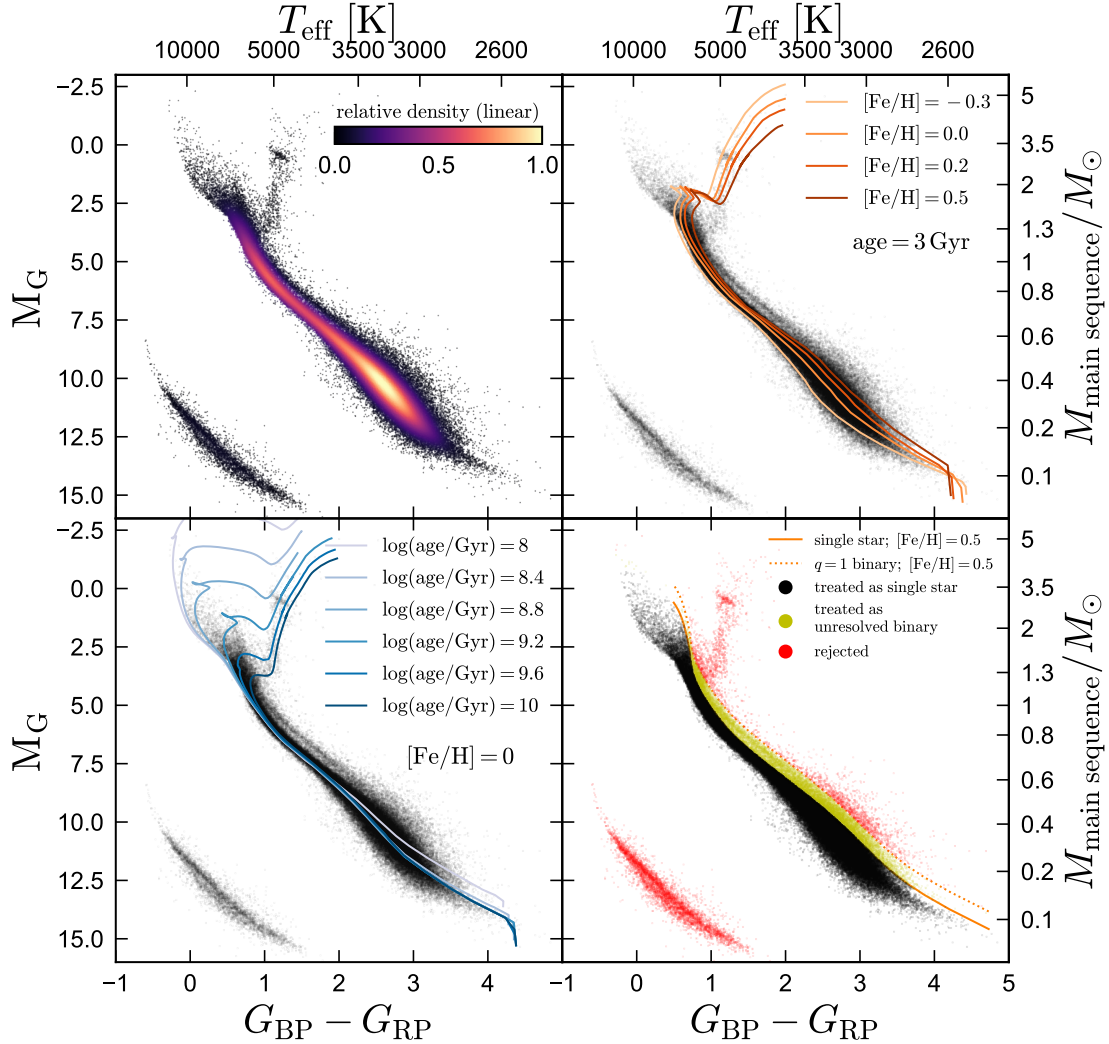


Figure 7.1: Color-magnitude diagram of all stars (both primaries and secondaries) in the ER18 catalog of resolved wide binaries. Right y-axes and upper x-axes show the mass and effective temperature corresponding to a given magnitude and color for main-sequence stars with $[\text{Fe}/\text{H}] = 0$. In the upper left panel, points are colored by a Gaussian kernel density estimate of the local density. We compare PARSEC isochrones with a range of metallicity (upper right) and age (lower left) to the data. The bottom-right panel divides the CMD into three regions. Black points (main-sequence components with no bright unresolved sub-components) comprise the bulk of our sample, and their masses are estimated using single-star isochrones. Yellow points (components of wide binaries with a bright unresolved companion) are kept, but their masses are estimated using unresolved binary isochrones. Red points (pre-main sequence stars, unresolved triples, and evolved stars) are discarded.

- We reject binaries in which the CMD-inferred mass of the primary falls outside $0.1 < M_1/M_\odot < 2.5$. The ER18 catalog contains fewer than 100 binaries with estimated primary masses above $2.5 M_\odot$, and none with estimated primary masses below $0.1 M_\odot$.

These additional cuts remove 23% of the ER18 catalog, leaving us with a sample of 42,338 MS/MS binaries.

Figure 7.1 shows the CMD of all stars in the ER18 catalog, with primaries and secondaries included on the same axes. Overplotted PARSEC isochrones show that the spread on the lower main sequence is primarily attributable to metallicity (upper right panel), while that on the upper main sequence and red giant branch is primarily due to age (lower left panel). A secondary sequence consisting mainly of unresolved binaries is visible above the main sequence. Wide “binaries” with one component in this sequence are primarily hierarchical triples. We do not remove these from our sample but account for them in our model when fitting for the mass ratio distribution in Section 7.4.¹

MS components that fall below a main sequence PARSEC isochrone with $[\text{Fe}/\text{H}] = 0.5$ (black points in the lower right panel of Figure 7.1) are treated as single stars. Those that fall above this isochrone but below an unresolved binary isochrone for two equal-mass stars with $[\text{Fe}/\text{H}] = 0.5$ (yellow points) are treated as unresolved binaries. Finally, sources that fall above this binary isochrone (red points; likely a mix of unresolved triples, pre-MS stars, and giants/subgiants) are rejected, as are white dwarfs. Of the 42,382 wide binaries in our catalog, 35,087 have two components consistent with having no unresolved sub-components (black points), and 7,295 have at least one component suspected to be an unresolved binary (yellow points).

We estimate masses for both components of each binary based on their location in the CMD. The ratios of these masses are not used explicitly in fitting the mass ratio distribution, but they are used to assign the primary vs. secondary components and to assign binaries to bins of primary mass. For MS stars that are suspected to be single, we estimate masses by interpolating from a grid of single-star PARSEC isochrones. This method is reasonably effective for single components, but it would yield biased results for the unresolved sub-components.

To estimate masses for components suspected to be unresolved binaries, we construct a population of synthetic unresolved binaries (see Section 7.4) and, for the subset of this population that falls in the region of the CMD colored in yellow in the bottom right panel of Figure 7.1, we calculate the median *primary* mass as a function of M_G of the unresolved binary. For observed sources in that region of the CMD, we estimate the primary mass by interpolating from M_G on this median relation. That is, the mass assigned to unresolved components represents the mass of the primary of the unresolved component, not the total mass.

This method of assigning masses is not without drawbacks: masses assigned to unresolved binaries are imprecise, because the mass ratio is not known. In addition, some pre-MS stars

¹We also experimented with removing binaries with suspected unresolved components from the sample; doing so does not qualitatively change any of our results.




Figure 7.2: SDSS images of a selection of binaries from our catalog. Top panel shows binaries in which the primary is a solar-type star ($0.9 < M_1/M_\odot < 1.1$) and the mass ratio varies from ≈ 0.1 to 1. The primary is at the center of each image, and the secondary is circled. Middle panel shows binaries in which the primary is a late K dwarf ($0.4 < M_1/M_\odot < 0.6$), again for a range of mass ratios. Bottom panels shows examples of “twin” binaries with near-identical components, each with $0.4 < M_1/M_\odot < 0.6$. Each image is 100 (top panel) or 45 (bottom panels) arcsec on a side. Our analysis uses photometry from *Gaia*, not SDSS. Because the SDSS photometry is ground-based, blending and source contamination affect it more severely.




Figure 7.3: Apparent G -band magnitude difference between the two components of binaries in our catalog as a function of the apparent magnitude of the primary (left), angular separation (middle), and projected physical separation (right). The sign of ΔG is randomized for easier visualization. There is a clear excess of equal-brightness binaries with $\Delta G \approx 0$. These “twin” binaries are found over a large range of apparent magnitudes, preferentially at closer physical and angular separations. The middle panel shows the contrast sensitivity limit for our sample; at fixed angular separation, the probability of a companion passing our photometric quality cuts drops rapidly outside this limit due to source contamination (Appendix 7.11.2). The lack of binaries with large ΔG at small separations is a selection effect; the narrow excess at $\Delta G \approx 0$ is not.

may be mistaken for unresolved binaries, and some low-mass ratio unresolved binaries may be mistaken for higher-metallicity single stars. However, we expect the typical accuracy to be $\lesssim 0.1M_{\odot}$, which is good enough for our purpose of assigning binaries to different bins of primary mass prior to fitting. In modeling the mass ratio distribution, it is not critical that the mass ratio of any one binary be measured accurately, but rather that the *distribution* of magnitude difference be predicted self-consistently.

Figure 7.2 shows “postage-stamp” images² of some of the binaries in our catalog from the Sloan Digital Sky Survey (SDSS; York et al. 2000). We note that SDSS photometry is not used in our analysis or in fitting the mass ratio distribution; we show it here because raw images from *Gaia* are not publicly available. Source contamination in the SDSS images is generally expected to be more severe than in *Gaia* photometry due to atmospheric seeing. To showcase the diversity of binaries in the catalog, we choose a selection of binaries with roughly solar-mass primaries and a range of mass ratios (top panel), binaries with primary masses of $\approx 0.5M_{\odot}$ and a range of mass ratios (middle panel), and some examples of equal-mass “twin” binaries with component masses of $\approx 0.5M_{\odot}$ (bottom panel). The twins are easily recognized because their magnitudes and colors are very similar.

An observable quantity that is closely related to the mass ratio is the difference in apparent magnitude of the two components, $\Delta G = |G_1 - G_2|$. Figure 7.3 shows the distribution of ΔG as a function of apparent magnitude of the primary, angular separation, and projected physical separation. The sign of ΔG is randomized, such that the distribution about $\Delta G = 0$

²skyserver.sdss.org/dr14/en/tools/chart/listinfo.aspx

is symmetric. This is helpful in making the equal-brightness population stand out, since it would otherwise be squished against the x-axis. An excess of equal-brightness binaries is evident over a wide range of apparent magnitudes; it is strongest at close separations. As we will show, the feature is quite narrow: the density of binaries is enhanced primarily at magnitude differences of $\Delta G < 0.25$ mag. This is much narrower than the range over which the selection function varies significantly: at $\theta > 2.5$ arcsec, the contrast sensitivity is basically constant over $0 < \Delta G < 2$ (see ER18; their Appendix A).

The middle panel of Figure 7.3 shows that there are no binaries with small separations and large ΔG in the catalog. This is a consequence of photometric contamination at small angular separations. The dashed line shows the contrast sensitivity limit derived in Appendix 7.11.2; this is the value of ΔG at which the sensitivity is 50% of its value at asymptotically large separations for a given θ . The contrast limit is derived from the correlation function of chance alignment sources subject to similar quality cuts as real binaries. The fall-off in sensitivity is quite steep beyond the contrast limit, leading to an envelope in $\Delta G(\theta)$ beyond which no binaries are found. The drop-off toward larger ΔG is less steep as a function of physical separation s , since binaries with similar s have different angular separations at different distances.

7.3.1 Is the twin excess real?

To test whether the narrow excess of binaries with $\Delta G \approx 0$ is a real astrophysical effect (as opposed to a data artifact), we produced a control sample of chance alignments with similar observable properties to the real binaries. This sample was produced by repeating the procedure used to produce the real binary catalog (applying the same quality cuts and limits on photometric and astrometric precision), but requiring that the two stars have parallaxes and proper motions that are *inconsistent* rather than consistent. This selects pairs of stars that are close together on the sky (and thus are affected by contamination and blending in the same way as real binaries) but are not physically associated. We applied the same procedure for removing resolved higher-order multiples and potential members of bound and dissolving clusters that was used for real binaries. Because chance alignments are intrinsically rare at close angular separations, we searched out to 400 pc (rather than 200 pc for the fiducial binary catalog) to obtain better statistics. We verified that our conclusions are unchanged when only the sample within 200 pc is considered.

In Figure 7.4, we compare the distributions of magnitude and color difference for real binaries in our catalog (left panels) and chance alignments (right panels). Chance alignments are more common at large angular separations. The broad distribution and outer envelope of ΔG and $\Delta(G_{\text{BP}} - G_{\text{RP}})$ at a given θ is similar for binaries and chance alignments, reflecting *Gaia*'s contrast sensitivity. For the chance alignments, there is no sharp excess of pairs with nearly-equal magnitude and color. Because chance alignments are subject to same cuts on astrometric and photometric quality and signal-to-noise as the real binaries, any aspects of the *Gaia* source detection algorithm that might be expected to produce a bias toward equal-brightness pairs should affect real binaries and chance alignments very similarly. We




Figure 7.4: Difference in G -band magnitude (top) and color (bottom) between the two components of pairs with a range of angular separations. We compare genuine binaries (left) to chance-alignments (right). The chance alignments are required to pass the same quality cuts as the true binaries. A clear excess of pairs with $\Delta G \approx 0$ and $\Delta(G_{\text{BP}} - G_{\text{RP}}) \approx 0$ is evident for the true binaries, but not for the chance alignments. The absence of the excess among chance alignments bolsters our confidence that the feature found in the real binaries is due to a true excess of equal-mass twins and is not a selection effect or data artifact.




Figure 7.5: Ratio of the number of binaries with nearly equal magnitudes ($\Delta G < 0.25$) to the number with slightly different magnitude ($0.25 < \Delta G < 0.5$) as a function of physical separation (left) and angular separation (right). This ratio quantifies the excess of equal-brightness “twins”. Error bars are 1 sigma. We compare binaries at different distances, as well as a control sample of chance-alignments. At fixed *physical* separation, the twin excess is consistent with being independent of distance. At fixed *angular* separation, it decreases with distance. Together, these trends strongly suggest that the $\Delta G \approx 0$ excess (Figures 7.3 and 7.4) reflects a real excess of equal-mass binaries, not a selection effect or photometric issue at close angular separations.

therefore interpret the lack of a thin excess at $\Delta G \approx 0$ and $\Delta(G_{\text{BP}} - G_{\text{RP}}) \approx 0$ for chance alignments as strong evidence that the twin excess among real binaries is astrophysical.

An initial worry was that apparent twins might be duplicate *Gaia* sources that were not properly removed: if a source were observed twice in two different scans without being identified as a duplicate, it could manifest in our catalog as an apparent binary pair in which the two components had essentially identical astrometry and photometry. We verified that unrecognized duplicate sources are *not* the source of the twin signal using the SDSS images: we visually inspected the postage stamps of several hundred equal-brightness binaries in the SDSS footprint, and all the systems indeed contain two stars.

Another useful test to verify that the $\Delta G \approx 0$ excess is physical is to determine whether its strength depends primarily on physical or angular separation. This can be accomplished by comparing the ΔG distributions of binaries at different distances. If the excess of equal-brightness pairs were due to an observational bias against binaries that are not nearly equal-mass (or if there were an issue with the *Gaia* photometry causing close pairs to erroneously have the same reported magnitude), one would expect the twin excess to depend on angular separation, manifesting itself at different *physical* separations for binaries at different distances. If the excess is due to an astrophysical preference for equal-mass binaries, then it should be primarily a function of s , manifesting itself at the same physical separation but

different θ for binary samples at different distances.³

In Figure 7.5, we compare the excess of equal-brightness pairs for binaries in three different distance bins as a function of physical separation (left) and angular separation (right). We measure the excess as the ratio of the number of binaries with $0 < \Delta G < 0.25$ (nearly-equal magnitudes) to the number with $0.25 < \Delta G < 0.5$ (slightly different, but still similar magnitudes). In order to include binaries at distance larger than 200 pc, we repeated the binary search from ER18, this time searching out to 400 pc. The additional binaries with $200 < d/\text{pc} < 400$ are considered for this test only and are not included in the sample used to fit for the mass ratio distribution.

The left panel of Figure 7.5 shows that at fixed *physical* separation, the fractional excess of twins increases toward closer separations in a manner that is consistent across the three different distance samples. The right panel shows that as expected, the twin excess at fixed *angular* separation varies with distance. At fixed θ , larger distances correspond to wider physical separation. Because the twin excess decreases with physical separation, it decreases with distance at fixed θ . The excess of equal-brightness chance alignments is also shown as a function of angular separation in the right panel of Figure 7.5. As expected, this is nearly consistent with 0 at all angular separations, meaning that there is no strong bias toward equal-brightness pairs compared to pairs with slightly different brightness. At the closest angular separations ($\theta \lesssim 5$ arcsec), there is a slight excess of equal-brightness chance alignments, suggesting that contrast sensitivity begins to play a role. The excess for the chance alignments is less than that found for the true binaries at all distances and is self-consistently accounted for in the selection function (Appendix 7.11.2).

Figure 7.6 shows distributions of ΔG for binaries in 5 different bins of primary mass. Both the overall shape of the distribution and the strength of the excess at $\Delta G \approx 0$ vary substantially between mass bins. For the highest-mass bin, $1.2 < M_1/M_\odot < 2.5$, the excess at $\Delta G \approx 0$ is weak, and the observed distribution (without accounting for incompleteness) peaks at $\Delta G \approx 8$, corresponding to $q \approx 0.3$ (see Figure 7.8). For $0.8 < M_1/M_\odot < 1.2$ and $0.6 < M_1/M_\odot < 0.8$, there is a clear excess of twins out to $s \approx 1000$ AU, with the observed peaks in the ΔG distribution corresponding to $q \approx 0.45$ and $q \approx 0.6$, respectively. In the two lowest-mass bins, the visible twin excess appears to extend to larger separations, and there is no secondary peak in the distributions of ΔG . For all mass bins, the binaries with the largest ΔG have secondaries with $M \approx 0.1M_\odot$. The maximum ΔG in each panel is probably set largely by observational incompleteness, since at moderately old ages ($\gg 100$ Myr), objects below the hydrogen burning limit are several magnitudes fainter than those just above it. Incompleteness at small separations due to the angular resolution and contrast sensitivity of our catalog becomes significant at different projected physical separations for different mass bins, because the low-mass binaries that pass our quality cuts are on average at closer distances than the higher-mass binaries.

³An implicit assumption here is that the fractional excess of twins is independent of the intrinsic properties of a binary (e.g. mass), since binaries that pass our quality cuts at larger distances will be more massive on average than those that are nearby. We show in Section 7.5 that this assumption does not hold up in detail, which could lead to modest variation with distance in the observed twin excess at fixed physical separation.




Figure 7.6: Magnitude difference as a function of projected physical separation for binaries with different primary masses. The sign of ΔG is randomized for easier visualization, and masses are computed from the CMD. The thin band of “twin” binaries with $\Delta G \approx 0$ is clearly visible in all but the highest-mass bin. In the top three bins, there is a clear preference for unequal-mass (large ΔG , low- q) binaries. Incompleteness to faint companions prevents the detection of large ΔG companions to low-mass secondaries.

7.4 Modeling

We now turn to the intrinsic mass ratio distribution, $p(q)$. Because we expect that $p(q)$ may vary with primary mass and/or separation, we split the observed sample into 5 bins of primary mass and 7 bins of projected physical separation, fitting each of the resulting 35 subsamples independently. Our choice of bins is designed to balance the number of binaries in each bin while still covering a large dynamic range of mass and separation. We constrain $p(q)$ for each subsample by forward-modeling a simulated population of binaries with a given distribution of primary mass, age, metallicity, distance, physical separation, and mass ratio, passing them through the selection function, and comparing to the data. The “data” we consider is the distribution of the observed binaries in the 3-dimensional space of angular separation θ , magnitude difference ΔG , and parallax ϖ . The best-fitting $p(q)$ is then the one that best matches the observed data, and uncertainties in $p(q)$ are estimated via Markov chain Monte Carlo from the range of $p(q)$ that adequately reproduce the observed data. This approach requires a parameterized form of $p(q)$ (Section 7.4.1), knowledge of the selection function (Section 7.4.2), and a Galactic model from which simulated binaries are drawn (Section 7.4.3).

To generate a model prediction for a given set of fitting parameters, we generate a population of $N = 10^6$ synthetic binaries and forward-model their distribution into the space of observables. Masses, ages, metallicities, and distances are Monte Carlo sampled from their respective distribution functions. For mass ratios and separations, we use a regular 1000×1000 grid, weighting the synthetic binary at each gridpoint by the mass ratio and separation distributions. This approach is chosen because in the fiducial model, we fit for the mass ratio and separations distributions but leave the distributions of age, mass, distance, and metallicity fixed. The number of synthetic binaries generated must be large enough that Poisson noise is negligible. We verified that $N = 10^6$ binaries in each mass and separation bin is large enough that our constraints are converged and insensitive to the random seed.

Synthetic photometry is calculated for both components from isochrones (including the effects of unresolved companions in hierarchies; see Section 7.4.3), and the observable properties of each binary are passed through the selection function. In constructing the distribution of mock-observables to be compared with the data, each synthetic binary is weighted by the selection function evaluated for its observables. Finally, the observed and simulated distributions are binned on a regular 3D grid. We default to using 100 bins in ΔG , 100 bins in $\log \theta$, and 5 bins in ϖ (because typical errors in ϖ are larger than those in ΔG or θ). The maximum value of the ΔG grid for a given primary mass bin is chosen to include the largest ΔG value for the data in that bin, and the 5 bins in ϖ are chosen so that roughly the same number of observed binaries fall in each bin. Our constraints are not sensitive to the choice of bins, since they are small compared to the scale on which the data exhibit substructure.

We re-scale the binned model prediction such that the total counts match the observed data. We calculate the likelihood for a particular set of model parameters by summing over all cells in the 3D distribution, assuming that the distribution of counts in each cell is set by

a Poisson process. The log-likelihood function is

$$\ln L = \sum_{m_i \neq 0} [d_i \ln m_i - m_i - \ln(d_i!)], \quad (7.1)$$

where m_i and d_i are the counts in the i th cell of the binned model and data (“!” denotes a factorial). We sample from the posterior using `emcee` (Foreman-Mackey et al. 2013a), using priors described in Section 7.4.1. We use 200 walkers and draw 20,000 samples for each bin of mass and separation after a burn-in period of 200 steps per walker. Inspecting the chains, we find this to be sufficient for convergence in all cases. We carried out tests with mock data that was drawn from a known mass ratio distribution assigned with realistic observational uncertainties in order to verify that our approach yields unbiased constraints on $p(q)$.

This fitting procedure is qualitatively very similar to the method commonly used to constrain population properties such as the IMF, star formation history, unresolved binary fraction, or initial-final mass relation from CMDs (e.g. Dolphin 2002; Bonatto et al. 2012; Geha et al. 2013; El-Badry et al. 2018d). The difference between our approach and these studies is that we are forward-modeling the distribution of binaries in the space of angular separation, magnitude difference, and parallax rather than color and magnitude. The approach can in principle be generalized to include other observables, such as color difference or apparent magnitude of the primary, but the computational expense increases rapidly with the dimension of the space in which the likelihood function is calculated.

We note that the data uncertainties do not enter Equation (7.1). The implicit assumption (which does hold for our problem setup) is that the uncertainties are small compared to the scale of the bins in all quantities. We also note that for a fine grid, a majority of grid cells will contain 0 or 1 real binaries. This is not a problem; Equation (7.1) does not make any assumptions about the magnitude of d_i .

7.4.1 Parameterization

We fit the mass ratio distribution by assuming a parameterized form of $p(q)$ and then obtaining constraints on the parameters. Our fiducial parameterization is motivated by the one used by Moe & Di Stefano (2017a) and is shown in Figure 7.7. $p(q)$ is parameterized as a broken power law with logarithmic slope $\gamma_{\text{small}q}$ at $q < q_{\text{break}}$ and $\gamma_{\text{large}q}$ at $q > q_{\text{break}}$. An possible excess (or deficit) of twins is added on top of the power law component at $q > q_{\text{twin}}$. This excess is modeled as a step function, with the magnitude such that the integrated excess of twins divided by the total number of binaries with $q > 0.3$ is F_{twin} . The reason for this choice (as opposed to normalizing relative to *all* binaries) is that, compared to the constraints at $q > 0.3$, the constraints on $p(q)$ at small q are often weak due to incompleteness. $p(q)$ is set to 0 at $q < 0.05$ to prevent divergence when $\gamma_{\text{small}q} < -1$. Because we are not sensitive to companions with $q < 0.05$, this has little effect on our results.

There is no a priori motivation for this particular parameterization: because the physics that set $p(q)$ are imperfectly understood, we simply require a functional form that is sufficiently flexible to reproduce the observed distributions of ΔG . We have experimented with




Figure 7.7: Parameterized mass ratio distribution. The broad part of the distribution is modeled with a broken power law of logarithmic slope $\gamma_{\text{small}q}$ at $q < q_{\text{break}}$ and $\gamma_{\text{large}q}$ at $q > q_{\text{break}}$. F_{twin} is the *excess* fraction of nearly equal-mass binaries with $q > q_{\text{twin}}$, relative to the underlying power-law distribution for $q > 0.3$. For this particular example, $F_{\text{twin}} = 0.04$, $q_{\text{twin}} = 0.95$, $\gamma_{\text{small}q} = 0.3$, $\gamma_{\text{large}q} = -1.3$, and $q_{\text{break}} = 0.5$.

Table 7.1: Summary of priors adopted in each primary mass bin. We use the same priors for all separation bins. $\mathcal{U}(a, b)$ represents a uniform distribution over $[a, b]$, and $\mathcal{N}(\mu, \sigma)$ represents a normal distribution with mean μ and dispersion σ . γ_s is the logarithmic slope of the local separation distribution; i.e., $p(s) \propto s^{\gamma_s}$. Other parameters are described in Figure 7.7.

	$0.1 < M_1/M_\odot < 0.4$	$0.4 < M_1/M_\odot < 0.6$	$0.6 < M_1/M_\odot < 0.8$	$0.8 < M_1/M_\odot < 1.2$	$1.2 < M_1/M_\odot < 2.5$
F_{twin}	$\mathcal{U}(-1, 1)$	$\mathcal{U}(-1, 1)$	$\mathcal{U}(-1, 1)$	$\mathcal{U}(-1, 1)$	$\mathcal{U}(-1, 1)$
q_{twin}	$\mathcal{U}(0.93, 1)$	$\mathcal{U}(0.93, 1)$	$\mathcal{U}(0.93, 1)$	$\mathcal{U}(0.93, 1)$	$\mathcal{U}(0.93, 1)$
γ_{largeq}	$\mathcal{N}(0.5, 1)$	$\mathcal{N}(-1, 1)$	$\mathcal{N}(-1, 1)$	$\mathcal{N}(-1, 1)$	$\mathcal{N}(-1, 1)$
γ_{smallq}	—	$\mathcal{N}(0.5, 0.5)$	$\mathcal{N}(0, 0.5)$	$\mathcal{N}(0, 0.5)$	$\mathcal{N}(0, 0.5)$
γ_s	$\mathcal{N}(-1.5, 1)$	$\mathcal{N}(-1.5, 1)$	$\mathcal{N}(-1.5, 1)$	$\mathcal{N}(-1.5, 1)$	$\mathcal{N}(-1.5, 1)$
q_{break}	—	0.5	0.5	0.5	0.3

other forms of $p(q)$, including adding an additional break point to the power law and modeling the twin excess as ramping up linearly instead of increasing stepwise. However, we find the form shown in Figure 7.7 to provide a good fit in all mass and separation bins, with more complicated models providing only marginal improvements. We thus use this functional form for our main analysis. In Appendix 7.8, we show results of using alternative parameterizations, including one that smoothly transitions between the two power laws instead of including a sharp break and one that leaves the shape of the twin excess flexible.

There are two differences between our parameterization of $p(q)$ and the one used by Moe & Di Stefano (2017a), who focused primarily on binaries with $M_1 \gtrsim M_\odot$. First, they fixed $q_{\text{break}} = 0.3$, a value that was appropriate for their analysis because several studies have found the mass ratio distribution to peak at $q \approx 0.2 - 0.3$ for binaries with $M_1 \gtrsim M_\odot$ (Duquennoy & Mayor 1991; Gullikson et al. 2016; Murphy et al. 2018). We also use $q_{\text{break}} = 0.3$ for our highest-mass bin, but we find $q_{\text{break}} = 0.5$ to provide a better fit at lower masses (see Appendix 7.8.2). Second, they fixed $q_{\text{twin}} = 0.95$, roughly the value found for spectroscopic binaries (e.g. Tokovinin 2000). In order to identify or rule out trends with mass and separation, we leave q_{twin} as a free parameter.

Along with the parameters of the mass ratio distribution, we also fit for γ_s , the local logarithmic slope of the separation distribution in each bin of primary mass and separation. Opik’s law (a uniform distribution of $\log(s)$) corresponds to $\gamma_s = 0$.

Our adopted priors are listed in Table 7.1. For the lowest-mass bin, the smooth component of $p(q)$ is modeled as a single power law, because no useful constraints can be obtained at $q \lesssim 0.5$. Our priors on γ_{smallq} and γ_{largeq} are loosely motivated by constraints from the literature (particularly Tokovinin 2014a and Moe & Di Stefano 2017a) but are fairly weak. We use a uniform flat prior on F_{twin} . We require $q_{\text{twin}} > 0.93$ in order to prevent cases where a large F_{twin} is combined with $q_{\text{twin}} \ll 1$ such that a broad “twin excess” simply modifies the overall shape of $p(q)$ but does not actually correspond to a sharp increase near $q = 1$. We show in Appendix 7.8.1 that, where a twin excess is significant, it is always narrow ($q_{\text{twin}} \gtrsim 0.94$).

7.4.2 Selection Function

Because both components of a binary must pass astrometric and photometric quality cuts, the binaries in our catalog are relatively bright. The median apparent magnitude of all stars in our fitting sample (considering primaries and secondaries together) is $\langle G \rangle = 14.5$, and 90% (99%) of stars fall in the range $9.1 < G < 17.4$ ($6.6 < G < 18.1$). For stars in this magnitude range, *Gaia* DR2 is nearly complete outside of crowded fields (Arenou et al. 2018a; Sollima 2019). The completeness is not quite 100% due to a variety of issues, but it is primarily a function of position on the sky, not color or magnitude. The selection function for our sample is thus determined by the cuts imposed on astrometric and photometric precision.

In order for a binary to appear in the catalog, (a) both components must be bright enough that they individually pass the cuts we impose on `parallax_over_error` and photometric signal to noise, and (b) they must not be so close on the sky that the photometry of either component is significantly contaminated. The selection function for binaries can thus be expressed as a product of the two components' single-star detection probabilities and a contrast sensitivity cross term:

$$s_{\text{binary}} = s_1 \times s_2 \times s_{\Delta G}(\theta). \quad (7.2)$$

Here s_1 and s_2 represent the independent probabilities of detecting an isolated star with the observable properties of star 1 and star 2; they depend primarily on apparent magnitude and color. $s_{\Delta G}(\theta)$ quantifies the reduction in the probability of detecting star 1 and star 2 together, *relative to the probability of detecting them at asymptotically large separation*. It is primarily a function of the angular separation of the two stars and their flux ratio. For example, $s_{\Delta G}(\theta) \approx 0$ at $\theta < 2$ arcsec and 1 at $\theta > 10$ arcsec; at intermediate separations, it depends strongly on ΔG . We calculate the single-star selection function given our quality cuts in Appendix 7.11.1 and $s_{\Delta G}(\theta)$ in Appendix 7.11.2. The derived selection functions are then validated in Appendix 7.12.

7.4.3 Model inputs

We draw primary masses and system ages assuming a Kroupa (2001b) IMF and a constant star formation history over the last 10 Gyr. Because suspected members of bound and dissolving clusters are removed from our binary catalog, we remove synthetic binaries with age < 100 Myr. We assume that the wide binary fraction, f_{wb} , scales with mass as $f_{\text{wb}} \propto M_1^{\alpha_{\text{wb}}}$, where $\alpha_{\text{wb}} = 0.4$ is a constant (i.e., higher-mass primaries are more likely to have wide binary companions). The effect of this assumption is that primary masses are drawn from a distribution with logarithmic slope $\alpha_{\text{IMF}} + \alpha_{\text{wb}}$, where α_{IMF} is the local logarithmic slope of the IMF ($\alpha_{\text{IMF}} = -2.3$ for $M_1/M_{\odot} > 0.5$ and $\alpha_{\text{IMF}} = -1.3$ for $M_1/M_{\odot} < 0.5$). We find that with this choice of α_{wb} , our model predicts a distribution of primary magnitudes in reasonably good agreement with that of the binary catalog when the selection function is taken into account.

We model the intrinsic spatial distribution of all stars as a plane-parallel exponential with the Sun at the midplane. The exponential scale height increases with stellar age (e.g.

(Nordström et al. 2004; Seabroke & Gilmore 2007), because older stars (a) were born from kinematically hotter gas and (b) have been dynamically heated more since their formation (e.g. Ting & Rix 2018). We use a fit to the empirical age-scale height relation recently measured by Sollima (2019) using *Gaia* star counts:

$$\log(h_z/\text{pc}) = 0.53 \log(\text{age}/\text{yr}) - 2.65. \quad (7.3)$$

I.e., the scale height increases from 40 pc for stars of age 100 Myr, to 130 pc at age 1 Gyr, to 450 pc at age 10 Gyr. We show in Appendix 7.12.1 that this leads to a predicted distance distribution in good agreement the data.

We use the tabulated empirical metallicity distribution function (MDF) for our binary catalog that was measured in El-Badry & Rix (2019b) by considering a subset of ~ 8000 binaries in the catalog for which $[\text{Fe}/\text{H}]$ was measured spectroscopically for at least one component. Most of the binaries in our catalog are disk stars, with a median metallicity of $\langle [\text{Fe}/\text{H}] \rangle \approx -0.1$ and tails extending to $[\text{Fe}/\text{H}] \approx -1.0$ and $[\text{Fe}/\text{H}] \approx +0.4$. We do not include any variation in the MDF with age or distance.

We generate synthetic photometry in the *Gaia* DR2 bands from Evans et al. (2018) using PARSEC⁴ isochrones (Bressan et al. 2012; Chen et al. 2014). Just as for the real data, we remove synthetic binaries in which either component has evolved off the main sequence. For companions with $M_2 < 0.1 M_\odot$, we supplement the PARSEC models with BT-Settl models for very low-mass stars and brown dwarfs (Allard et al. 2012; Allard 2014).⁵ We include companions with masses as low as $0.01 M_\odot$ in our model for completeness. However, these have very little effect on our results, because substellar companions are too faint and red to pass our quality cuts except at very young ages.

Figure 7.8 shows the relation between q and ΔG predicted for PARSEC isochrones. We show relations for 5 different primary masses, corresponding approximately to the median masses of primaries in each of our five mass bins, and use a minimum companion mass of $0.1 M_\odot$ in all cases. Comparing to Figure 7.6, it is evident that the lowest-mass secondaries have $M_2 \approx 0.1 M_\odot$ in all bins of primary mass. As for the excess of equal-brightness binaries, it is primarily manifest over $0 < \Delta G < 0.25$ (though there is some variation with mass; see Figure 7.10), which corresponds to $0.93 \lesssim q < 1$.

Some components of wide binaries have their own spatially-unresolved close companions. We incorporate these in our models following observed binary statistics. The probability that a star has a close companion is taken to be a function of its mass, increasing from 20% at $M < 0.2 M_\odot$, to 30% at $0.2 < M/M_\odot < 0.5$, to 35% at $0.5 < M/M_\odot < 0.8$, to 45% at $M/M_\odot > 0.8$ (Duchêne & Kraus 2013). For components that are assigned a binary companion, we draw the unresolved companion mass from a mass ratio distribution that is uniform between $q_{\min} = 0.1 M_\odot / M_1$ and 1. We assume that the two components' probabilities of having a close companion are independent and neglect dynamical stability

⁴<http://stev.oapd.inaf.it/cgi-bin/cmd>

⁵Synthetic photometry in *Gaia* DR2 bands is computed for the BT-Settl models using the Phoenix web simulator, available at phoenix.ens-lyon.fr/simulator-jsf22-26/index.faces.




Figure 7.8: Magnitude difference vs. mass ratio for binaries of different primary mass, computed from PARSEC isochrones. We assume both components have $[\text{Fe}/\text{H}] = 0$, adopting an age of 1 Gyr for the highest primary mass and 5 Gyr for the others. Because the mass-luminosity relation varies with mass, the relation between ΔG and q does as well. A magnitude difference of 0.25 mag (roughly where the twin excess becomes significant in our catalog) corresponds to a mass ratio of 0.93 to 0.97, depending on primary mass.

constraints. This prescription reproduces the morphology of the observed CMD reasonably well. In particular, the fraction of stars within 200 pc that fall in the “suspected unresolved binary” region of the CMD (yellow points in the bottom right panel of Figure 7.1) at a given M_G is reproduced within $\sim 25\%$ over $0.1 < M/M_\odot < 1.0$.

We do not include extinction or reddening due to dust in our model and do not attempt to correct the data for it. Because the stars in our catalog are nearby ($d < 200$ pc), the effects of extinction are expected to be modest. The morphology of the CMD (e.g. the compactness of the red clump and main sequence; see Figure 7.1) validates this assumption. Moreover, because the two components of a binary have similar position on the sky (within an arcminute in most cases) and similar distance, the extinction toward both components is expected to be similar.

We discuss the sensitivity of our constraints on the mass ratio distribution to various model ingredients in Appendix 7.9. The largest systematic uncertainties come from the choice of stellar models. All the systematics we consider primarily affect constraints at low q . This is also true for uncertainties in the completeness function: because the two components of binaries with $q \approx 1$ have similar magnitudes and colors, incompleteness affects them similarly. The translation between the distribution of ΔG and $p(q)$ is thus more straightforward at $q \approx 1$ than at $q \ll 1$.

7.5 Results

Our fitting produces samples from the posterior distribution of free model parameters for each bin of mass and separation. These samples translate to marginalized constraints on each parameter of the mass ratio distribution and the covariances between them. An example for a single mass and separation bin is shown in Figure 7.21. Marginalized constraints for all bins are listed in Appendix 7.14.

The 35 panels of Figure 7.9 show median and 2 sigma constraints on the mass ratio distribution for 35 bins of primary mass and separation. We do not show the y-axis ticks to avoid cluttering the figure, as the limits are different in each panel. The uncertainties (shaded regions) are derived by drawing 1000 samples from the posterior, calculating a normalized $p(q)$ for each sample, and then shading the middle 95.4% range of these samples at each q . Solid black lines show the median at each q . Because $p(q)$ is normalized, uncertainty in the mass ratio distribution at *any* q translates to uncertainty in the normalization of $p(q)$ at *all* q . This reflects the fact that it is impossible to know the total fraction of all binaries that fall in some mass ratio range if one does not know how many binaries with low mass ratios are missed. However, $\gamma_{\text{large}q}$, the slope of the mass ratio distribution at $q > q_{\text{break}}$, is usually well-constrained. Dashed vertical lines in each panel show the lowest-mass ratio observed binary included in that bin. This roughly corresponds to the observational completeness limit and marks the mass ratio below which meaningful constraints cannot be obtained.

The fact that $p(q)$ is modeled as a double power law with a sharp break leads to some unphysical artifacts in Figure 7.9, including a sharp change in slope at $q = q_{\text{break}}$ and

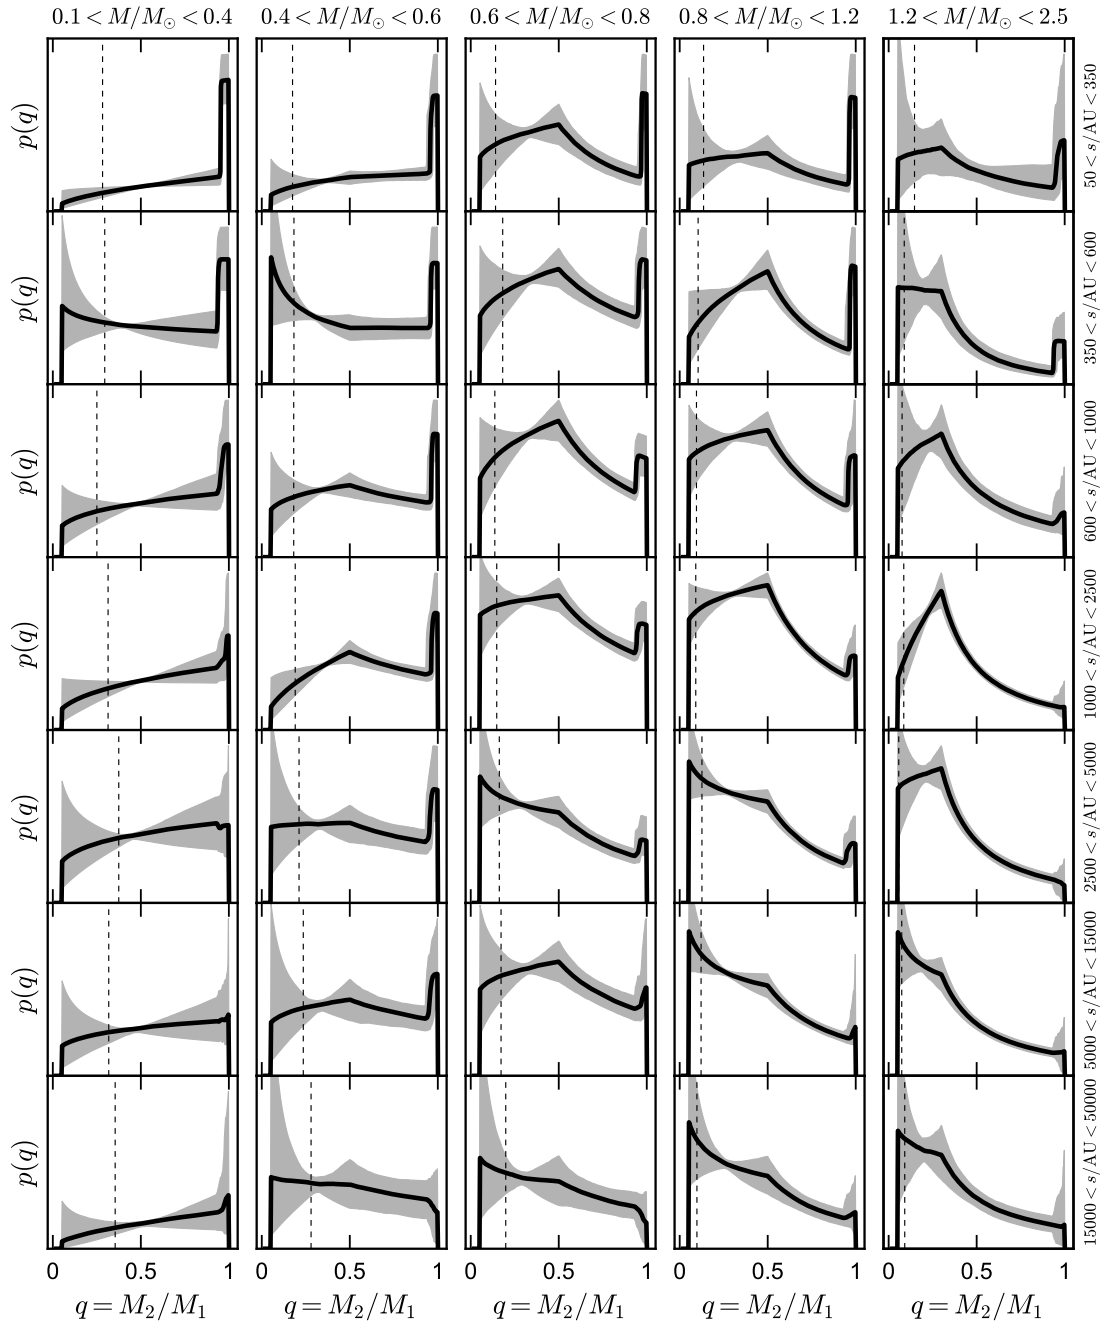


Figure 7.9: Mass ratio distribution constraints. Each panel corresponds to a single bin of primary mass (increasing left to right) and physical separation (increasing top to bottom). Gray shaded regions show 95.4% probability. Normalization is arbitrary; the scale of the y-axis is linear and begins at 0. Vertical dashed lines show the lowest mass ratio of observed binaries in each bin.

artificially low uncertainty just below q_{break} , which is caused by $p(q)$ “pivoting” about this point as $\gamma_{\text{small}q}$ varies and $p(q)$ is renormalized. We show in Appendix 7.8.3 that these features are not present when we fit a more flexible “smoothly-broken” power law model. However, doing so introduces parameter covariances that are not present for the fiducial form of $p(q)$. We use the simpler sharply broken power law as our fiducial model to facilitate easier comparison between different mass and separation bins and comparison with the literature.

Figure 7.10 compares the observed 1D distributions of ΔG to the predictions of the best-fit model; this is useful for assessing the quality of the fits. In each panel, we note the number of observed binaries and the marginalized 1 sigma constraints on F_{twin} . In generating the model predictions, we use the median of the 1D marginalized posterior distribution for each free parameter. We then generate a Monte Carlo realization of the binary population, drawing primary masses, mass ratios, separations, ages, metallicities, and distances from the appropriate joint distributions, calculating synthetic photometry, and weighting each binary by the selection function evaluated for its observables.

Overall, the model predictions are in good agreement with the observed distributions of ΔG . This indicates that our parameterization of $p(q)$ is suitable and sufficiently flexible. In panels where F_{twin} is inconsistent with 0, there is a clear excess of equal-brightness binaries. The distributions of ΔG near $\Delta G = 0$ are also adequately reproduced by the model, indicating that the simple “step function” model for the twin excess is consistent with the data. Although it is not shown here, we also find the best-fit models to predict distributions of other observables (angular separation, parallax, and apparent magnitude of the primary) in good agreement with the data.

Figure 7.11 shows the constraints on F_{twin} as a function of mass and projected physical separation. Light and dark error bars show middle 95.4% and 68.2% of the marginalized posterior distributions. Note that the y-axis scale is different in the top panel. In all bins of primary mass, the excess twin fraction falls with increasing separation and is negligible at the largest separations we consider ($s > 15,000$ AU). At fixed separation, the magnitude of the excess varies with primary mass. For close separations ($50 < s/\text{AU} < 350$), F_{twin} is largest in the lowest bin of primary mass ($0.1 < M_1/M_\odot < 0.4$). This may be in part because low-mass primaries in our sample are at closer distances on average, such that the median separation within the ($50 < s/\text{AU} < 350$) bin is smaller than for higher-mass primaries.

The maximum separation out to which there is a significant twin excess also varies with primary mass. For the highest-mass bin, F_{twin} is consistent with 0 at $s > 600$ AU. For solar-type primaries, F_{twin} is only consistent with 0 at $s > 5000$ AU, but it is $< 3\%$ for $s > 1000$ AU. The fall-off in F_{twin} with increasing separation is shallowest for $0.4 < M_1/M_\odot < 0.6$ and $0.6 < M_1/M_\odot < 0.8$, with a larger normalization for $0.4 < M_1/M_\odot < 0.6$. Here F_{twin} is almost independent of separation over $350 < s/\text{AU} < 15,000$. Finally, the fall-off steepens again in the lowest primary mass bin, where $F_{\text{twin}} \approx 0$ beyond 2500 AU. We discuss possible interpretations of these trends in Section 7.6.2.

In the bottom two panels of Figure 7.11, we compare our constraints on F_{twin} to the 1 sigma constraints obtained by Moe & Di Stefano (2017a, their Tables 8 and 11) for binaries in similar mass and separation ranges. The constraints for solar-type stars ($0.8 < M_1/M_\odot <$

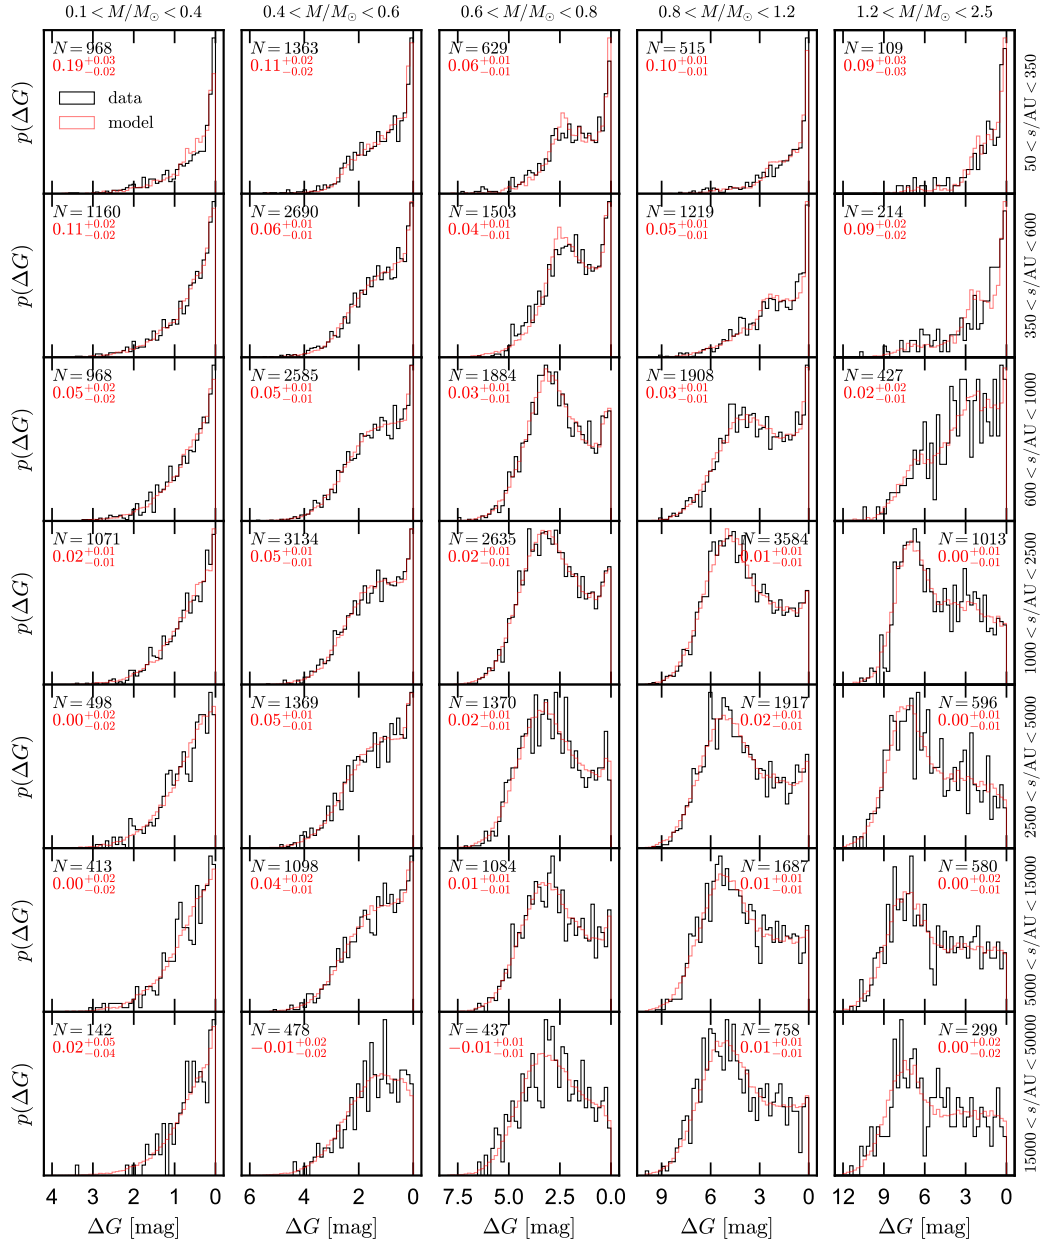


Figure 7.10: Magnitude difference between the two components of a binary (linear y-axis; normalization is arbitrary). Black histograms show real binaries, split into bins of primary mass (increasing left to right) and physical separation (increasing top to bottom). The twin excess can be seen in many panels as an excess of binaries with $\Delta G \approx 0$. Red histograms show Monte Carlo populations generated from the best-fit model. The number of observed binaries (black) and the marginalized 1 sigma constraints on F_{twin} (red) are listed in each panel. Overall, the model is quite successful in matching the observed distributions.




Figure 7.11: Marginalized constraints on F_{twin} , the fractional excess of binaries with nearly equal mass, vs. physical separation. Each panel shows a separate bin of primary mass. Dark and light error bars show 1 and 2 sigma constraints. In all mass bins, F_{twin} declines with increasing separation and is consistent with 0 at the largest separations. The twin excess reaches the widest separations for $0.4 < M_1/M_\odot < 0.6$, where F_{twin} is $\sim 5\%$ out to 15,000 AU. We compare to constraints from the nearby solar-type binary sample of [Raghavan et al. \(2010, panel 4\)](#) and the adaptive optics imaging survey from [De Rosa et al. \(2014, panel 5\)](#); these were derived by [Moe & Di Stefano \(2017a\)](#).

1.2) were obtained from the solar neighborhood sample of [Raghavan et al. \(2010\)](#). Those shown in the bottom panel were obtained from the AO-assisted survey of visual binaries with A-type primaries ($1.7 < M_1/M_\odot < 2.3$) described in [De Rosa et al. \(2014\)](#). Reassuringly, constraints for both mass bins are consistent with those obtained from our catalog. Because the binary sample we analyze is larger than the [Raghavan et al. \(2010\)](#) and [De Rosa et al. \(2014\)](#) samples, we can tighten the uncertainties on F_{twin} at large separations, showing, for instance, that F_{twin} for solar-type primaries is inconsistent with 0 out to $s \approx 5000$ AU.

Figure 7.12 shows constraints on γ_{largeq} as a function of separation, again separating binaries by primary mass. The dashed line in each panel shows the slope that would be expected for random pairings from the IMF. This is obtained by fitting a power law mass ratio distribution over the same range of primary masses and mass ratios to a simulated population of binaries in which the masses of both components are drawn from a [Kroupa \(2001b\)](#) IMF and paired randomly.

Consistent with previous work (e.g. [Lépine & Bongiorno 2007a](#); [Reggiani & Meyer 2011](#); [Duchêne & Kraus 2013](#); [Moe & Di Stefano 2017a](#)), we find that the mass ratio distribution is *not* consistent with random pairings from the IMF, but is weighted toward higher mass ratios than would be expected in such a scenario. Whether this is an imprint of the binary formation process or in part reflects the fact that binaries with higher mass ratio have higher binding energy and are thus more difficult to disrupt is an open question. Any trends in γ_{largeq} with separation are weak over the separation range we probe: at the 2 sigma level, our constraints are consistent with a separation-independent γ_{largeq} over $100 \lesssim s/\text{AU} < 50,000$ in all mass bins. However, they are not consistent with being independent of primary mass: $p(q)$ becomes increasingly bottom-heavy (lower γ_{largeq}) with increasing M_1 .

The fact that γ_{largeq} does not vary much with separation can serve as a strong constraint on formation models for wide binaries. It has frequently been argued that while binaries with separations of $100 \lesssim s/\text{AU} \lesssim 5000$ formed primarily by core fragmentation, those with $s \gtrsim 5000$ AU (the size of typical cloud cores) formed by another process. Candidate processes include cluster dissolution ([Kouwenhoven et al. 2010](#); [Moeckel & Bate 2010](#)), unfolding of hierarchical triples ([Reipurth & Mikkola 2012](#)), or pairing of adjacent cores ([Tokovinin 2017b](#)). One might naively expect a change in the mass ratio distribution at $s \sim 5000$ AU if the binary formation mechanism changes there, but none is observed.

The mass ratio distribution for wide solar-type binaries is not uniform, but is weighted towards low mass ratios. For example, companions with $q \approx 0.5$ are roughly twice as common as those with $q \approx 0.9$. The dominance of low-mass ratio companions can be seen clearly in the data at wide separations (Figure 7.10) and cannot be due to selection effects, which all work *against* low-mass ratio binaries. $p(q)$ is thus more bottom-heavy at wide separations than at close separations, where it is basically uniform ([Mazeh et al. 1992](#); [Tokovinin 2014a](#)). Analyzing the 25-pc [Raghavan et al. \(2010\)](#) sample of solar-type binaries, [Moe & Di Stefano \(2017a\)](#) found the mass ratio distribution to transition from $\gamma_{\text{largeq}} = -0.4 \pm 0.3$ (close to uniform) across $10 < s/\text{AU} < 200$ to $\gamma_{\text{largeq}} = -1.1 \pm 0.3$ across




Figure 7.12: Marginalized constraints on $\gamma_{\text{large}q}$, the logarithmic slope of the power law component of the mass ratio distribution at large q , vs. projected separation. Dark and light error bars show 1 and 2 sigma constraints. The range of q over which $\gamma_{\text{large}q}$ is fit in each mass bin is noted in the upper left of each panel. Dashed lines show the slope expected if binary component masses were drawn from the IMF and paired randomly. Except in the lowest mass bin, $p(q)$ is more bottom-heavy than uniform, but it is always more top-heavy than expected for random IMF pairings.

$200 < s/\text{AU} < 5000$.⁶ Combined with our constraints at wide separations, this implies that the transition between a uniform mass ratio distribution at close separations and a bottom-heavy distribution at wide separations occurs relatively abruptly at $s \sim 100$ AU. Several other binary population properties are observed to change at $s \sim 100$ AU (see [El-Badry & Rix 2019b](#) and references therein), perhaps due to a transition in the dominant binary formation mechanism at this separation. This sharp transition, and the fact that $\gamma_{\text{large}q}$ is nearly constant over $300 \lesssim s/\text{AU} \lesssim 50,000$, provides a useful constraint for star formation models. Because the effects of dynamical processing after formation on the mass ratio distribution are imperfectly understood, similar constraints obtained in star-forming environments will prove useful for disentangling the primordial mass ratio distribution from the effects of dynamical processing.

We show constraints on other parameters obtained from our fitting in [Figure 7.13](#). The left column shows q_{twin} . This parameter is only meaningful if F_{twin} is nonzero, so we do not show constraints for bins of mass and separation where F_{twin} is consistent with 0 at the 2 sigma level. There are no strong trends in q_{twin} with mass or separation; the typical best-fit value is $q_{\text{twin}} \approx 0.95$. In a few bins (e.g. $1.2 < M_1/M_\odot < 2.5$ and $350 < s/\text{AU} < 600$), the constraint runs up against the prior, implying that a broader excess may be preferred. However, the width of the observed $\Delta G \approx 0$ excess is still reproduced reasonably well in these cases ([Figure 7.10](#)).

Constraints on $\gamma_{\text{small}q}$ are shown in the middle column. Most constraints are consistent with $\gamma_{\text{small}q} = 0$ (a flat distribution at small q), and trends with separation are weak. For the lowest-mass bin ($M_1 < 0.4M_\odot$), we fit a single power law, so $\gamma_{\text{small}q} = \gamma_{\text{large}q}$. Unlike for $\gamma_{\text{large}q}$, there are no strong trends with primary mass. We note that our data cannot strongly constrain $\gamma_{\text{small}q}$ for low-mass primaries, so the constraints are influenced somewhat by the priors ([Section 7.4.1](#)).

Finally, the right column of [Figure 7.13](#) shows constraints on the local slope of the separation distribution γ_s , which is defined such that $p(s) \propto s^{\gamma_s}$. The dashed line shows $\gamma_s = -1.6$, which is approximately the value that has been found for wide binaries when marginalizing over a larger range of primary masses and separations ([Andrews et al. 2017](#); [ER18](#)). The constraints we find here are similar to this value on average but show some evidence for a steepening in $p(s)$ with increasing separation. Any trends with primary mass at fixed separation are weak. We emphasize that these constraints come from the gradient in binary counts as a function of s *measured within a narrow bin of s* , which is necessarily noisy: trends in γ_s with separation represent the second derivative of $p(s)$.

⁶Because we use $q_{\text{break}} = 0.5$ for solar-type primaries and [Moe & Di Stefano \(2017a\)](#) used $q_{\text{break}} = 0.3$, our measurements should not be directly compared. $p(q)$ flattens at $q < q_{\text{break}}$, so at fixed $\gamma_{\text{large}q}$, a lower q_{break} corresponds to a more bottom-heavy mass ratio distribution.




Figure 7.13: Constraints on q_{twin} (left; the mass ratio above which the twin excess begins), γ_{smallq} (middle; the logarithmic slope of the mass ratio distribution at small q), and γ_s (right; the logarithmic slope of the local separation distribution). Dark and light error bars show 1 and 2 sigma constraints. Primary mass increases from top to bottom. Constraints on q_{twin} are only shown for mass and separation bins where F_{twin} is inconsistent with 0 at the 2 sigma level. Dashed line in the left panels shows $q = 0.93$, the lower limit set on q_{twin} by the prior. Dashed lines in the right panels show $\gamma_s = -1.6$, the value found by ER18 when averaging over all masses and separations.

7.6 Discussion

7.6.1 Comparison to previous work

Although most work on twin binaries to date has focused on spectroscopic binaries, hints of a twin excess at wide separation have been reported in several previous works. [Trimble \(1987\)](#) found a significant excess of equal-brightness pairs among a sample of bright, nearby visual binaries. She suggested that this excess might point toward a formation mechanism that favors equal-mass systems but found significantly different distributions of magnitude difference when comparing different samples of wide binaries and thus did not rule out the possibility that the $q \sim 1$ peak was the result of selection effects. Similar conclusions were reached by [Giannuzzi \(1987\)](#). [Halbwachs \(1988\)](#) argued that the mass ratio distribution of wide, common proper motion binaries was likely consistent with random pairings from the IMF once selection effects were corrected for.

Larger and more homogeneous samples of bright wide binaries were identified using astrometry from the Hipparcos satellite for one or both components (e.g. [Söderhjelm 2000, 2007](#); [Eggenberger et al. 2004](#); [Lépine & Bongiorno 2007a](#); [Shaya & Olling 2011](#)). The mass ratio distribution of Hipparcos binaries with A and F star primaries was modeled in detail by [Söderhjelm \(2007\)](#), who found evidence for a $q \approx 1$ peak at $100 \lesssim s/\text{AU} \lesssim 1000$. The strength of the peak decreased with primary mass. He argued that the twin excess was not the result of any known selection effect, although he did not reject the possibility that an unknown bias in the Hipparcos input catalog could explain it. The twin feature identified by [Söderhjelm \(2007\)](#) is likely the same feature apparent in our catalog. We note however, that most of the binaries observed by Hipparcos fall in our highest-primary mass bin (where the twin excess is weaker than at lower masses), as Hipparcos only observed bright stars ($G \lesssim 13$).

A preference for equal-brightness pairs was also noticed among wide common proper motion disk and halo binaries identified by [Chanamé & Gould \(2004\)](#) and [Dhital et al. \(2010\)](#). The authors could not rule out the possibility that it was the result of selection effects, which were not well understood for their samples. The twin excess in our catalog is visually quite striking (e.g. [Figure 7.4](#)), so a natural question is why it was not as clear in earlier binary catalogs. In [Appendix 7.10](#), we show that a twin excess *is* apparent in the large, low-mass wide binary catalog produced by [Dhital et al. \(2015\)](#) using SDSS photometry, but it only becomes obvious once objects with blended photometry are removed. It is also clearer in the *Gaia* photometry than in the ground-based SDSS photometry.

More broadly, the twin excess we identify is visually striking because it is narrow, but it represents only a few percent of the total twin population. This means that it will only become obvious when (a) the photometry is sufficiently precise and uncontaminated that the difference in the components' apparent magnitude can be measured with precision that is good compared to the intrinsic width of the twin excess, and (b) the sample considered is large, containing (at least) hundreds of objects. A twin excess among solar-type visual binaries is also observed in the [Raghavan et al. \(2010\)](#) 25-pc and [Tokovinin \(2014a\)](#) 67-pc

samples, extending out to $s \approx 200$ AU at a statistically significant level [Moe & Di Stefano \(2017a\)](#). This excess is consistent with our constraints (e.g. [Figure 7.11](#)); at wider separations, these samples did not contain enough objects to detect or rule out a few-percent excess with high significance.

Width of the twin excess

Among spectroscopic binaries, there has been some disagreement in the literature over whether the twin excess is limited to a narrow peak in the mass ratio distribution at $q \gtrsim 0.95$ (e.g. [Tokovinin 2000](#)) or is a broader feature, corresponding simply to a positive slope in $p(q)$ at $q \gtrsim 0.8$ (e.g. [Halbwachs et al. 2003](#); see [Lucy 2006](#) for further discussion). In our sample, the twin feature is unambiguously narrow, only becoming significant above $q_{\text{twin}} \approx 0.95$ ([Figure 7.13](#)). In [Appendix 7.8.1](#), we show that $p(q)$ is always consistent with a flat or negative power law at $q < 0.94$; for the majority of the mass and separation bins, the twin excess only becomes strong at $q > 0.97$ (see [Figure 7.18](#)). To allow better comparison between close and wide binaries, we now re-examine the twin excess among spectroscopic binaries.

In the volume-limited 67-pc sample of FG dwarfs, [Tokovinin \(2014a\)](#) identified 98 double-lined spectroscopic binaries (SB2s) with periods $P = 1\text{--}100$ days, 70 of which have dynamical mass ratios $0.76 < q < 1$. Nearly all binaries with $P < 100$ days and $q > 0.76$ will appear as SB2s, so the 70 observed systems in this parameter space represent a relatively complete subsample. In [Figure 7.14](#), we plot the cumulative mass ratio distribution of the 70 short-period SB2s with $q > 0.76$, about half of which have $q > 0.95$. We model a uniform mass ratio distribution across $0.76 < q < 1$ with an excess twin fraction F_{twin} above $q > q_{\text{twin}}$. We use the maximum-likelihood method described in [Moe & Di Stefano \(2017a\)](#) to fit the two free parameters F_{twin} and q_{twin} and draw 1,000 bootstrap samples to estimate their uncertainties. We show in [Figure 7.14](#) the best-fit models obtained when fixing $q_{\text{twin}} = 0.93, 0.95,$ and 0.97 . The value $q_{\text{twin}} = 0.93$ is inconsistent with the data ($p = 0.014$), while $q_{\text{twin}} = 0.95$ and 0.97 both provide reasonable fits. We formally measure $q_{\text{twin}} = 0.964 \pm 0.013$ (1 sigma uncertainties).⁷ By using a larger and more complete sample of short-period SB2s, we thus confirm the conclusions of [Tokovinin \(2000\)](#) and [Moe & Di Stefano \(2017a\)](#) that close solar-type binaries with $a < 0.5$ AU exhibit a large excess twin fraction and that the twins are narrowly distributed above $q_{\text{twin}} \gtrsim 0.95$.

In [Figure 7.15](#), we show constraints on F_{twin} and q_{twin} for solar-type binaries across a wide range of periods and separations. At wide separations, the constraints from *Gaia* wide binaries are reproduced from [Figure 7.11](#) and [7.13](#). At closer separations, we show the constraints on F_{twin} obtained by [Moe & Di Stefano \(2017a\)](#) from the [Raghavan et al. \(2010\)](#) sample, as well as the constraint on q_{twin} at $P = 1\text{--}100$ days derived above for the [Tokovinin \(2014a\)](#) sample ([Figure 7.14](#)). At intermediate periods ($P = 10^{2-6}$ days), we show $q_{\text{twin}} \approx 0.95 \pm 0.02$ based on the [Moe & Di Stefano \(2017a\)](#) analysis of the [Raghavan et al.](#)

⁷We do not present constraints on F_{twin} for this sample, because F_{twin} depends on $p(q)$ at $0.3 < q < 1$, and many lower-mass ratio binaries will not be double-lined.




Figure 7.14: Cumulative distribution function of short-period double-lined spectroscopic binaries with $q > 0.76$ in the volume-limited 67-pc sample of FG dwarfs from Tokovinin (2014a). We compare predictions for models with three values of q_{twin} (see Section 7.6.1). For this sample, we find $q_{\text{twin}} = 0.964 \pm 0.013$, consistent with the values we find for wide binaries with periods $P = 10^{5-9}$ days.




Figure 7.15: F_{twin} (top) and q_{twin} (bottom) for solar-type binaries ($0.8 \lesssim M_1/M_\odot \lesssim 1.2$). We compare results from *Gaia* wide binaries (dark and light error bars show 1 and 2 sigma constraints) to 1 sigma constraints at closer separations from the samples of Raghavan et al. (2010) and Tokovinin (2014a, see Figure 7.14). F_{twin} decreases with separation and is consistent with 0 at $a \gtrsim 5,000$ AU. However, $q_{\text{twin}} \approx 0.95$ is basically constant at all separations.

(2010) sample. This constraint is not the result of formal fitting but provided a good fit to the data (see Figure 30 of Moe & Di Stefano 2017a). Figure 7.15 shows that while F_{twin} decreases with separation, $q_{\text{twin}} \approx 0.95$ is consistent with being constant over all separations in this mass range. Similar values of q_{twin} are also found for massive binaries at the short periods where there is a significant twin excess (Moe & Di Stefano 2013).

7.6.2 Origin of twin binaries

It is typically assumed that the components of binaries wider than a few hundred AU formed nearly independently of one another (e.g. White & Ghez 2001; Moe & Di Stefano 2017a; Tokovinin 2017b; Moe et al. 2018; El-Badry & Rix 2019b) during turbulent core fragmentation (for binaries with separation less than a few thousand AU; e.g. Offner et al. 2010) or by becoming bound at slightly later times (for those with the widest separations; e.g. Parker et al. 2009; Moeckel & Bate 2010; Kouwenhoven et al. 2010; Tokovinin 2017a).

The existence of a narrow twin excess at $q \gtrsim 0.95$ suggests that the components of a fraction of binaries with $s \gg 100$ AU formed at closer separations in a highly correlated way. We do expect that the dynamical process of becoming and staying bound may lead to a preference for roughly equal-mass binaries (say, $q > 0.5$), because these have higher binding energy. This could quite reasonably explain, at least in part, why the power-law component of the mass ratio distribution is shallower than expected for random pairings from the IMF (Figure 7.12). But we do not expect dynamical processes to produce a *sharp* twin feature like what is observed: at fixed primary mass, the binding energy at $q = 0.9$ is not much less than that at $q = 1$. And indeed, simulations of binary formation during cluster dissolution find larger typical mass ratios at wide separations than predicted for random pairings, but they do not predict a narrow excess of twins (Kroupa 1998; Moeckel & Bate 2010; Kouwenhoven et al. 2010).

The excess twin fraction uniformly decreases with separation and eventually goes to 0 at $s > 15,000$ AU in all mass bins. The shape of the twin excess (i.e. q_{twin} and the slope of $p(q)$ at $q > q_{\text{twin}}$) does not vary much between 50 and 15,000 AU in our catalog. Moreover, it is effectively the same for spectroscopic binaries (with separations as close as 0.01 AU) and wide binaries (Figure 7.15). Invoking Occam’s razor, it seems more likely that the wide binary twin phenomenon is an extension of the phenomenon that has previously been observed at $s \lesssim 100$ AU than that it is produced by a qualitatively different process.

Even for close binaries, there is not a clear consensus in the literature about the physical origin of the twin phenomenon. Some models for the formation of twins can only apply to very close binaries. In the first paper to highlight the twin phenomenon for spectroscopic binaries, Lucy & Ricco (1979) suggested that twins were formed by fragmentation of rapidly rotating pre-main sequence stars during the late stages of dynamical collapse, at scales of $a \ll 1$ AU. Alternatively, Krumholz & Thompson (2007) proposed that twins could be produced by mass transfer between stars of initially different masses during pre-main sequence evolution.⁸ It seems implausible that such mechanisms can explain the twin phenomenon

⁸These authors sought to explain the observed twin excess in massive stars. The specific mechanism they

among wide binaries, because there is no known mechanism to widen the orbits of twins from the separations at which they operate – a few, or at most a few tens of, solar radii – to the separations at which they are observed today. Such widening would require a very strong velocity kick, the magnitude of which would have to be fine-tuned in order to not unbind the binaries completely.

A more plausible formation mechanism for equal-mass twins at wider separations is through competitive accretion from a circumbinary disk. Many studies have found that the accretion rate from a circumbinary disk is usually higher for the secondary than the primary (e.g. [Bate & Bonnell 1997](#); [Bate 2000](#); [Farris et al. 2014](#); [Young & Clarke 2015](#); [Nelson & Marzari 2016](#); [Matsumoto et al. 2019](#)). Because the secondary’s orbit is larger than that of the primary, it sweeps out a larger radius in the disk and can accrete more rapidly than the primary, unless the material being accreted has low angular momentum.⁹ Preferential accretion onto the secondary will necessarily drive the mass ratio towards unity. If such accretion continues for long enough, binaries within circumbinary disks should thus end up with $q \approx 1$. An appeal of this formation mechanism is that it can operate at scales comparable to the size of circumbinary disks, $s \lesssim 100$ AU.

It is plausible but not obvious that preferential accretion onto the secondary will give rise to a sharp twin feature like the one found observationally. In order to end up at $q \approx 1$, a binary that initially had an intermediate mass ratio must accrete a large fraction of its mass from a circumbinary disk, such that there is enough time to drive the mass ratio to 1 even while the primary continues to accrete. If twins are formed by accretion from circumbinary disks, then the width of the twin feature (i.e., q_{twin}) can tell us about the fraction of the total mass accreted from the disk, as well as the mass ratio above which the accretion rate onto the two components becomes nearly equal. A sharp twin feature could be expected if a fraction of binaries accrete most of their mass from a circumbinary disk (becoming twins) and the rest either do not develop circumbinary disks or only accrete a subdominant fraction of their total mass from them.

In accretion-driven explanations of the twin excess, the increase in F_{twin} towards close separations has been interpreted as evidence that gravitational torques within circumbinary disks (e.g. [Artymowicz et al. 1991](#); [Shi et al. 2012](#)) shrink the orbits of twins ([Young & Clarke 2015](#)). Indeed, the observed twin excess is largest at separations $a < 0.2$ AU (Figure 7.15), too close for binaries to have formed at their current separations. This implies that some

proposed, which relies on deuterium shell burning causing protostars to expand and overflow their Roche lobes, cannot operate in solar-type or lower mass stars. Moreover, the twin excess for massive stars appears to be limited to close separations ([Moe & Di Stefano 2017a](#)), so mass transfer may adequately explain it. Here we simply suppose, for the sake of argument, that there is some mechanism through which stable mass transfer in lower mass stars could drive the mass ratio to unity.

⁹It is worth noting that there has not been a full consensus in the literature whether this mechanism works: some simulations of accretion from a circumbinary disk have actually predicted the opposite trend, with accretion favoring the primary ([Ochi et al. 2005](#); [Hanawa et al. 2010](#); [de Val-Borro et al. 2011](#)). These simulations assumed a higher gas temperature than those which have found accretion to favor the secondary; [Young & Clarke \(2015\)](#) showed that accretion only favors the secondary when accreted gas is cold, as is appropriate for stellar binaries.

combination of gravitational torques, viscous dissipation, and dynamical interactions (e.g. [Bate 2012](#); [Moe et al. 2018](#)) must have shrunk the orbits of twins at very close separations.

However, several recent simulations of circumbinary disks have found that, contrary to previous results in the literature, accretion can also *widen* binaries within circumbinary disks, when the advective torque dominates over the gravitational torque ([Miranda et al. 2017](#); [Muñoz et al. 2019](#); [Moody et al. 2019](#)). Whether gravitational or advective torques dominate depends on details such as the sink prescription used for accretion ([Tang et al. 2017](#)). Further work is needed to clarify the effects of circumbinary disks on orbital evolution. However, the fact that the twin excess extends to very wide separations suggests that orbit shrinkage is not a ubiquitous outcome of accretion from circumbinary disks.

High-resolution studies of the dynamics of individual circumbinary disks are generally too idealized, run for too short a time, and are focused on too narrow a range of initial conditions to make ab-initio predictions of the full mass ratio distribution. However, they do typically find that gas is preferentially accreted onto the secondary for realistic accretion geometries once a steady state is reached. On the other hand, global simulations of the fragmentation of molecular clouds (e.g. [Bate 2009, 2014, 2019](#)) are reaching the point where they can make realistic predictions of the mass ratio distribution with minimal fine-tuning. Such studies have lower resolution than idealized simulations of individual binaries, so it is not guaranteed that the dynamics within disks are well resolved, but they are able to predict the accretion rate and angular momentum distribution of accreted material, and the mass and size distribution of disks (e.g. [Bate 2018](#)). Such global calculations predict an excess of equal-mass binaries out to separations of order 100 AU (see e.g. [Bate 2014](#), Figure 18). Because they typically only form a few dozen binaries in a cloud, such calculations do not yet have the statistical power to predict or rule out a few-percent effect at wider separations.

Accretion can plausibly explain an excess of twins out to significantly larger separations than mass transfer or late-stage fragmentation, but it alone cannot explain a signal reaching out to thousands of AU. Observed circumstellar and circumbinary gas disks have typical radii of order 100 AU ([Ansdell et al. 2018](#); [Eisner et al. 2018](#)). The *largest* observed circumbinary disks have radii of order 500 AU and host binaries with separations of 50-200 AU (e.g. [Hioki et al. 2007](#); [Brinch et al. 2016](#); [Tobin et al. 2016](#); [Takakuwa et al. 2017](#); [Comerón et al. 2018](#); [Czekala et al. 2019](#)); these preferentially host relatively massive binaries. There are no observed disks with radii exceeding 1000 AU, and simulations also predict the largest circumbinary disks to have radii of several hundred AU ([Bate 2018](#)). It thus seems exceedingly unlikely that twin binaries with $s \gtrsim 200$ AU formed at their present-day separation by accretion from a circumbinary disk. This implies that either (a) twin binaries formed at closer separations and their orbits were subsequently widened, or (b) some other process is responsible for producing twins at very wide separations. Lacking a good candidate mechanism for (b), we here consider the plausibility of orbit widening.

Dynamical orbit widening in young clusters

In the Galactic field, dynamical interactions have a negligible effect on most binaries with $s \lesssim 10,000$ AU (e.g. [Weinberg et al. 1987](#)). However, dynamical interactions are expected to be more efficient in binaries' birth environments, where the typical stellar density is higher. The dynamical evolution of binaries within their birth clusters has been the subject of considerable study. As a general rule, interactions within birth clusters are expected to widen the orbits of binaries with orbital velocities lower than the cluster velocity dispersion, and to tighten the orbits of those with orbital velocities greater than it ([Heggie 1975](#); [Hills 1975](#)). There are several complicating factors in real clusters. For example, the mass distribution within clusters at early times is not smooth but clumpy ([Dorval et al. 2017](#)), both stars and gas can be dynamically important, and cluster density and velocity dispersion change as clusters age, in part because energy is redistributed among binaries (see e.g. [Kroupa & Burkert 2001](#); [Parker et al. 2009](#); [Goodwin 2010](#)). Observed binary populations provide constraints on models for binaries' dynamical evolution, but a unified model to explain the diversity of binary populations found in different environments does not exist. We summarize some relevant observational constraints on disruption below.

1. Over a wide range of separations ($10 \lesssim s/\text{AU} < 3000$), the binary fraction in low-density star forming regions (e.g. Taurus) is higher than in the field by roughly a factor of 2 (e.g. [Leinert et al. 1993](#); [Duchêne 1999](#)). This is true especially at wider separations, where the separation distribution of young binaries is roughly log-uniform, but that of field stars declines more steeply (e.g. [Connelley et al. 2008](#); [Kraus et al. 2011](#)).
2. In dense young clusters (e.g. the ONC), the binary fraction at separations of $s \gtrsim 100$ AU is similar to the field (lower than in low-density star forming regions) and declines steeply at wide separations ([Reipurth et al. 2007](#)). At closer separations ($10 \lesssim s/\text{AU} \lesssim 60$) the binary fraction in dense clusters is comparable to that in low-density star forming regions, and higher than that found in the field ([Duchêne et al. 2018](#)).
3. At very close separations ($s \lesssim 5$ AU), the binary fraction in star forming regions (over a range of densities) is consistent with that in the field (e.g. [Kounkel et al. 2019](#)).

Some models (e.g. [Kroupa 1995](#); [Marks et al. 2011](#); [Marks & Kroupa 2012](#)) postulate that the initial binary fraction and separation distribution are insensitive to local properties, such that observed variation in binary populations must be due to disruption. These models interpret the higher wide binary fraction in low-density clusters as the primordial population, which is transformed into the field population by dynamical widening and disruption. Because disruption of wide systems is more rapid in dense clusters, these models also predict the binary fraction at $s \gtrsim 100$ AU to decrease with cluster density, in agreement with observations. However, if such models are correct, it is not clear what happens to the excess of relatively tight binaries ($10 \lesssim s/\text{AU} \lesssim 100$), which are over-represented relative to the field in both high- and low-density star forming regions. Such systems have high enough

binding energies that they can only be disrupted in very dense clusters. Thus the lower binary fraction in the field at $10 \lesssim s/\text{AU} \lesssim 100$ would imply that a large fraction of field stars formed in dense environments.

In any case, the fact that the binary fraction in the field is lower than that in star forming regions down to fairly close separations implies that a significant fraction of young binaries undergo quite energetic interactions, the cumulative effect of which is sufficient to disrupt binaries with initial separations as close as 10-100 AU. In many cases, binaries will be disrupted by such interactions, but in some cases, they will only be widened (e.g. [Kroupa & Burkert 2001](#)). It is in this latter context that the twin fraction at wide separations is informative about the fraction of stars at a given present-day separation that formed at significantly closer separations.

The fact that the dependence of F_{twin} on separation varies with binary mass in a non-monotonic way (Figure 7.11) suggests that, if the twin excess at $s \gg 100$ AU is due to dynamical widening of orbits, then the primordial twin statistics (e.g. the twin fraction and the range of separations over which twins are form) must also vary with mass. This may not be unreasonable, since the physical properties of disks do vary with mass ([Bate 2018](#); [Eisner et al. 2018](#)), but it means that disentangling the effects of dynamical widening and the primordial separation distribution of twins is nontrivial.

We use a simple toy model to explore how the separation-dependence of the twin excess may depend on the initial twin population and the strength of dynamical interactions. We suppose that twin formation is efficient out to a separation of $a_{\text{twin, max}} \approx 100$ AU and does not operate at wider initial separations. We consider primordial populations of equal-mass binaries with components of mass M and a uniform distribution of $\log a$ between 10 and 10^4 AU. At $a < a_{\text{twin, max}}$, we assign a random subset of 20% of the binaries the label of “twin”. We then assume that dynamical interactions add an energy E_{int} to each orbit (for both twins and non-twins), such that the new orbital energy is $E_{\text{new}} = E_{\text{initial}} + E_{\text{int}}$, where $E_{\text{initial}} = -GM^2/a_{\text{initial}}$. Values of E_{int} for each binary are drawn from a uniform distribution between 0 and $E_{\text{max, interactions}}$. If E_{new} is positive, then the binary is considered unbound and is discarded. The new separation after interactions is $a_{\text{new}} = -GM^2/E_{\text{new}}$. Finally, we measure what fraction of the surviving binaries bear the twin label as a function of final separation. This fraction is proportional to the fraction of binaries at a given present-day separation that formed with $a < a_{\text{twin, max}}$. We do not model dynamical hardening of close binaries, and thus implicitly assume that all binaries we consider are soft; i.e., that they reside in clusters with velocity dispersion larger than the highest orbital velocities of binaries being widened. For an initial separation of 100 AU and binaries with typical component mass of $0.5 M_{\odot}$, this corresponds to a dispersion of $\sigma \approx 1.5 \text{ km s}^{-1}$; for an initial separation of 10 AU, to $\sigma \approx 5 \text{ km s}^{-1}$.

We plot the results of this experiment in Figure 7.16, varying $a_{\text{twin, max}}$ (top), $E_{\text{max, interactions}}$ (middle), and M (bottom). The final separation-dependence of the twin fraction depends significantly on all of these parameters. The twin excess extends to larger separations, and falls off less steeply with separation, when (a) the initial separation out to which twins form is larger, (b) dynamical interactions are more energetic, or (c) the binding energy of twins is




Figure 7.16: Predictions of the toy model for orbit widening described in Section 7.6.2. We assume that the initial twin fraction is 20% for $a < a_{\text{twin,max}}$ (varied in the top panel; 100 AU by default) and 0 at larger separations. We consider an initial separation distribution that is flat in log space. We then widen the orbits of all binaries by adding energy, representing the cumulative effect of gravitational interactions within a birth cluster. The total added energy is drawn from a uniform distribution between 0 and $E_{\text{max,interaction}}$. Binaries that are disrupted are discarded. We plot the final twin fraction as a function of separation, varying $E_{\text{max,interaction}}$ (middle panel; default is $G(0.5M_{\odot})^2 / (100 \text{ AU})$) and mass (bottom panel; default is $0.5 M_{\odot}$).

lower. In order to obtain a nonzero twin fraction at very wide separations, it is necessary that a fraction of binaries undergo dynamical interactions energetic enough to unbind binaries at separation at which twins form. Once this is satisfied, twins can contribute significantly to the binary population at wider separations, because the same interactions that widen close binaries will unbind a large fraction of initially wider non-twin binaries. Wide twins are produced most efficiently in clusters where the velocity dispersion is comparable to the orbital velocity at a separation of $a_{\text{twin,max}}$: significantly denser clusters produce few wide binaries, since most binaries that are not close are disrupted completely.

The toy model generically predicts that F_{twin} decreases with separation and that low-mass twins can be more efficiently widened than high-mass twins due to their lower binding energy. It thus predicts that the excess twin fraction will fall off less steeply for lower-mass binaries. This trend is found in the observed binaries over $0.4 < M_1/M_\odot < 2.5$ (Figure 7.11), but it is reversed for the lowest-mass subsample: F_{twin} falls off significantly more steeply for $0.1 < M_1/M_\odot < 0.4$ than for $0.4 < M_1/M_\odot < 0.6$. This is not easily explained in the context of the toy model.

It is thought that a fraction of wide binaries (likely at $a \gg 1000$ AU) form at later times (perhaps during cluster dissolution) than wide binaries at closer separations. Dilution of the twin excess due to these binaries is not accounted for in the toy model. If the fraction of binaries at fixed separation that formed during cluster dissolution were higher at lower masses, this could explain the observed steeper decline in F_{twin} with separation at low masses.

An additional complication is that the observed population of field binaries is an average over a wide range of formation environments, from low-density regions to dense clusters. The trends in Figure 7.16 rely on a number of crude approximations and should not be used to directly interpret the observed trends in F_{twin} (and we have not attempted to tune the model to match observed trends). Here we simply emphasize that the separation distribution of twins is quite sensitive to dynamical processing. We conclude that dynamical orbit widening provides a plausible explanation for the existence of the wide twin excess, but more realistic theoretical modeling is needed to determine whether the observed trends in F_{twin} with primary mass and separation can be reproduced when averaging over a realistic population of star-forming environments.

Widening of unstable triple systems?

Another possible mechanism for dynamically widening twins is through unfolding of unstable triple systems, in which initially close companions can be scattered to much wider orbits (e.g. Reipurth & Mikkola 2012). Triples are not uncommon: more than a third of all wide binaries contain subsystems (e.g. Tokovinin & Smekhov 2002a; Tokovinin 2014a), and a large fraction of binaries are thought to have formed as higher-order multiples that subsequently decayed (e.g. Sterzik & Durisen 1998).

Widening of the outer orbit in triples generally comes at the expense of shrinking of the inner orbit. Since the outer orbit is the more fragile one, a triple-origin of the wide twin excess would imply that one of the component stars in wide twins had, and in most cases still




Figure 7.17: Magnitude difference between the two resolved components of all wide binaries in which one or both components is suspected to have a bright, unresolved companion ($N = 7,295$; yellow points in Figure 7.1). There is no obvious excess at $\Delta G \approx 0$ (compare to the right panel of Figure 7.3).

has, a close companion, which in most cases would be spatially unresolved. However, the twin excess at $\Delta G \lesssim 0.25$ mag cannot be due to systems in which either resolved component has an unresolved companion any less than ≈ 2.5 magnitudes fainter than it; otherwise the light from the unresolved companion would make it more than 0.25 mag brighter than the other resolved component. Figure 7.17 shows that there is indeed no obvious excess at $\Delta G \approx 0$ for wide binaries in which either component has a bright unresolved companion. There is thus a photometric upper limit on the mass ratio of any unresolved companions to twins with $\Delta G \approx 0$; it ranges from $q < 0.45$ to $q < 0.7$, depending on primary mass (see Figure 7.8).

Unresolved companions with lower mass ratios cannot be ruled out based on photometry, but at close separations, they can be detected from radial velocity (RV) variability. In Appendix 7.13, we show that the *Gaia* visit-to-visit radial velocity scatter implies a *lower* close binary fraction for the components of wide twins than for components of non-twins with similar separations and masses. This speaks against a triple-origin of wide twins. It is, however, consistent with our expectations if twins are formed through accretion in a circumbinary disk, since the presence of a third close companion in the disk would complicate the mechanism through which accretion drives the mass ratio of two stars in a disk toward unity.

7.7 Summary and conclusions

We have analyzed a pure and homogeneous sample of $\sim 42,000$ main-sequence wide binaries selected from *Gaia* DR2 to constrain the mass ratio distribution of binaries with projected separations $50 \lesssim s/\text{AU} < 50,000$ and primary masses $0.1 < M_1/M_\odot < 2.5$. High-precision photometry allows us to measure mass ratios with unprecedented accuracy, and a well-understood selection function makes it possible to account for biases arising from the magnitude and contrast sensitivity limits of *Gaia* DR2 given our quality cuts.

A striking feature of the catalog is a high-significance excess of “twin” binaries with nearly equal brightness (Figure 7.3). The excess is present over a narrow range of magnitude differences, $0 < \Delta G \lesssim 0.25$, corresponding to mass ratios $0.95 \lesssim q \lesssim 1$, and extends over a wide range of masses and separations. The twin excess is reminiscent of the excess of equal-mass binaries historically reported at very close separations ($a < 0.2$ AU) and recently found to extend to $a \sim 100$ AU, but it extends to separations of several thousand AU, where binary formation models do not predict strongly correlated component masses.

We have done a variety of tests to confirm that the twin feature is caused by an astrophysical excess of equal-mass binaries, not selection effects or data artifacts. We first repeated the search that produced the binary catalog, but required that the two components of a pair have *inconsistent* rather than consistent astrometry. This yields a catalog of physically unassociated “chance alignments” that are subject to essentially the same selection function as the real binaries and have similar distributions of most observable properties. A narrow excess of equal-brightness pairs is not found among chance alignments (Figure 7.4).

By considering binaries at different distances, we verified that the strength of the twin feature varies primarily with physical, not angular, separation (Figure 7.5). This speaks to the physicality of the twin feature, since most observational biases are expected to depend on angular separation, which is the observable quantity. Finally, we verified that the twin feature is not an artifact of the *Gaia* photometry; it is visible in photometry from other surveys. We also identified the same twin excess in archival data from another binary catalog after removing objects with contaminated photometry (Appendix 7.10).

In order to measure the intrinsic mass ratio distribution, we forward-model the joint distribution of magnitude difference, angular separation, and parallax given a Galactic model, a selection function, and a parameterized mass ratio distribution. We measure the selection function for our catalog empirically, using chance alignments to quantify the contrast sensitivity as a function of angular separation. Our primary results are as follows:

1. *Twin fraction*: we quantify the twin excess as F_{twin} , the fractional excess of binaries with $q > q_{\text{twin}}$ relative to the full population with $q > 0.3$ (Figure 7.7; $q_{\text{twin}} \approx 0.95$ quantifies the width of the excess). Typical values of F_{twin} are 10% at 100 AU and 3% at 1000 AU, with some dependence on mass (Figure 7.11). These values are lower than the value found for spectroscopic binaries (Figure 7.15), which is $F_{\text{twin}} \approx 25\%$, but they are clearly inconsistent with 0.
2. *Width of the twin excess*: at all masses and separations where it is statistically significant, the twin excess we find is quite narrow, only becoming significant at $q \gtrsim 0.95$ (Figure 7.13). We experimented with different functional forms for the enhancement near $q \approx 1$, but we find a step function increase at $q \gtrsim 0.95$ to perform as well as more complicated models (Figure 7.18). We also re-measured q_{twin} at close separations using a volume-limited sample of spectroscopic binaries (Figure 7.14). The width of the twin excess, q_{twin} , is basically constant over 6 orders of magnitude in separation, from the closest spectroscopic binaries to wide binaries at $s \sim 10,000$ AU.
3. *Mass and separation dependence*: the fractional excess of twins declines with separation and is consistent with 0 at $s > 15,000$ AU for all mass bins (Figure 7.11). The twin excess at $s > 1,000$ AU is strongest for primary masses $0.4 < M_1/M_\odot < 0.6$. At closer separations ($50 < s/\text{AU} < 1,000$), it is strongest for low-mass primaries ($M_1 < 0.4M_\odot$) and roughly consistent for other mass bins. The slope of the fall-off in F_{twin} with separation varies non-monotonically with mass; it is shallowest at $0.4 < M_1/M_\odot < 0.6$.
4. *Mass ratio distribution at lower q* : we provide broken power law fits to the full mass ratio distribution for all mass and separation bins (Figure 7.9; Appendix 7.14). These fits reproduce the observed data well (Figure 7.10). For solar-type stars, the power law slope is $\gamma_{\text{large}q} \approx -1.3$ for $q > 0.5$ (i.e., weighted toward lower mass ratios than the uniform distribution found at closer separations) and $\gamma_{\text{small}q} \approx 0$ for $q < 0.5$. $p(q)$ becomes more bottom-heavy with increasing primary mass but is always flatter than expected for random pairings from the IMF. Besides variation in the excess twin frac-

tion, $p(q)$ does not vary much with separation over $100 \lesssim s/\text{AU} < 50,000$ (Figures 7.12 and 7.13).

5. *Origin of the twin excess:* no theoretical models that have been proposed to explain twin binaries at close separations predict them to form at $s \gtrsim 100$ AU. Dynamical processes may lead to a formation bias against low- q binaries, but they are not expected to produce a sharp peak at $q \approx 1$. Given the monotonic fall-off in F_{twin} with separation, the similar width of the twin feature between close and wide binaries, and the lack of a plausible mechanism for forming twins at very wide separations, we conjecture that the excess twins must have formed at closer separations ($s \lesssim 100$ AU; likely through accretion from a circumbinary disk) and subsequently been widened by dynamical interactions.

In this scenario, the separation-dependence of the twin fraction is an imprint of dynamical orbit widening in binaries' birth environments (see Section 7.6). The plausibility of this explanation is not straightforward to assess because (a) present-day field binaries formed in a wide range of environments and (b) existing models for widening and disruption of binaries in star forming regions do not fully explain the diversity of observed binary populations in young clusters and in the field. A simple toy model suggests that a separation-dependence in F_{twin} comparable to that which is observed is predicted if typical dynamical interactions are strong enough to disrupt binaries at the separation inside which twins are expected to form ($a \lesssim 100$ AU). However, the mass-dependence of F_{twin} at wide separations is not fully explained in such models. Further theoretical work is required to (a) predict the primordial separation distribution of twins at different masses and (b) constrain the efficiency of dynamical orbit widening for a realistic ensemble of star forming environments. In future work, we will search for a wide twin excess in observed star forming regions to shed more light on the primordial separation distribution of twins and the density-dependence of orbit widening.

Acknowledgements

We are grateful to the referee, Andrei Tokovinin, for thoughtful comments. We thank Matthew Bate, Anthony Brown, Eugene Chiang, Ian Czekala, Paul Duffel, Morgan Fouesneau, Harshil Kamdar, Tomoaki Matsumoto, Chris Mckee, Eliot Quataert, and Dan Weisz for helpful discussions. KE was supported in part by an NSF graduate research fellowship and by SFB 881. HJT acknowledges the National Natural Science Foundation of China (NSFC) under grants 11873034. MM acknowledges financial support from NASA under Grant No. ATP-170070. This project was developed in part at the 2019 Santa Barbara Gaia Sprint, hosted by the Kavli Institute for Theoretical Physics at the University of California, Santa Barbara. This research was supported in part at KITP by the Heising-Simons Foundation and the National Science Foundation under Grant No. NSF PHY-1748958. This work has made use of data from the European Space Agency (ESA) mission *Gaia*

(<https://www.cosmos.esa.int/gaia>), processed by the *Gaia* Data Processing and Analysis Consortium (DPAC, <https://www.cosmos.esa.int/web/gaia/dpac/consortium>). Funding for the DPAC has been provided by national institutions, in particular the institutions participating in the *Gaia* Multilateral Agreement.

7.8 Functional form of $p(q)$

7.8.1 How sharp is the twin feature?

In our default model, the twin feature is characterized by two number: F_{twin} , the fractional excess of binaries with near-equal masses, and q_{twin} , the mass ratio above which the excess manifests itself. The implicit assumption in this model is that (a) the increase in the mass ratio distribution near $q = 1$ occurs abruptly, not gradually, and (b) the amount of twin excess is constant between $q = q_{\text{twin}}$ and $q = 1$. In this Section, we assess the validity of this assumption by using a more flexible model for the twin excess.

We model the twin excess between $q = 0.85$ and $q = 1$ as a histogram with 5 equally-spaced bins of width $\Delta q = 0.03$. As in the fiducial model, we assume the smooth underlying distribution in this region is $p(q) \sim q^{\gamma_{\text{large}q}}$. We then introduce a set of uncorrelated weights, w_i with $i = 0 \dots 5$, and multiply $p(q)$ in the i -th bin by w_i . With this parameterization, $w_i = 1$ entails no deviation from a smooth power law in the i -th bin, $w_i = 2$ corresponds to a factor-of-two enhancement relative to the underlying power law, and $w_i = 0$ means that there are no binaries at all with mass ratios in the i -th bin. With a sufficiently large number of bins, this parameterization is flexible enough to represent any arbitrary shape of $p(q)$. 5 bins of width 0.03 is a pragmatic choice given the size of our binary catalog, since the shot noise uncertainty increases as the bin size is reduced.

Figure 7.18 shows results of fitting a histogram model. Consistent with Figure 7.9, a statistically significant excess of twins is found out to several thousand AU for $0.1 < M_1/M_\odot < 1.2$ and $s \lesssim 2500$ AU, with the excess reaching $s = 15,000$ AU for $0.4 < M_1/M_\odot < 0.6$. In most mass and separation bins, the data is consistent with a smooth power law distribution (with no twin excess) all the way up to $q = 0.97$. In a few mass bins, there is also a significant excess relative to the underlying power law at $0.94 < q < 0.97$, but never at $q < 0.94$. That is, in all bins where it is significant, the twin feature is “thin”: there is not a broad peak at $q \gtrsim 0.8$, but a narrow excess only at $q \gtrsim 0.95$. It is because of the narrowness of the twin peak that the excess of equal-mass pairs is so clearly apparent in the observed data (e.g. Figure 7.3), even though twins only make up a small fraction ($< 5\%$; see Figure 7.11) of the total binary population at wide separations.

7.8.2 Choice of q_{break}

As discussed in Section 7.4, we model the mass ratio distribution in all but the lowest-mass bin as a broken power law whose slope changes at $q = q_{\text{break}}$ (Figure 7.7). In order to make comparison between different separation bins more straightforward, we do not leave

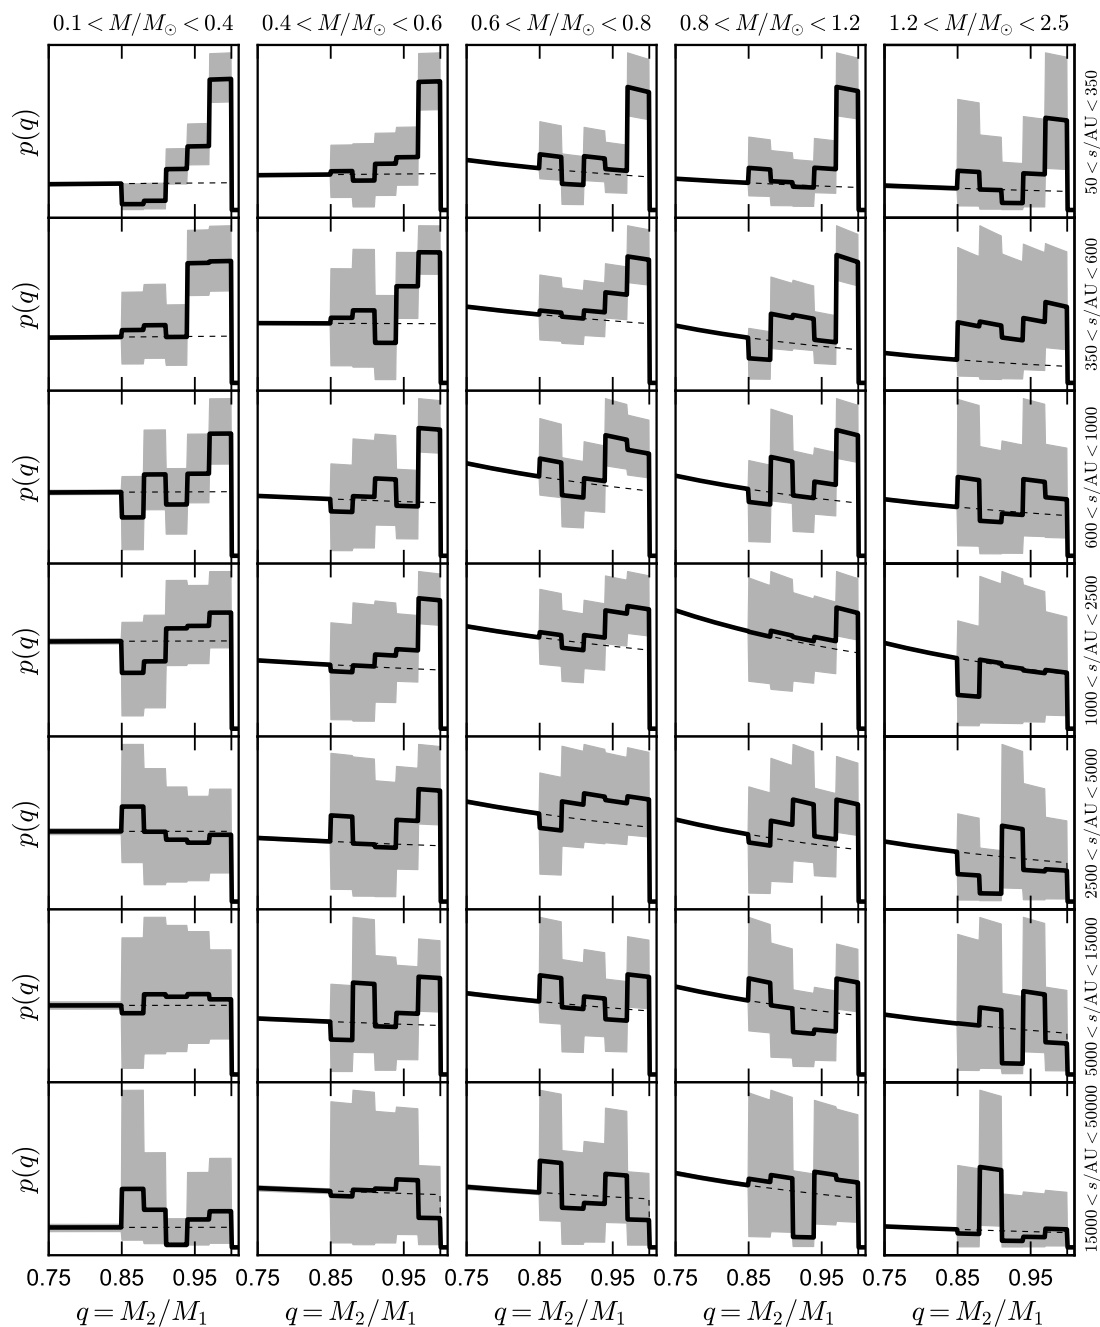


Figure 7.18: Constraints on $p(q)$ near $q = 1$. Solid lines and gray shaded regions represent median and 2σ (middle 95.4%) constraints; y-axis scale is linear. We model the mass ratio distribution over $0.85 < q < 1$ as a histogram with 5 bins of width 0.03, fitting for the height of each bin as a free parameter. The twin excess generally becomes statistically significant only at $0.97 < q < 1$, in some cases with a modest enhancement at $0.94 < q < 0.97$. At lower q , the mass ratio distribution is generally consistent with a smooth power law (dashed line).




Figure 7.19: Mass ratio distribution constraints (right) and predicted ΔG distributions (left) for fits with $q_{\text{break}} = 0.5$ (red) and $q_{\text{break}} = 0.3$ (cyan). We show a single separation bin ($1000 < s < 2500$). Top panels show solar-type stars ($0.8 < M_1/M_\odot < 1.2$); bottom panels show A and F stars ($1.2 < M_1/M_\odot < 2.5$). For each choice of q_{break} , we show the best-fit constraints when q_{break} is fixed at that value and other parameters are left free. For solar-type stars, a significantly better fit can be obtained with $q_{\text{break}} = 0.5$ than with $q_{\text{break}} = 0.3$. The opposite is true for A and F stars, where the data strongly favor $q_{\text{break}} = 0.3$.

q_{break} free, but instead fix it to a single value in a given mass bin: $q_{\text{break}} = 0.3$ for $1.2 < M_1/M_\odot < 2.5$, and $q_{\text{break}} = 0.5$ in all other mass bins. These values were chosen by trial and error, but it *is* necessary to use different values of q_{break} in different mass bins, as we now show.

Figure 7.19 shows constraints on the mass ratio distribution (right) and corresponding predicted ΔG distributions (left) for two choices of q_{break} and two bins of primary mass. A single separation bin is shown for illustrative purposes: $1000 < s < 2500$, where the twin excess is weak in both mass bins. In the left panels, we compare the observed distributions of ΔG to Monte Carlo binary populations produced for 5 draws from the posterior.

The top panels show that for solar-mass primaries ($0.8 < M_1/M_\odot < 1.2$), the best-fit model obtained while assuming $q_{\text{break}} = 0.5$ provides a significantly better fit to the observed ΔG distribution than the best-fit model with $q_{\text{break}} = 0.3$: in the latter case, the observed distribution is poorly reproduced both at $0 < \Delta G \lesssim 2$ and at $\Delta G \approx 5$. On the other hand, the bottom panels show that $q_{\text{break}} = 0.3$ provides a much better fit for A and F star primaries.

7.8.3 Smoothly-broken power law

As discussed in Section 7.5, our choice to model the mass ratio distribution with a broken power law leads to an unphysical sharp break in the best-fit mass ratio distribution at intermediate q . We test the sensitivity of our constraints on $p(q)$ to the assumed functional form below. As an alternative to a sharply broken power law, we fit a “smoothly broken power law” with the functional form

$$p(q) \propto \left(\frac{q}{q_{\text{break}}}\right)^{\gamma_{\text{small}q}} \left[1 + \left(\frac{q}{q_{\text{break}}}\right)^{1/\Delta}\right]^{(\gamma_{\text{large}q} - \gamma_{\text{small}q})\Delta}. \quad (7.4)$$

This function approaches $p \propto q^{\gamma_{\text{small}q}}$ at $q/q_{\text{break}} \ll 1$, and $p \propto q^{\gamma_{\text{large}q}}$ at $q/q_{\text{break}} \gg 1$. The parameter Δ controls the sharpness of the transition between the two regimes; in the limit of $\Delta \rightarrow 0$, it reduces to a simple double power law with a sharp transition between the two slopes.

In Figure 7.20, we compare the constraints on $p(q)$ obtained for a single mass and separation bin when using the fiducial model (gray) and a smoothly broken power law (hatched red). Overall, the constraints are quite similar, but as expected, fitting a broken power law model smooths the best-fit profile for $p(q)$ and removes the region near $q \sim 0.3$ where the nominal uncertainty is suppressed. We find qualitatively similar results for other bins of mass and separation. However, we note that increasing Δ shifts the peak of the distribution towards lower q (upper panel of Figure 7.20). Since Δ itself is often only weakly constrained, this can lead to parameter covariances between Δ , $\gamma_{\text{small}q}$, and q_{break} . We therefore use the simpler broken power law as our fiducial model.




Figure 7.20: Top panel shows example smoothly broken power laws for different values of the smoothing parameter, Δ . For all cases, $\gamma_{\text{small}q} = 0.3$, $\gamma_{\text{large}q} = -1.2$, $F_{\text{twin}} = 0.04$, and $q_{\text{twin}} = 0.95$. Increasing Δ softens the transition between $p \sim q^{\gamma_{\text{small}q}}$ and $p \sim q^{\gamma_{\text{large}q}}$. Bottom: middle 95.4% constraints on $p(q)$ for $0.8 < M_1/M_\odot < 1.2$ and $600 < s/\text{AU} < 1000$. We compare constraints obtained by fitting a smoothly broken power law (red hatched) to those obtained by fitting our default model, a simple double power law (i.e., forcing $\Delta = 0$). The constraints we obtain with the two models are similar, with the most significant difference being that the smoothly broken model lacks (a) the sharp break at $q_{\text{break}} = 0.5$ and (b) the artificial suppression of uncertainty near a pivot point at $q = 0.3$.

7.9 Sensitivity to systematics

The uncertainties we report on mass ratio distribution parameters represent formal fitting uncertainties due to Poisson errors, but they do not include various systematic uncertainties due to modeling choices we make that are held fixed during fitting. Here we vary several aspects of the model to assess how sensitively our constraints depend on them. Figure 7.21 shows the effects of varying the assumed IMF, stellar age distribution, and the isochrones used to generate synthetic photometry, on our parameter constraints. We show a single primary mass and separation bin with typical uncertainties and sensitivity to systematics.

Varying the IMF (blue) has very weak effects on our constraints, primarily because we fit narrow bins in primary mass independently. Our constraints are also relatively insensitive to the assumed star formation history (brown). Varying the age distribution has several effects: it changes the number of low-mass companions that are still undergoing Kelvin-Helmholtz contraction and are thus brighter, it changes the mass-luminosity relation at $M \gtrsim 0.7M_{\odot}$, where age effects are non-negligible; and it changes the distance distribution implied in our model because we adopt an age-dependent scale height (Equation 7.3). Varying the model for unresolved subsystems has similarly weak effects. The green lines in Figure 7.21 show constraints obtained when unresolved subsystems are ignored entirely in the model (i.e. none of the synthetic wide binary components are assigned an unresolved companion). The thus-obtained constraints are nearly indistinguishable from those obtained when subsystems are included in the model.

Changing the adopted stellar models from PARSEC to MIST (Choi et al. 2016) has the strongest effect on our constraints. Because the two models predict modestly different mass-luminosity relations, a given ΔG corresponds to different q in the two models. This has weak effects on the constraints on F_{twin} and q_{twin} : twins necessarily have similar masses, and the mapping between luminosity and mass varies less over a small range of masses than over a large one. Differences between the models are largest for low-mass stars, where isochrones are known to be more uncertain.

We note that while there clearly *are* uncertainties in our results associated with the stellar models, the PARSEC models appear to fit the *Gaia* data for low-mass stars significantly better than the MIST models. Particularly on the lower main sequence $M_G \gtrsim 10$, we are unable to reproduce the morphology of the CMD unless we assume a higher-than-expected mean metallicity of $\langle [\text{Fe}/\text{H}] \rangle \approx +0.2$. This is why we use the PARSEC synthetic photometry in our fiducial model. We note that the PARSEC models use surface boundary conditions that are empirically calibrated to match the observed mass-radius relation at low masses (Chen et al. 2014); without this calibration, the tension between their predictions and the observed lower main sequence is similar to that found for MIST models (e.g. Choi et al. 2018).

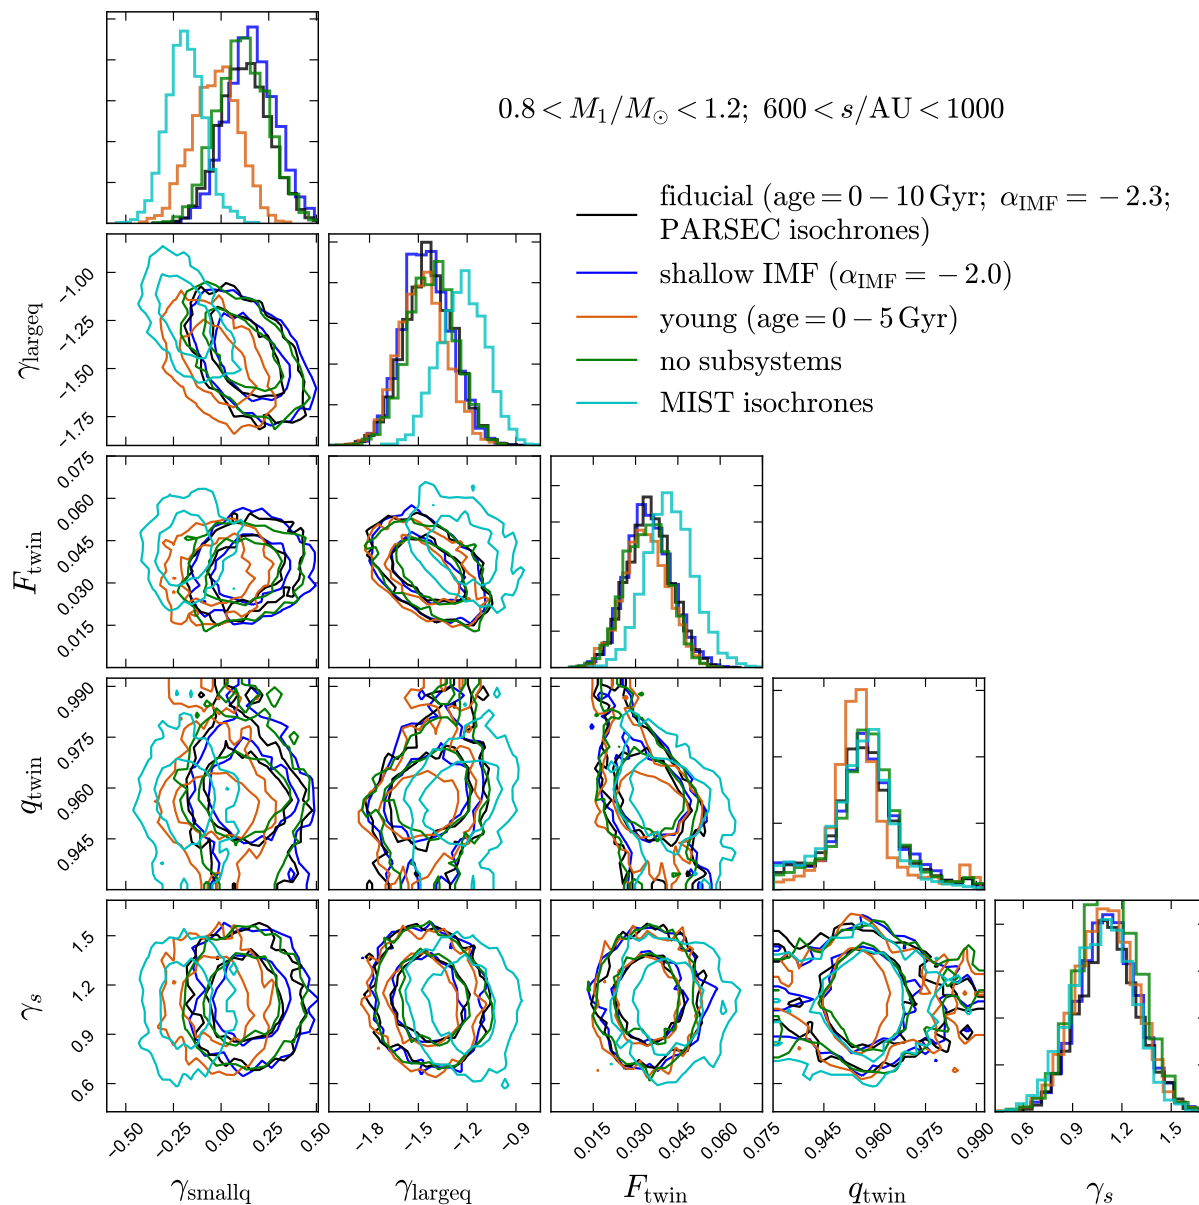


Figure 7.21: Sixty-eight and ninety-five per cent probability contours for one representative bin of primary mass and physical separation. Panels on the diagonal show marginalized probability distributions. Black contours show results for fitting the fiducial model. Blue, brown, green, and cyan contours show constraints obtained when the assumed IMF, star formation history, model for unresolved binaries, and stellar models are varied. Overall, constraints are not very sensitive to the assumed IMF, SFH, or unresolved binary model; this is particularly true for F_{twin} and q_{twin} . They are moderately sensitive to the choice of stellar models, particularly at low mass ratios.

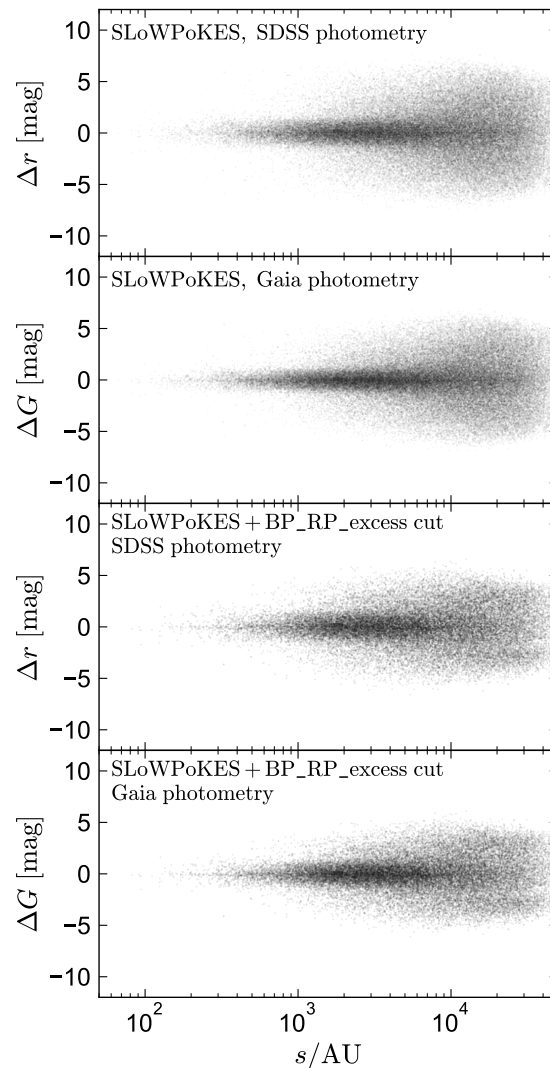


Figure 7.22: Excess of twins in the SLOWPoKES-II catalog (Dhital et al. 2015). Top panel shows the full catalog with SDSS photometry. Here there is no obvious excess of equal-brightness twins. Middle panel shows Gaia photometry from the full SLOWPoKES-II catalog. The Gaia photometry is significantly more precise than the ground-based SDSS photometry, so a slight excess of twins is apparent, but the signal is somewhat weaker than that found in our catalog. The bottom two panels show SDSS and Gaia photometry for the subset of the SLOWPoKES-II catalog for which the Gaia photometry passes our cut on `phot_bp_rp_excess_factor`. This cut removes binaries in which the photometry for either component is contaminated (either by the binary companion or by a background star). Once sources with contaminated photometry have been removed, an excess of twins is visible in both the SDSS and Gaia photometry.




Figure 7.23: Ratio of the number of binaries with nearly equal magnitudes ($\Delta G < 0.25$) to the number with slightly different magnitude ($0.25 < \Delta G < 0.5$) as a function of projected physical separation (right). Error bars are 1 sigma. We compare the fiducial sample from this work (solid black line) to the subsample of the SLoWPoKES-II catalog (Dhital et al. 2015) with uncontaminated Gaia photometry (dashed red line; bottom panel of Figure 7.22) and the sample of Gaia binaries from El-Badry & Rix (2019b), which reaches closer angular separations than our fiducial sample but passes less stringent quality cuts.

7.10 Twins in other catalogs

Another large sample of low-mass wide binaries is the SLoWPoKES-II catalog created by (Dhital et al. 2015) using SDSS photometry. This catalog was produced without astrometry and extends to larger distances than our catalog ($d \lesssim 1$ kpc). Despite the lack of parallaxes and proper motions Dhital et al. (2015) were able to ensure relatively low contamination by limiting their search to close angular separations ($\theta < 20$ arcsec) and using an isochrone prior (i.e., requiring both components of a binary to have similar photometric distance). The selection function of their catalog is quite different from ours, so we do not attempt a full probabilistic model of the mass ratio distribution. However, we check whether a twin excess is visible in the magnitude difference distributions of their catalog and whether its strength at fixed physical separation is consistent with that found in our catalog.

Figure 7.22 (top panel) shows the distribution of Δr , the difference in apparent r -band magnitude between the two components, as a function of physical separation for the full SLoWPoKES-II catalog. Projected physical separations are calculated using the mean photometric distance of the two stars estimate by Dhital et al. (2015). We verified that for stars that are bright enough to have precise *Gaia* parallaxes, the photometric and geometric distances are usually in reasonably good agreement. The SLoWPoKES-II catalog primarily contains binaries with $q \gtrsim 0.5$: there are few pairs with $\Delta r > 5$, and at physical separations $s \lesssim 2000$ AU, most pairs have $\Delta r < 2$. This is primarily a consequence of the catalog selection function. No clear excess of equal-brightness twins is apparent in the top panel of Figure 7.22: the distribution of Δr at fixed separation is reasonably smooth.

To check whether the lack of obvious twin excess in the full SLoWPoKES-II catalog is a consequence of the SDSS ground-based photometry being poorer, we cross-matched the catalog with *Gaia* DR2. A *Gaia* source was found with 1 arcsecond for both components for $\sim 90\%$ of the binaries in the catalog. In the 2nd panel of Figure 7.22, we plot all of these sources (many of which do not pass the photometric and astrometric quality cuts imposed in our catalog), showing the difference in G -band magnitude. Here, an excess of equal-brightness binaries *is* visible, but it appears somewhat weaker than in our catalog.

Finally, in the 3rd and 4th panels of Figure 7.22, we show the SDSS and *Gaia* photometry for the subset of the SLoWPoKES-II catalog in which the *Gaia* photometry for both components passes the quality cut on `phot_bp_rp_excess_factor` that is imposed on our catalog. This cut removes objects in which the sum of the flux in the BP and RP bands is not consistent with the flux in the G -band (see Evans et al. 2018). Because the fluxes in the BP and RP bands are dispersed over several arcsec while the G -band flux is obtained by profile fitting a narrower image, this cut efficiently selects objects in which the *Gaia* photometry (both BP/RP and G -band) is contaminated by a nearby source. Once this cut is applied, a stronger excess of equal-brightness binaries is apparent in both the *Gaia* and SDSS photometry. We show below that for the subsample of the SLoWPoKES-II catalog with uncontaminated *Gaia* photometry, the twin excess at fixed separation is consistent with that found in our catalog.

A natural worry is that the cut on `phot_bp_rp_excess_factor` could somehow select

against pairs that are close on the sky and do not have nearly identical brightness or color, thus erroneously producing an apparent excess of equal-brightness binaries. However, the fact that no excess of equal-brightness pairs is found for chance alignments subject to the same cut speaks against this possibility. Inspecting the SDSS images of binaries in which the *Gaia* photometry for either component does not pass the `phot_bp_rp_excess_factor` cut, we find that in most cases the contamination is quite strong: either the light from the two stars is blended, or another source is blended with one of the components. In such cases, the individual magnitudes of the components cannot be measured with high fidelity, likely leading to underestimated photometric errors.

In Figure 7.23, we compare the strength of the twin excess found in the *Gaia* photometry for the SLoWPoKES-II catalog (once the `phot_bp_rp_excess_factor` cut is applied) at fixed separation (red) to that found in our catalog (black). The y-axis is similar to that in Figure 7.5. The strength of the twin excess is consistent between the two catalogs at all separations. We note that at close physical separations, binaries are only spatially resolvable if they are nearby, so at $s \lesssim 500$ AU, there is substantial overlap between the two catalogs. At $s \gtrsim 1000$ AU, there is little overlap, because most of the SLoWPoKES-II binaries are too distant and faint to enter our catalog.

We also show in Figure 7.23 the excess of equal-brightness binaries found in the catalog of *Gaia* binaries constructed in El-Badry & Rix (2019b), which targeted closer separations. This catalog only contains binaries with $s < 500$ AU. Unlike our primary catalog, it did not require the components to have measured G_{BP} and G_{RP} magnitudes, and it did not apply any cuts on `phot_bp_rp_excess_factor`. This makes it possible to reach binaries with a factor of ~ 4 closer angular separations than our primary catalog. Figure 7.23 shows that in this catalog, the twin excess continues to increase toward smaller separations, at least to $s \approx 50$ AU. The twin excess is slightly weaker at fixed separation in this catalog, likely because the photometry is more contaminated, but it is strongly inconsistent with 0. This provides further evidence that the twin excess is not the result of the `phot_bp_rp_excess_factor` cut creating a bias against binaries with unequal brightness or color.

In addition to the comparisons to other surveys described above, we have also verified that the twin excess we find is not an artifact of the *Gaia* photometry. Cross-matching our full binary catalog with the Pan-STARRS1 survey (Chambers et al. 2016; Flewelling et al. 2016), we find a twin excess of similar strength in its photometry. The excess is also apparent in SDSS photometry, but it is somewhat less narrow there due to the larger typical photometric errors.

7.11 Selection function

The selection function of *Gaia* DR2 is known to exhibit spatial variation on small scales as a result of the scanning law (e.g. Arenou et al. 2018a). Modeling all the small-scale structure in the selection function is beyond the scope of this work, but doing so is not necessary for our purposes. First, spatial variations are almost negligible for sources within

the magnitude range of our catalog ($G \lesssim 18$). More generally, since we do not expect the *intrinsic* properties of the binary population to vary much with on-sky position (particularly on small scales), one can construct an effective selection function averaged over the whole sky without introducing biases (e.g. [Bovy 2016a](#)).

We model the selection function for binaries as the product of two single-star terms and a cross-term that depends on the angular separation and magnitude difference of the two stars (Equation 7.2). Both terms are described below.

7.11.1 Single-star term

The single-star term (s_1 and s_2 in Equation 7.2; here we default to s_1) is set primarily by the following cuts we imposed in creating the catalog:

1. 5% parallax uncertainty: `parallax_over_error > 20`
2. precise photometry: `phot_bp_mean_flux_over_error > 20`, `phot_rp_mean_flux_over_error > 20`, and `phot_g_mean_flux_over_error > 20`
3. good astrometric model fits; $\sqrt{\chi^2/(\nu' - 5)} < 1.2 \times \max(1, \exp(-0.2(G - 19.5)))$, where χ^2 and ν' are respectively referred to as `astrometric_chi2_al` and `astrometric_n_good_obs_al` in the *Gaia* archive.
4. uncontaminated photometry; $1.0 + 0.015(G_{\text{BP}} - G_{\text{RP}})^2 < \text{phot_bp_rp_excess_factor} < 1.3 + 0.06(G_{\text{BP}} - G_{\text{RP}})^2$.

The motivation for these cuts is described in [Lindgren et al. \(2018\)](#). We discuss the effects of each cut below. We note that [ER18](#) also required difference in total proper motion of the two stars to be precise, satisfying $\sigma_{\Delta\mu} < 1.5 \text{ mas yr}^{-1}$. We find that for the subset of the binaries studied in this work, that cut has a negligible effect on the selection function, as sources that satisfy (i) and (ii) already satisfy it.

For (i), parallax error is expected to depend mainly on apparent magnitude. Figure 7.24 shows that the parallax error at a given G -band magnitude roughly follows a lognormal distribution, the mean value of which increases for fainter stars. This allows us to calculate the fraction of stars at a given distance and magnitude that will have `parallax_over_error > 20`. In particular, the distribution of ϖ/σ_ϖ at a given distance and magnitude is

$$P(\varpi/\sigma_\varpi) = P(\log \sigma_\varpi) \left| \frac{d \log \sigma_\varpi}{d(\varpi/\sigma_\varpi)} \right|, \quad (7.5)$$

where $P(\log \sigma_\varpi)$ is the (Gaussian) distribution of \log parallax error. The fraction of stars with $\varpi/\sigma_\varpi > 20$ is then

$$s_\varpi(G) = \int_{20}^{\infty} P(\varpi/\sigma_\varpi) d(\varpi/\sigma_\varpi) \quad (7.6)$$

$$= \frac{1}{2} \left[1 - \text{erf} \left(\frac{\mu_f \ln 10 + \ln(20/\varpi)}{\sqrt{2}\sigma_f \ln 10} \right) \right]. \quad (7.7)$$

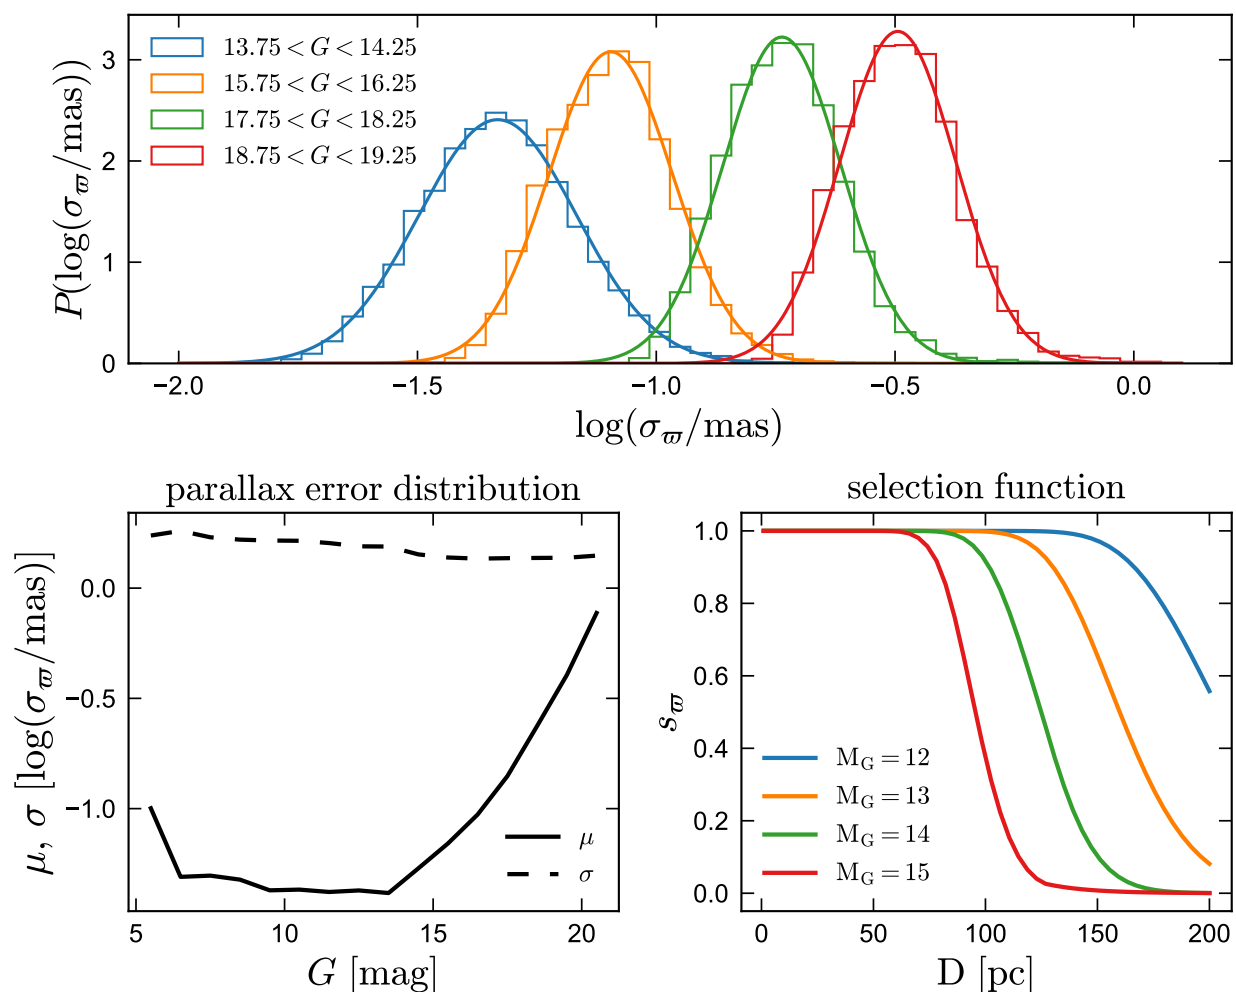


Figure 7.24: Top panel shows the distributions of log parallax error for nearby stars ($D < 100$ pc) in 0.5-mag wide bins of apparent magnitude. We find that the distribution of σ_{ϖ} at fixed apparent magnitude is roughly lognormal; solid curves show Gaussian fits to the observed histograms. Bottom left panel shows the mean and dispersion of the log parallax error distribution as a function of G -band magnitude. Modeling the distribution of parallax error as a lognormal with μ and σ following this panel allows us to calculate what fraction of stars with a given absolute magnitude and distance will pass the cut of `parallax_over_error` > 20 (Equation 7.6; bottom right panel).

Here μ_f and σ_f represent the mean and dispersion of the log parallax error distributions at magnitude G ; i.e., the quantities plotted in the lower left panel of Figure 7.24. Given Equation (7.6), we can calculate the probability that a star with given absolute magnitude at a given distance will satisfy `parallax_over_error` > 20 ; this is shown in the lower right panel of Figure 7.24.

A similar strategy can be used to calculate the fraction of stars passing the photometric precision cuts. We find that the cut on `phot_bp_mean_flux_over_error` is the most important: once it is applied and white dwarfs are excluded, all the sources that pass it also pass the other cuts in (ii). This quantity depends primarily on the apparent magnitude in the G_{BP} band. As Figure 7.25 shows, the distribution of `phot_bp_mean_flux_over_error` at a given G_{BP} is also roughly lognormal. We use the empirical dependence of the distributions of `phot_bp_mean_flux_over_error` on G_{BP} from the bottom left panel of Figure 7.25 to predict the distribution of `phot_bp_mean_flux_over_error` for a hypothetical star with absolute BP-band magnitude at a given distance, and then calculate s_{BP} as the fraction of that distribution that exceeds the adopted threshold of 20:

$$s_{\text{BP}}(G_{\text{BP}}) = \int_{20}^{\infty} P(\log(F_{\text{BP}}/\sigma_{\text{BP}})) \, d \log(F_{\text{BP}}/\sigma_{\text{BP}}) \quad (7.8)$$

$$= \frac{1}{2} \left[1 + \operatorname{erf} \left(\frac{\mu_f - \log 20}{\sqrt{2}\sigma_f} \right) \right]. \quad (7.9)$$

Here $F_{\text{BP}}/\sigma_{\text{BP}}$ represents `phot_bp_mean_flux_over_error`, and μ_f and σ_f represent the mean and dispersion of the log `phot_bp_mean_flux_over_error` distributions at apparent BP-band magnitude G_{BP} ; i.e., the quantities plotted in the lower left panel of Figure 7.25.

It is also worth considering whether the quality cuts (iii) and (iv) change the single-star selection function. We assess the effects of each quality cut as follows. We begin with the full sample of stars within 200 pc satisfying $\varpi/\sigma_{\varpi} > 10$. We then apply each quality cut and monitor the effects on the sample. For this nearby sample, it is straightforward to determine whether most of the sources removed by a particular cut are erroneous (generally distant faint stars with incorrect parallaxes) or real, because distant sources with incorrect parallaxes generally fall in a cloud below the main sequence (see Lindegren et al. 2018). What we wish to quantify is the fraction of real sources that are removed by each cut, and whether these cuts introduce systematic biases.

Figure 7.26 shows the fraction of suspected genuine sources removed by cuts (iii) and (iv). We conclude that cut (iv) has a negligible effect on the single-star selection function. The effects of cut (iii) are also weak, but it does appear to preferentially remove nearby sources. This may be because unresolved astrometric binaries produce large deviations from linear motion of the light centroid at close distances. We tabulate s_{cut} due to cut (iv) as a function of parallax. The total single-star selection function is then

$$s_1 = s_{\varpi}(G) s_{\text{BP}}(G_{\text{BP}}) s_{\text{cut}}(\varpi), \quad (7.10)$$

where $s_{\varpi}(G)$ and $s_{\text{BP}}(G_{\text{BP}})$ are given by Equations (7.6) and (7.8), respectively.

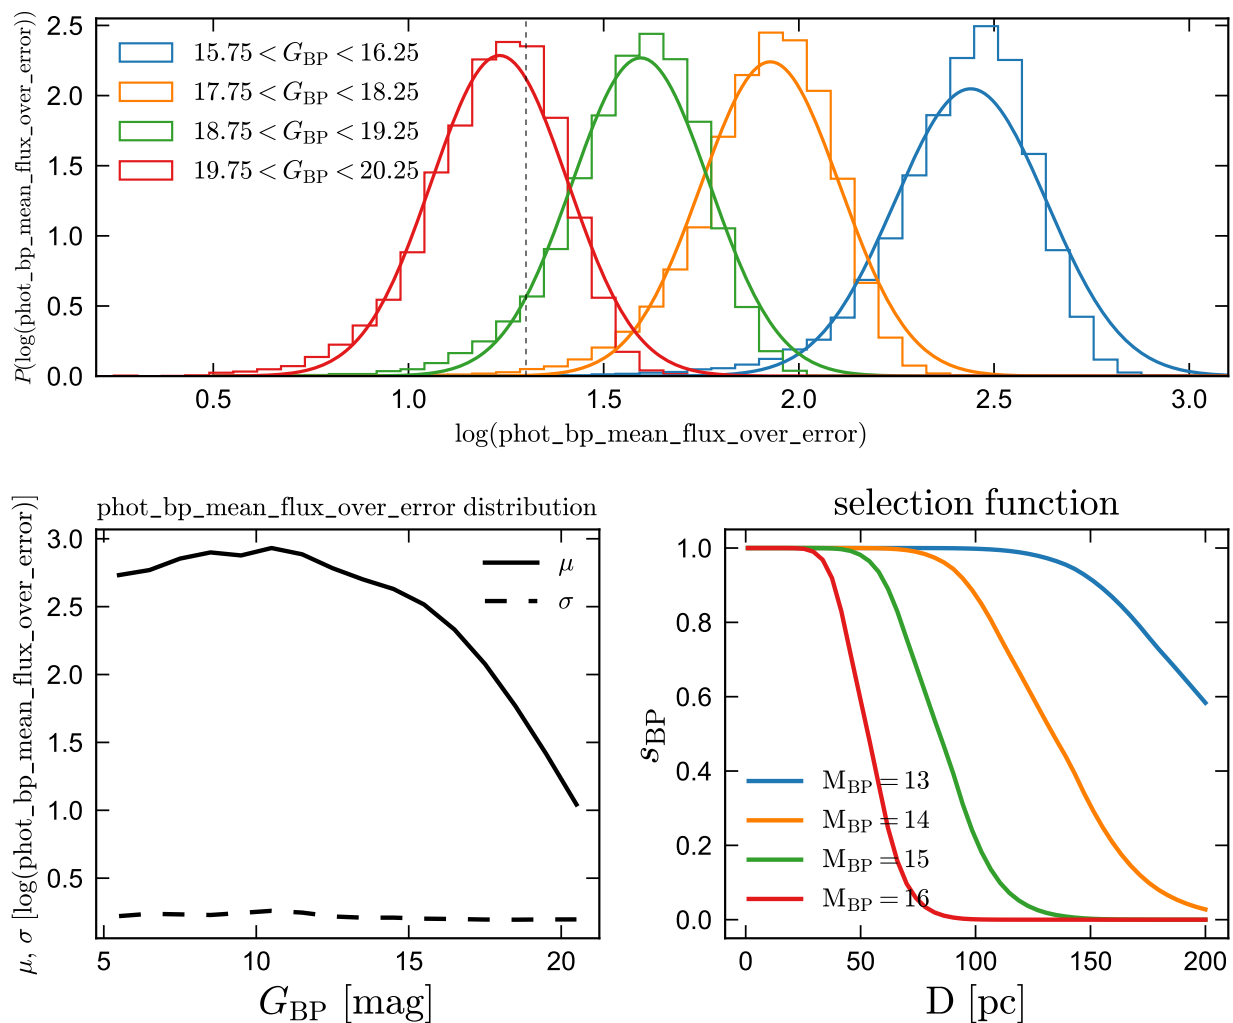


Figure 7.25: Top panel shows distributions of `phot_bp_mean_flux_over_error` for nearby stars ($D < 100$ pc) in different 0.5 mag wide bins of G_{BP} apparent magnitude. Solid lines show lognormal fits; vertical dashed line shows the cut of `phot_bp_mean_flux_over_error` > 20 that both components of a binary must satisfy to enter the catalog. Bottom left panel shows μ and σ of lognormal fits to these distributions as a function of G_{BP} . From these, we calculate the probability that a star at a given distance and absolute magnitude M_{BP} will have `phot_bp_mean_flux_over_error` > 20 (bottom left panel).




Figure 7.26: Fraction of suspected genuine sources within 200 pc that survive the cuts we impose on `astrometric_chi2_al` and `bp_rp_excess_factor`, as a function of apparent magnitude (left) and distance (right). Both cuts remove only a small fraction of genuine sources. The cut on `bp_rp_excess_factor` removes sources with contaminated photometry (section 7.11.2). The cut on `astrometric_chi2_al` preferentially removes nearby sources, perhaps because perturbations from a companion lead to the larger angular deviations from linear motion for nearby sources.




Figure 7.27: Contrast sensitivity given our quality cuts. θ_0 is the angular separation below which the sensitivity to a companion with magnitude difference ΔG drops off rapidly (see Equation 7.11). Companions can be detected at closer separations when the magnitude difference is small. We compare the values of θ_0 found by ER18 (for all stars passing our quality cuts in a dense field) to those derived from our all-sky catalog of nearby chance alignments with precise parallaxes.

7.11.2 Contrast sensitivity for close pairs

A critical aspect of the selection function for binaries is the reduction in sensitivity to a companion at close angular separations. Whether a binary enters our catalog depends both on the angular separation of the two stars and on their flux ratio, as a secondary is more likely to be outshone or contaminated by light from the primary at fixed angular separation if the flux ratio is large than if it is small.

The sensitivity to a companion can be measured by comparing the two point correlation function of chance alignments to what would be expected in the absence of crowding/blending effects (e.g. Arenou et al. 2018a; Brandeker & Cataldi 2019a). Of course, the sensitivity to companions in a particular catalog depends on the quality cuts imposed: cuts that remove objects with somewhat contaminated photometry will lower the sensitivity.

We quantified the reduction in sensitivity to a companion as a function of angular separation θ and magnitude difference ΔG in ER18. There we found that, given our cuts on astrometric χ^2 and photometric quality (primarily `bp_rp_excess_factor`), the sensitivity to a companion goes to 0 at $\theta \ll \theta_0$, where $\theta_0 \approx 2$ arcsec for sources with similar magnitude, and θ_0 increases with ΔG . We quantified this dependence by fitting a function:

$$s_{\Delta G}(\theta) = \frac{1}{1 + (\theta/\theta_0)^{-\beta}}, \quad (7.11)$$

where $\beta \approx 10$ and we fit for θ_0 as a function of ΔG . We note that the angular resolution of *Gaia* DR2 is actually significantly better than 2 arcsec: most companions are detected down 1 arcsec separations, and the detection fraction only drops to zero at $\theta < 0.5$ arcsec (Arenou et al. 2018a; Ziegler et al. 2018). The ~ 2 arcsec limit for our catalog is a result of our requirement that both stars have a measured `bp_rp` color and the cut on `bp_rp_excess_factor`.

The dependence of θ_0 on ΔG calculated in ER18 was derived from the source counts of all sources passing our quality cuts (iii) and (iv) in a dense field, most of which are fainter than the objects in the binary catalog. Here, we improve slightly on the ER18 measurement by repeating their analysis, but using the sources from the chance alignment catalog described in Section 7.3 instead of all sources in a dense field. The advantage of this approach is that the sources in the chance alignment catalog, being nearby and having precise parallaxes, are more representative of objects in the binary catalog. Unsurprisingly – since it was verified in ER18 that there are no strong trends in contrast sensitivity with apparent magnitude or source density; see also Brandeker & Cataldi (2019a) – the improved constraints are fairly similar to those derived in ER18. We compare them in Figure 7.27. On average, we find that the angular resolution at fixed separation is marginally better for the bright chance alignments than was found in ER18.

7.12 Model validation

To test whether the assumptions of the model we use to fit for the intrinsic mass ratio distribution are valid, we use it to predict properties of the population of single stars that pass the same quality cuts as the binaries in our catalog (Section 7.12.1) and the separation distribution of chance alignments (Section 7.12.2).

7.12.1 Single stars

We query the *Gaia* catalog for all stars with $\varpi > 5$ mas that pass the cuts on astrometric and photometric quality and precision that we require both members of the binary catalog to pass. Distributions of their distance, apparent magnitude, absolute magnitude, and color are shown in Figure 7.28 with a black histogram. To compare to the model predictions, we draw masses, ages, metallicities, distances from the fiducial distributions described in Section 7.4, compute synthetic photometry using PARSEC isochrones, and pass the observables through




Figure 7.28: Validation of the single-star selection function and underlying Galactic model. Black histograms show distributions of distance, apparent magnitude, absolute magnitude, and color for all stars (not just binaries) within 200 pc that pass the quality cuts of the binary catalog. Black line show the distributions predicted by our model; i.e., assuming the same single-star selection function, IMF, Galactic scale height, star formation history, and metallicity distribution that we use when fitting the binary population. The reasonably good agreement with the observed distributions suggests that our empirical selection function and Galactic model are reasonable. Our fiducial model slightly overpredicts the number of faint stars. This tension can be resolved if a shallower IMF is assumed (with a logarithmic slope of -0.9 instead of -1.3 at $M < 0.5M_{\odot}$; gold). Overall, we regard the agreement between the fiducial model and observations as quite satisfactory.

the single-star selection function (Appendix 7.4.2). The resulting selection-function weighted distributions are compared to the data in Figure 7.28 (red histogram). In the upper left panel, we also show the distance distributions predicted for a uniform spatial distribution and for an exponential disk with scale height of 300 pc (comparable to the Milky Way; see Jurić et al. 2008a), both assuming no incompleteness. The distance distribution of single stars is not that different from the exponential disk prediction, but it begins to deviate at $d \gtrsim 130$ pc, which is where incompleteness effects become significant for $M_G \approx 14$ (Figure 7.24).

The agreement between model and data is reasonably good. However, the fiducial model predicts slightly too many faint stars and too few bright stars. The gold histograms shows that the agreement can be improved if a slightly shallower IMF is assumed, with a logarithmic slope of -0.9 (instead of -1.3, as assumed in the fiducial model) at $M < 0.5 M_\odot$.¹⁰ This is consistent with the recent result from Sollima (2019), who used the *Gaia* nearby star sample to measure the IMF. For the sake of this work, we are agnostic of whether Figure 7.28 indicates that the low-mass IMF in the solar neighborhood is slightly shallower than that assumed in our fiducial model or points towards a systematic in some other aspect of the model (such as the metallicity distribution function or stellar models). We note that changing the assumed IMF has very little effect on our inferred mass ratio distribution (Figure 7.21).

We also note that the color distributions predicted by our model do not exactly match the observed distribution for either choice of IMF. This could be due to reddening, the adopted MDF (as the subset of stars with spectroscopic metallicity measurements is not guaranteed to be an unbiased sample), or imperfect stellar models.

7.12.2 Chance alignments

To validate our model for the contrast sensitivity as a function of angular separation, we predict the distributions of magnitude difference at fixed separation using our model and compare to the chance alignment catalog described in Section 7.3. This is accomplished as follows. We draw masses, ages, distances, and metallicities for both “components” of a chance alignment independently, and we draw angular separations assuming $P(\theta) d\theta \propto 2\pi\theta$. Just as for true binaries, we compute the selection function for each pair by multiplying the two single-star selection functions and the angular contrast sensitivity term. The results are shown in Figure 7.29 σ_0 . The agreement with the real chance alignment catalog is good.

7.13 RV variability

Although *Gaia* DR2 does not contain multi-epoch radial velocity measurements, the published radial velocity uncertainties contain information that can be exploited to detect a large fraction of RV-variable close binaries. The `radial_velocity_error` reported in the *Gaia* archive represents the uncertainty on the median of velocity measurements from several

¹⁰This slope was found by fitting the IMF from the CMD, as described in Appendix C of El-Badry et al. (2018a).




Figure 7.29: Distributions of apparent magnitude difference for our observed sample of chance alignments within 400 pc (black) and a simulated sample based on the same selection function and distributions of mass, age, metallicity, and distance assumed in modeling genuine binaries. Agreement between data and model validates the model assumptions, particularly the sensitivity to a companion as a function of angular separation and magnitude difference.

transits (see [Katz et al. 2019](#)). It is calculated as

$$\epsilon_{\text{RV}} = \left[\frac{\pi}{2N_{\text{obs}}} \sigma_{\text{RV}}^2 + \sigma_0^2 \right]^{1/2}, \quad (7.12)$$

where $\sigma_0 = 0.11 \text{ km s}^{-1}$ is a constant term that represents the minimum achievable RV uncertainty due to calibration issues, N_{obs} is the number of radial velocity transits (`rv_nb_transits` in the *Gaia* archive), and σ_{RV} is the standard deviation of the RVs measured in individual transits. The standard deviation of the measured transit RVs can thus be reconstructed as

$$\sigma_{\text{RV}} = [2N_{\text{obs}} (\epsilon_{\text{RV}}^2 - \sigma_0^2) / \pi]^{1/2}. \quad (7.13)$$

We expect σ_{RV} to be larger than usual if the variation in the true radial velocity of a star between transits is large compared to the observational RV precision. This suggests that close binaries could be identified as stars with unusually large σ_{RV} for their color and apparent magnitude.

To quantify this, we queried the *Gaia* archive for all stars within 200 pc that pass the quality cuts imposed on binary components and additionally have `rv_nb_transits` > 2. Selecting main-sequence stars in bins of `bp_rp` color and *G*-band magnitude, we calculated the median σ_{RV} in each bin (Figure 7.30). As expected, the typical RV error increases with increasing *G* magnitude and is larger for bluer stars, which have weaker and broader absorption lines, at fixed magnitude.

We designate sources that have σ_{RV} larger than 2.5 times the median for their color and apparent magnitude as likely close binaries. The factor of 2.5 is a practical choice to balance the number of false-positives and false-negatives. To assess the false-negative rate for this designation, we cross-matched *Gaia* DR2 with the catalog of RV-variable main-sequence SB1s identified by [El-Badry et al. \(2018g\)](#) using APOGEE spectra and calculated σ_{RV} for the subsample of that catalog that passes our quality cuts. Among SB1s for which [El-Badry et al. \(2018g\)](#) found the radial velocity to vary by at least 5 km s^{-1} between visits, 62% are correctly identified as binaries based on the *Gaia* σ_{RV} . This fraction climbs to 82% for sources whose APOGEE RVs varied by at least 15 km s^{-1} . On the other hand, the false-positive rate is relatively low: only 2.6% of the stars classified by [El-Badry et al. \(2018g\)](#) as likely to be single are classified as binaries based on the *Gaia* σ_{RV} . This means that, although some true binaries will be missed and there will be some false positives, the *Gaia* σ_{RV} can be used to obtain an estimate of the close binary fraction in a population (where “close” means $a \lesssim$ few AU).

To assess whether the components of wide twins are more likely to have an unresolved close companion than the components of non-twins, we use the σ_{RV} distributions to estimate the fraction of components of twin ($\Delta G < 0.25$) and non-twin ($0.25 \leq \Delta G < 1$) wide binaries with an unresolved close companion. We only consider binaries with $100 < s/\text{AU} < 500$, since at wider separations, the number of “excess” twins is subdominant relative to the underlying population (Figure 7.10). We consider all components with masses in the range $0.5 < M/M_{\odot} < 1.2$ that have *Gaia* RVs with `rv_nb_transits` > 2 (i.e., we do not require




Figure 7.30: Median σ_{RV} (Equation 7.13) for main-sequence stars within 200 pc that pass our quality cuts and `rv_nb_transits` > 2. Since a majority of stars are not intrinsically RV-variable, this represents the typical radial velocity precision for stars of a given brightness and spectral type. Most stars that have σ_{RV} substantially larger than the median for their color and magnitude are close binaries.

both resolved components to have measured RVs). Among 546 eligible components of wide twins, 29 have σ_{RV} consistent with having unresolved close companion, implying a close companion fraction per wide binary component of

$$f_{\text{close companion, twins}} = 0.053 \pm 0.012. \quad (7.14)$$

Of the 780 eligible components of non-twins, 76 have σ_{RV} consistent with having unresolved close companion. This implies

$$f_{\text{close companion, non-twins}} = 0.097 \pm 0.013. \quad (7.15)$$

That is, the fraction of wide binary components that have a close unresolved companion is *higher* at the 2 sigma level for non-twins than for twins. This would seem to speak against a scenario in which the excess of wide twins is causally linked to hierarchical triples. Indeed, given the non-zero false positive rate of the σ_{RV} -based close binary identification and the fact that “excess” twins represent only about half of the population with $\Delta G < 0.25$, the data are consistent with the excess twins having no close companions at all.

7.14 Full fitting constraints

Constraints on fitting parameters for all bins of primary mass and separation are listed in Table 7.2. Error bars are 2 sigma (middle 95.4%).

Table 7.2: Marginalized 2 sigma constraints on fitting parameters for all bins of primary mass and separation. For $0.1 < M_1/M_\odot < 0.4$, we fit a single power law, so $\gamma_{\text{smallq}} = \gamma_{\text{largeq}}$. The units of M_1 are M_\odot in all columns.

	$0.1 < M_1 < 0.4$	$0.4 < M_1 < 0.6$	$0.6 < M_1 < 0.8$	$0.8 < M_1 < 1.2$	$1.2 < M_1 < 2.5$
F_{twin}					
$50 < s/\text{AU} < 350$	$0.189 \pm_{-0.049}^{+0.042}$	$0.112 \pm_{-0.038}^{+0.034}$	$0.059 \pm_{-0.021}^{+0.019}$	$0.101 \pm_{-0.029}^{+0.027}$	$0.086 \pm_{-0.075}^{+0.058}$
$350 < s/\text{AU} < 600$	$0.105 \pm_{-0.035}^{+0.036}$	$0.062 \pm_{-0.020}^{+0.018}$	$0.043 \pm_{-0.019}^{+0.016}$	$0.054 \pm_{-0.018}^{+0.014}$	$0.088 \pm_{-0.047}^{+0.037}$
$600 < s/\text{AU} < 1,000$	$0.049 \pm_{-0.032}^{+0.039}$	$0.054 \pm_{-0.020}^{+0.017}$	$0.032 \pm_{-0.019}^{+0.016}$	$0.035 \pm_{-0.015}^{+0.015}$	$0.015 \pm_{-0.034}^{+0.024}$
$1,000 < s/\text{AU} < 2,500$	$0.023 \pm_{-0.033}^{+0.025}$	$0.049 \pm_{-0.022}^{+0.018}$	$0.024 \pm_{-0.016}^{+0.016}$	$0.013 \pm_{-0.011}^{+0.010}$	$0.002 \pm_{-0.020}^{+0.015}$
$2,500 < s/\text{AU} < 5,000$	$-0.001 \pm_{-0.048}^{+0.039}$	$0.047 \pm_{-0.025}^{+0.022}$	$0.018 \pm_{-0.020}^{+0.019}$	$0.023 \pm_{-0.017}^{+0.021}$	$-0.003 \pm_{-0.023}^{+0.023}$
$5,000 < s/\text{AU} < 15,000$	$0.005 \pm_{-0.043}^{+0.049}$	$0.043 \pm_{-0.031}^{+0.030}$	$0.014 \pm_{-0.021}^{+0.018}$	$0.007 \pm_{-0.015}^{+0.015}$	$0.003 \pm_{-0.035}^{+0.023}$
$15,000 < s/\text{AU} < 50,000$	$0.024 \pm_{-0.120}^{+0.094}$	$-0.009 \pm_{-0.043}^{+0.039}$	$-0.012 \pm_{-0.032}^{+0.028}$	$0.009 \pm_{-0.032}^{+0.024}$	$0.001 \pm_{-0.057}^{+0.037}$
q_{twin}					
$50 < s/\text{AU} < 350$	$0.953 \pm_{-0.006}^{+0.010}$	$0.959 \pm_{-0.014}^{+0.023}$	$0.969 \pm_{-0.005}^{+0.019}$	$0.962 \pm_{-0.009}^{+0.014}$	$0.950 \pm_{-0.029}^{+0.018}$
$350 < s/\text{AU} < 600$	$0.937 \pm_{-0.011}^{+0.006}$	$0.959 \pm_{-0.025}^{+0.017}$	$0.954 \pm_{-0.016}^{+0.021}$	$0.967 \pm_{-0.010}^{+0.027}$	$0.936 \pm_{-0.018}^{+0.006}$
$600 < s/\text{AU} < 1,000$	$0.954 \pm_{-0.025}^{+0.022}$	$0.963 \pm_{-0.012}^{+0.025}$	$0.939 \pm_{-0.018}^{+0.008}$	$0.957 \pm_{-0.028}^{+0.023}$	$0.954 \pm_{-0.040}^{+0.022}$
$1,000 < s/\text{AU} < 2,500$	$0.975 \pm_{-0.017}^{+0.043}$	$0.963 \pm_{-0.015}^{+0.029}$	$0.941 \pm_{-0.030}^{+0.011}$	$0.959 \pm_{-0.030}^{+0.025}$	$0.957 \pm_{-0.037}^{+0.026}$
$2,500 < s/\text{AU} < 5,000$	$0.955 \pm_{-0.039}^{+0.025}$	$0.953 \pm_{-0.019}^{+0.015}$	$0.950 \pm_{-0.029}^{+0.030}$	$0.945 \pm_{-0.039}^{+0.030}$	$0.958 \pm_{-0.035}^{+0.027}$
$5,000 < s/\text{AU} < 15,000$	$0.959 \pm_{-0.038}^{+0.028}$	$0.951 \pm_{-0.030}^{+0.020}$	$0.965 \pm_{-0.027}^{+0.034}$	$0.970 \pm_{-0.027}^{+0.038}$	$0.960 \pm_{-0.035}^{+0.029}$
$15,000 < s/\text{AU} < 50,000$	$0.961 \pm_{-0.031}^{+0.030}$	$0.961 \pm_{-0.033}^{+0.029}$	$0.961 \pm_{-0.032}^{+0.029}$	$0.957 \pm_{-0.038}^{+0.026}$	$0.957 \pm_{-0.039}^{+0.026}$
γ_{largeq}					
$50 < s/\text{AU} < 350$	$0.52 \pm_{-0.45}^{+0.45}$	$0.17 \pm_{-0.39}^{+0.40}$	$-1.43 \pm_{-0.51}^{+0.48}$	$-1.22 \pm_{-0.67}^{+0.66}$	$-0.89 \pm_{-1.00}^{+0.90}$
$350 < s/\text{AU} < 600$	$-0.16 \pm_{-0.43}^{+0.43}$	$-0.01 \pm_{-0.27}^{+0.27}$	$-0.89 \pm_{-0.33}^{+0.33}$	$-1.83 \pm_{-0.42}^{+0.42}$	$-1.88 \pm_{-0.54}^{+0.54}$
$600 < s/\text{AU} < 1,000$	$0.23 \pm_{-0.45}^{+0.41}$	$-0.44 \pm_{-0.27}^{+0.28}$	$-1.19 \pm_{-0.30}^{+0.30}$	$-1.43 \pm_{-0.30}^{+0.30}$	$-1.16 \pm_{-0.28}^{+0.33}$
$1,000 < s/\text{AU} < 2,500$	$0.39 \pm_{-0.41}^{+0.32}$	$-0.55 \pm_{-0.40}^{+0.28}$	$-0.90 \pm_{-0.26}^{+0.26}$	$-1.54 \pm_{-0.22}^{+0.22}$	$-1.55 \pm_{-0.19}^{+0.21}$
$2,500 < s/\text{AU} < 5,000$	$0.23 \pm_{-0.52}^{+0.32}$	$-0.43 \pm_{-0.38}^{+0.40}$	$-1.04 \pm_{-0.34}^{+0.36}$	$-1.52 \pm_{-0.33}^{+0.32}$	$-1.52 \pm_{-0.27}^{+0.29}$
$5,000 < s/\text{AU} < 15,000$	$0.22 \pm_{-0.59}^{+0.51}$	$-0.42 \pm_{-0.42}^{+0.44}$	$-0.86 \pm_{-0.35}^{+0.37}$	$-1.35 \pm_{-0.31}^{+0.33}$	$-1.31 \pm_{-0.25}^{+0.28}$
$15,000 < s/\text{AU} < 50,000$	$0.53 \pm_{-0.84}^{+0.86}$	$-0.43 \pm_{-0.67}^{+0.67}$	$-0.72 \pm_{-0.63}^{+0.65}$	$-1.39 \pm_{-0.58}^{+0.55}$	$-1.22 \pm_{-0.43}^{+0.46}$
γ_{smallq}					
$50 < s/\text{AU} < 350$		$0.33 \pm_{-0.56}^{+0.60}$	$0.20 \pm_{-0.58}^{+0.51}$	$0.13 \pm_{-0.69}^{+0.64}$	$0.13 \pm_{-1.00}^{+0.98}$
$350 < s/\text{AU} < 600$		$-0.37 \pm_{-0.47}^{+0.44}$	$0.24 \pm_{-0.47}^{+0.41}$	$0.40 \pm_{-0.43}^{+0.41}$	$-0.01 \pm_{-0.86}^{+0.79}$
$600 < s/\text{AU} < 1,000$		$0.17 \pm_{-0.46}^{+0.44}$	$0.24 \pm_{-0.36}^{+0.33}$	$0.12 \pm_{-0.26}^{+0.26}$	$0.15 \pm_{-0.72}^{+0.55}$
$1,000 < s/\text{AU} < 2,500$		$0.53 \pm_{-0.44}^{+0.44}$	$0.06 \pm_{-0.27}^{+0.25}$	$0.12 \pm_{-0.17}^{+0.16}$	$0.56 \pm_{-0.47}^{+0.42}$
$2,500 < s/\text{AU} < 5,000$		$0.02 \pm_{-0.58}^{+0.56}$	$-0.15 \pm_{-0.34}^{+0.33}$	$-0.14 \pm_{-0.21}^{+0.21}$	$0.09 \pm_{-0.43}^{+0.36}$
$5,000 < s/\text{AU} < 15,000$		$0.13 \pm_{-0.69}^{+0.59}$	$0.12 \pm_{-0.40}^{+0.39}$	$-0.20 \pm_{-0.22}^{+0.21}$	$-0.18 \pm_{-0.40}^{+0.40}$
$15,000 < s/\text{AU} < 50,000$		$-0.02 \pm_{-0.88}^{+0.84}$	$-0.13 \pm_{-0.60}^{+0.69}$	$-0.24 \pm_{-0.41}^{+0.42}$	$-0.09 \pm_{-0.66}^{+0.64}$
γ_s					
$50 < s/\text{AU} < 350$	$-1.36 \pm_{-0.21}^{+0.22}$	$-1.55 \pm_{-0.20}^{+0.18}$	$-1.22 \pm_{-0.38}^{+0.32}$	$-1.02 \pm_{-0.46}^{+0.40}$	$-0.84 \pm_{-0.93}^{+0.77}$
$350 < s/\text{AU} < 600$	$-1.10 \pm_{-0.42}^{+0.43}$	$-1.20 \pm_{-0.27}^{+0.29}$	$-0.89 \pm_{-0.38}^{+0.39}$	$-0.80 \pm_{-0.43}^{+0.47}$	$-1.28 \pm_{-1.00}^{+1.00}$
$600 < s/\text{AU} < 1,000$	$-1.48 \pm_{-0.33}^{+0.48}$	$-1.16 \pm_{-0.14}^{+0.31}$	$-1.75 \pm_{-0.15}^{+0.34}$	$-1.13 \pm_{-0.26}^{+0.36}$	$-0.03 \pm_{-0.54}^{+0.74}$
$1,000 < s/\text{AU} < 2,500$	$-1.67 \pm_{-0.26}^{+0.33}$	$-1.56 \pm_{-0.14}^{+0.13}$	$-1.52 \pm_{-0.16}^{+0.15}$	$-1.47 \pm_{-0.13}^{+0.13}$	$-1.41 \pm_{-0.26}^{+0.26}$
$2,500 < s/\text{AU} < 5,000$	$-2.05 \pm_{-0.46}^{+0.51}$	$-1.71 \pm_{-0.30}^{+0.30}$	$-1.75 \pm_{-0.29}^{+0.30}$	$-1.67 \pm_{-0.26}^{+0.30}$	$-1.40 \pm_{-0.43}^{+0.42}$
$5,000 < s/\text{AU} < 15,000$	$-1.74 \pm_{-0.36}^{+0.34}$	$-1.74 \pm_{-0.21}^{+0.30}$	$-1.62 \pm_{-0.21}^{+0.21}$	$-1.84 \pm_{-0.17}^{+0.26}$	$-1.35 \pm_{-0.27}^{+0.43}$
$15,000 < s/\text{AU} < 50,000$	$-1.61 \pm_{-0.83}^{+0.90}$	$-1.68 \pm_{-0.49}^{+0.47}$	$-2.21 \pm_{-0.50}^{+0.54}$	$-2.00 \pm_{-0.37}^{+0.40}$	$-1.50 \pm_{-0.63}^{+0.66}$

Chapter 8

Not so fast: LB-1 is unlikely to contain a $70 M_{\odot}$ black hole

An earlier version of this article was previously published in El-Badry, K. and Quataert, E., MNRASL, 2020, 493, 22.

8.1 Abstract

The recently discovered binary LB-1 has been reported to contain a $\sim 70 M_{\odot}$ black hole (BH). The evidence for the unprecedentedly high mass of the unseen companion comes from reported radial velocity (RV) variability of the $H\alpha$ emission line, which has been proposed to originate from an accretion disk around a BH. We show that there is in fact no evidence for RV variability of the $H\alpha$ emission line, and that its apparent shifts instead originate from shifts in the luminous star’s $H\alpha$ absorption line. If not accounted for, such shifts will cause a stationary emission line to appear to shift in anti-phase with the luminous star. We show that once the template spectrum of a B star is subtracted from the observed Keck/HIRES spectra of LB-1, evidence for RV variability vanishes. Indeed, the data rule out periodic variability of the line with velocity semi-amplitude $K_{H\alpha} > 1.3 \text{ km s}^{-1}$. This strongly suggests that the observed $H\alpha$ emission does not originate primarily from an accretion disk around a BH, and thus that the mass ratio cannot be constrained from the relative velocity amplitudes of the emission and absorption lines. The nature of the unseen companion remains uncertain, but a “normal” stellar-mass BH with mass $5 \lesssim M/M_{\odot} \lesssim 20$ seems most plausible. The $H\alpha$ emission likely originates primarily from circumbinary material, not from either component of the binary.

8.2 Introduction

Characterization of the Milky Way’s stellar-mass black hole (BH) population is a primary aim of numerous stellar surveys. As part of a search for stellar-mass BHs with normal-

star companions, [Liu et al. \(2019\)](#) recently identified an intriguing BH-candidate (“LB-1”) consisting of a B star with an unseen companion. The B star is on a nearly circular orbit with $P = 78.9 \pm 0.3$ days and velocity semi-amplitude $K_B = 52.8 \pm 0.7$ km s⁻¹, yielding a mass function (which represents the absolute minimum companion mass compatible with the orbit of the star) of $f(m) = 1.2 M_\odot$. If the B star’s mass is taken to be $M_B \approx 8.2 M_\odot$, as [Liu et al. \(2019\)](#) estimated from a model fit to the stellar spectrum, this corresponds to a lower limit on the companion mass of $M_2 \gtrsim 6.3 M_\odot$.

However, [Liu et al. \(2019\)](#) argue for a companion mass much higher than this, $M_2 = 68_{-13}^{+11} M_\odot$, which implies a low inclination ($i \approx 16$ deg). The spectrum of LB-1 contains both the expected absorption lines of a B star, and a strong H α emission line, which the authors interpret as originating from an accretion disk around a BH. They find the H α line to be RV variable, with velocity semi-amplitude $K_{H\alpha} = 6.4 \pm 0.8$ km s⁻¹, and interpret the ratio of the velocity amplitudes of the absorption and emission lines, $K_B/K_{H\alpha} \approx 8.2$, as representing the mass ratio, M_2/M_B . Taking $M_B \approx 8.2 M_\odot$, this yields $M_2 \approx 68 M_\odot$.

The proposed BH mass is more than double that of the most massive BHs in dynamically confirmed X-ray binaries (e.g. [McClintock & Remillard 2006](#)) and is also more massive than the most massive pre-merger BHs thus far discovered via gravitational waves ([Abbott et al. 2019](#)). The existence of a $68 M_\odot$ BH formed by stellar collapse is unexpected on theoretical grounds, because would-be progenitors of BHs with masses between about 50 and $130 M_\odot$ are predicted to lose mass and then explode via pulsational pair-instability supernovae (e.g. [Woosley 2017](#)), leaving behind either lower-mass BHs or no remnants at all. The novel and unexpected nature of the object has already spurred several theoretical investigations into possible formation channels for such a binary (e.g. [Belczynski et al. 2019](#); [Shen et al. 2019](#); [Groh et al. 2019](#)).

In this Letter, we reconsider the evidence for the mass of the unseen companion in LB-1. We argue that the companion is probably a stellar-mass BH, but that there is no evidence for an unusually high mass.

8.3 Basic issue

The observed spectrum of LB-1 contains absorption lines, presumably formed in the photosphere of a B star, and emission lines (most notably, H α), with some other origin. B stars have strong H α absorption lines with broad wings, so the observed flux at the H α line is a sum of emission and absorption (i.e., a reduction in the stellar continuum). The stellar absorption lines are known to be RV-variable. If one wishes to determine the velocity of the H α *emission* alone, it is necessary to first account for the wavelength-dependent reduction in total observed flux due to the absorption line. This is illustrated schematically in [Figure 8.1](#), where we show the effects of an RV-variable absorption line on the apparent velocity of an emission line whose true velocity is constant.

The emission line is modeled with a Voigt profile centered on 0 with Gaussian $\sigma = 90$ km s⁻¹, Lorentzian $\gamma = 50$ km s⁻¹, and amplitude 2.8. As we discuss in [Section 8.4](#), we




Figure 8.1: Schematic illustration of the effects of an RV-variable absorption line on the apparent velocity of a stationary emission line. The velocity of the emission line is fixed at $\text{RV}_{\text{H}\alpha} = 0 \text{ km s}^{-1}$. We Doppler shift the spectrum of a B star (dotted lines) by $\pm 53 \text{ km s}^{-1}$, representing the luminous star’s orbital motion. This causes the total spectrum (emission plus absorption; solid lines) to shift in the opposite direction. We then fit for the mean velocity of the $\text{H}\alpha$ emission at each epoch using the wings of the total emission profile, without accounting for contamination from the absorption line. This causes the emission line to appear to shift by $\pm 4.4 \text{ km s}^{-1}$, exactly in anti-phase with the luminous star. This shift is not real; the emission line’s velocity is fixed at 0 by construction.




Figure 8.2: Left: comparison of Keck/HIRES spectra of LB-1 obtained at two different epochs, centered on the $H\alpha$ line. The velocity of the B star varies by 71 km s^{-1} between these epochs. The wings of the emission line profiles also suggest a velocity offset between them, and fitting yields an offset of 10 km s^{-1} . Right: the same two spectra after the template spectrum of a B star at the appropriate velocities has been subtracted. There is now no evidence of a velocity offset between the two emission lines: the apparent offset in the left panel was due to the RV variability of the stellar absorption line.

find this functional form to provide a good fit to the Keck/HIRES $H\alpha$ emission profile of LB-1, particularly its wings. We model the spectrum of the B star with a TLUSTY model (Hubeny & Lanz 1995) with $T_{\text{eff}} = 18,000 \text{ K}$, $\log g = 3.5$, and $Z = Z_{\odot}$, consistent with the fit to the spectrum of the B star obtained by Liu et al. (2019)¹. We broaden the spectrum to $v \sin i = 10 \text{ km s}^{-1}$, as found for the observed B star, using the rotational broadening profile from Gray (1992a). We shift the stellar spectrum by $\pm 53 \text{ km s}^{-1}$, representing changes in the velocity of the B star throughout its orbit, and then add the emission line (always centered on $\text{RV}_{H\alpha} = 0 \text{ km s}^{-1}$) to the shifted stellar spectrum at each epoch. We degrade the total spectrum to the resolution of the Keck/HIRES data analyzed in the next section ($R \approx 60,000$), assuming a Gaussian line spread function. The details of our treatment of the spectrum do not sensitively affect our results, because the $H\alpha$ absorption line is broad and basically featureless in the region used to measure its velocity.

Figure 8.1 shows that shifts of the absorption line cause the total emission profile to shift in anti-phase with the star. When we fit for the line’s mean velocity (using the profile wings, as described in the next section), we infer an erroneous shift of $\pm 4.4 \text{ km s}^{-1}$, *even though the line is stationary by construction*. When the stellar spectrum is subtracted from the total emission, we of course recover the true $\text{RV}_{H\alpha} = 0 \text{ km s}^{-1}$ at all epochs. One might hope that

¹Several works (e.g. Abdul-Masih et al. 2019; Simón-Díaz et al. 2019; Irrgang et al. 2019) have reported a cooler temperature, $T_{\text{eff}} \approx 13,000 \text{ K}$. None of our results change significantly if we adopt this value.

masking the line center during fitting would protect against contamination from absorption. It does not, because although the absorption line core is narrow and excluded from the region where RVs are fit, the Balmer lines in the B star have broad wings, such that the magnitude of emission and absorption in LB-1 are comparable at $|\text{RV}| \gtrsim 200 \text{ km s}^{-1}$.

8.4 Contamination from absorption in LB-1

We now investigate the effect of contamination from $\text{H}\alpha$ absorption on the spectrum of LB-1. We analyze the same 7 epochs of Keck/HIRES data analyzed by Liu et al. (2019, their Extended Data Table 1). The data were obtained and reduced using the standard California Planet Search setup (Howard et al. 2010) and are available through the Keck Archive. We focus on the Keck/HIRES data because it has the highest spectral resolution, allowing the most precise measurement of the $\text{H}\alpha$ line’s velocity, but we expect measurements based on spectra from other instruments (e.g. LAMOST and GTC) to be affected in the same way.

We barycenter-corrected the spectra and normalized them by dividing by the mean flux at $6580 < \lambda/\text{AA} < 6600$; this simple procedure works well because the spectra are featureless in this region and we only study a narrow wavelength range centered on $\text{H}\alpha$.

Figure 8.2 compares the $\text{H}\alpha$ line profiles of LB-1 in the two epochs between which the velocity of the B star varies most. In the left panel, we simply overplot the normalized spectra from the two epochs. Here, there appears to be a visible shift between the two line profiles, and fitting for the emission line velocities (see Section 8.4.1) yields a velocity shift of 10 km s^{-1} between them. However, the right panel of Figure 8.2 shows that once the TLUSTY spectrum of the B star is subtracted, the line profile wings in the two epochs become almost perfectly aligned, with no evidence for a velocity shift between them. We find very similar results when comparing all pairs of single-epoch spectra.

8.4.1 Fitting RVs

We now turn to measuring the radial velocity of the $\text{H}\alpha$ emission line at all epochs. As shown in Figure 8.2, the emission near line center (within $\pm 150 \text{ km s}^{-1}$ of the mean velocity) varies substantially between epochs, but the shape of the emission line wings is relatively stable. Following Liu et al. (2019), we measure the line velocity using the region of the spectrum with $|\text{RV}| < 500 \text{ km s}^{-1}$, while masking out the line center, which is defined as the region where the emission line flux exceeds 1/3 of its maximum value (shaded regions in Figure 8.3). We measure the velocity of the line at each epoch by fitting a Voigt profile, which we find can always provide a good fit to the line profile within these windows.

Each fit has 4 free parameters: a central velocity $\text{RV}_{\text{H}\alpha}$, an amplitude A (which we define as the profile’s value at line center), a Gaussian dispersion σ , and a Lorentzian HWHM γ . Figure 8.3 shows the results of these fits. We also carry out similar fits *without* first subtracting the spectrum of the B star. The formal fitting uncertainty in $\text{RV}_{\text{H}\alpha}$ is typically $(0.1 - 0.2) \text{ km s}^{-1}$. We inflate this uncertainty by adding to it a 1 km s^{-1} systematic error in

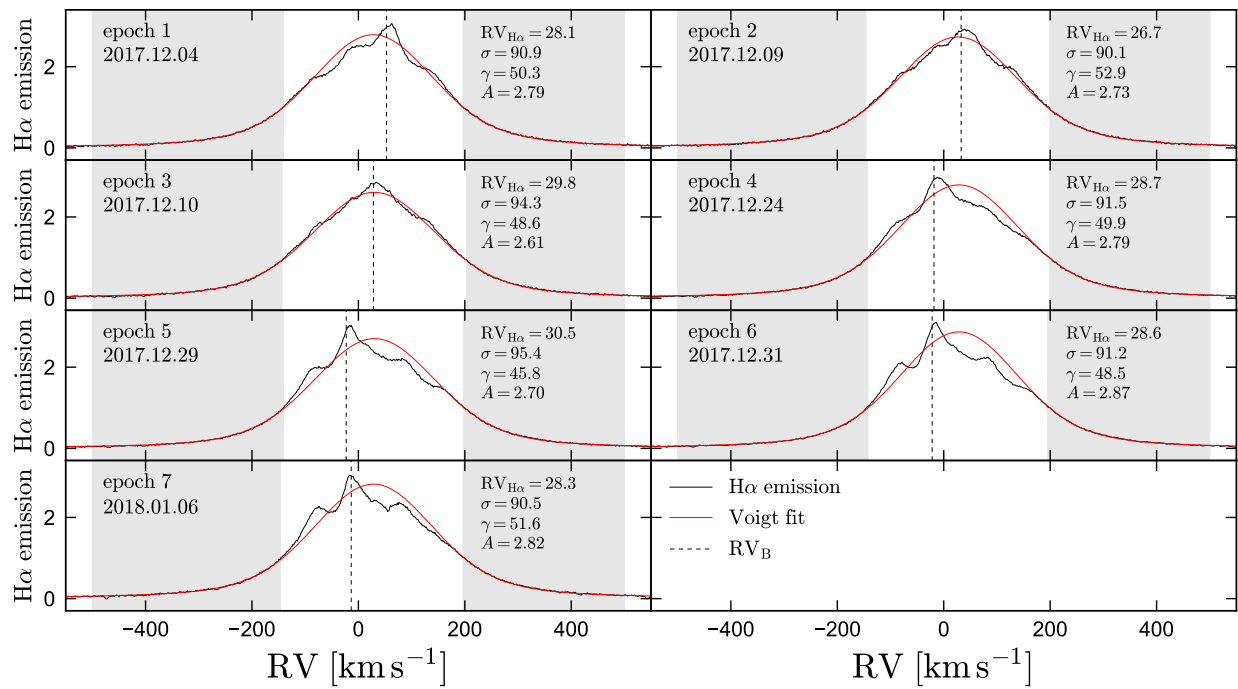


Figure 8.3: Voigt profile fits to all 7 epochs of Keck/HIRES data, after subtraction of the stellar spectrum. Gray shaded regions show the data used in fitting; the line centers are masked. Although the shape of the line center varies between epochs, the wings are relatively stable and can always be reasonably well-fit by a Voigt profile. Best-fit parameters are listed for each epoch. Dashed vertical lines show the velocity of the B star as measured from absorption lines. A narrow component of the emission appears to track the B star.

quadrature to account for uncertainty in the wavelength solution when it is not registered to lines of known wavelength (e.g. Griest et al. 2010). Our fitting procedure is not identical to the one used by Liu et al. (2019), but we find it to be robust and show below that it yields consistent results.

In Figure 8.4, we plot the results of this RV fitting as a function of the binary’s orbital phase. Red points show RV measurements when contamination from the B star’s absorption lines is not corrected for. As expected, these are in agreement with the model from Liu et al. (2019), which has $K_{\text{H}\alpha} = 6.4 \text{ km s}^{-1}$, and they are inconsistent with a constant velocity for the H α line. Black points show RV measurements after the B star template spectrum has been subtracted. These show no evidence for variation with orbital phase and appear consistent with a constant velocity for the H α line wings.

The data are strongly inconsistent with $K_{\text{H}\alpha} = 6.4 \text{ km s}^{-1}$ and in fact rule out any $K_{\text{H}\alpha} > 1.3 \text{ km s}^{-1}$ at the 2-sigma level. The best-fit sinusoidal velocity amplitude is formally $K_{\text{H}\alpha} = 0.25_{-0.3}^{+0.7} \text{ km s}^{-1}$, but a constant velocity provides essentially as good a fit as the best-fit sinusoidal model ($\Delta\chi^2 = 0.1$). Given the sinusoidal model’s higher complexity, there is no evidence that the H α line is RV-variable.

The dashed vertical lines in Figure 8.3 show the velocity of the B star at each epoch. Although the wings of the emission line are stationary, there is a narrow component of the emission (FWHM $\approx 50 \text{ km s}^{-1}$) that shifts between epochs and appears to trace the B star. Moreover, the emission flux in the inner wings (outside the narrow peak) is correlated with the velocity of the star: in epochs 4-8, when the star is blueshifted relative to the barycenter, there is excess emission in the left wing relative to the Voigt fit, and a deficit of emission in the right wing. The situation is reversed in epoch 1, when the star is redshifted relative to the barycenter.

8.5 Discussion

We have shown that the H α line in LB-1, the binary reported to contain a $70 M_{\odot}$ BH, does not actually display evidence of periodic RV variability, at least not at the $K_{\text{H}\alpha} > 1.3 \text{ km s}^{-1}$ level. This undermines the interpretation of the previously reported $K_{\text{H}\alpha}$ in terms of the system mass ratio, and thus the reported $70 M_{\odot}$ mass.

8.5.1 Nature of the unseen companion

One conceivable interpretation is that the companion is a BH with even higher mass than reported: assuming $M_B = 8.2$, our upper limit on $K_{\text{H}\alpha}$ would imply $M_2 \gtrsim 330 M_{\odot}$. We regard this scenario as exceedingly unlikely both because of the astrophysical challenge of producing such a system and because it would imply an improbably low inclination ($i \lesssim 9 \text{ deg}$, representing $\sim 1\%$ of randomly oriented orbits). We thus conclude that the H α line (at least its wings) is not moving with either component of the binary.

Although there is little evidence that LB-1 contains an unusually massive BH, the unseen

companion may still be a stellar-mass BH. If the mass of the B star is indeed $M_B \approx 8M_\odot$ as found by Liu et al. (2019), this would imply $6 \lesssim M_2/M_\odot \lesssim 20$ for inclinations $30 < i/\text{deg} < 90$ (representing 87% of randomly oriented orbits). Normal-star companions in this mass range are expected to contribute to the spectrum at a detectable level, so the lack of absorption lines from a luminous secondary would make a “normal” stellar-mass BH the most likely companion.

Two additional observations point toward the companion being a stellar remnant. First, the orbit is nearly circular. This is rare for ordinary main-sequence binaries with $P \gtrsim 10$ days, which have a wide range of eccentricities (Moe & Di Stefano 2017a; El-Badry et al. 2018g), but it is expected for binaries containing evolved stars, because even binaries with relatively long periods ($P \lesssim 200$ days) become tidally circularized when one component ascends the giant branch (Verbunt & Phinney 1995; Price-Whelan & Goodman 2018). If LB-1 was tidally circularized when the companion was a giant, then the fact that the orbit is still circular may set stringent limits on any mass loss or kicks that occurred during the death of the companion.

Second, the rotation velocity of the luminous component ($v \sin i \approx 10 \text{ km s}^{-1}$) is exceptionally low for a B star. This could be explained if the star was tidally synchronized when the companion was a giant, such that the rotation period is equal to the orbital period. Synchronization is expected to occur on a timescale similar to or shorter than circularization. The expected rotation velocity in this case is $v_{\text{rot}} = 2\pi R_B/P$, which ranges from 4.5 to 7 km s^{-1} for $7 < R_B/R_\odot < 11$, similar to the observed rotation velocity. If the B star is tidally synchronized, then a precise measurement of $v \sin i$ and the stellar radius could constrain the orbital inclination.

As discussed by Liu et al. (2019), their spectroscopically-inferred $M_B \approx 8M_\odot$ is inconsistent with the *Gaia* DR2 parallax, which implies a lower mass, $M_B \approx 5M_\odot$. The revised atmospheric parameters estimated by Simón-Díaz et al. (2019) and Abdul-Masih et al. (2019) also support such a scenario, which would yield a lower limit of $M_2 \gtrsim 5M_\odot$, still making a BH the most likely companion.

Eldridge et al. (2019) and Irrgang et al. (2019) proposed that the luminous star could be a $\sim 1M_\odot$ pre-subdwarf in a short-lived ($\sim 10^4$ yr) evolutionary phase. In this scenario, the companion could also have a significantly lower mass: for $M_B < 1.2M_\odot$, the minimum mass of the companion is less than $2.5M_\odot$, such that it could be a neutron star. The primary challenge to this scenario is its short lifetime: the probability is low that the LAMOST RV monitoring campaign, which observed ~ 3000 targets, would happen to observe such a rare object.

8.5.2 Source of the $\text{H}\alpha$ emission

The shape of the $\text{H}\alpha$ profile near line center varies with orbital phase, with a narrow component that tracks the velocity of the B star (Figure 8.3). This may originate at the star itself, perhaps from accretion or from interaction of the stellar wind with the circumbinary environment. It is possible that a component of the emission tracks the BH, but this is not




Figure 8.4: Red points show RV measurement of LB-1’s H α emission line obtained from Keck-/HIRES spectra without accounting for the stellar absorption line (e.g. left panel of Figure 8.4). These show clear evidence of RV variability and are consistent with the orbital solution for the H α line derived in Liu et al. (2019). Black points show RV measurements once contamination from stellar absorption lines is accounted for (e.g. right panel of Figure 8.4). Here there is no evidence of RV variability, and the model with $K_{\text{H}\alpha} = 6.4 \text{ km s}^{-1}$ is manifestly ruled out.

easily distinguishable with the presently available spectra.

As for the wings, the fact that their mean velocity does not vary with phase suggests that this emission must not originate from either component of the binary, but rather from material in the barycentric rest frame. Because the present-day orbit is smaller than the radius of plausible progenitors to the companion when it was a giant, the system likely went through a period of common envelope evolution. Fallback material from the companion's envelope could form a circumbinary disk in this scenario (Muno & Mauerhan 2006; Kashi & Soker 2011). If the luminous star is a recently-stripped pre-subdwarf (e.g. Irrgang et al. 2019), fallback material from its envelope could also contribute. Another possibility is that the low-velocity component of winds from the B star are captured by the binary (e.g. Taam & Spruit 2001).

For a $5M_{\odot}$ BH and a $5M_{\odot}$ star, $a \approx 0.75$ AU, and the circular velocity at the inner rim of a circumbinary disk ($r \approx 2a$) would be $v_c \approx 75 \text{ km s}^{-1}$. For an optically thin Keplerian disk, one would thus expect a double-peaked emission line with peaks separated by $\sim 150 \text{ km s}^{-1}$. The $\text{H}\alpha$ line is somewhat wider than this, with $\text{FWHM} \approx 260 \text{ km s}^{-1}$ and shallow wings that extend at least $\pm 500 \text{ km s}^{-1}$ from line center. The wings likely do not trace disk kinematics, but are instead broadened by a combination of non-coherent scattering (i.e., repeated absorption and re-emission of line photons in an optically thick medium; see Hummel & Dachs 1992) and electron scattering (in the outermost wings; see Poeckert & Marlborough 1979). The effects of these processes in broadening $\text{H}\alpha$ are well-studied in the disks around classical Be stars (e.g. Hanuschik et al. 1996). Spiral density waves, interactions with the B star's winds, and fluctuations of the disk's rim temperature due to the B star's orbital motion can all complicate the emission line profile. In the future, a more dense sampling of spectra in orbital phase may enable reconstruction of the emission region and more detailed modeling of the circumbinary environment.

There are several similarities between LB-1 and the recently-identified binary AS 386 (Khokhlov et al. 2018), which contains a $7M_{\odot}$ B-star primary and a detached $\gtrsim 7M_{\odot}$ unseen companion suspected to be a BH. That system also has a strong $\text{H}\alpha$ emission line whose shape changes between epochs. Although the mean velocity of the system's $\text{H}\alpha$ emission does vary with phase, its velocity amplitude is lower than that of the star, and Doppler tomography suggested that it originates from a shock where the wind from the B star interacts with a circumbinary disk. The orbit of the B star in AS 386 is also circular, despite a long period of 131 days. Like LB-1, and in contrast to most classical Be stars, AS 386 is slowly rotating. Also in contrast to classical Be stars (Dachs et al. 1990), both LB-1 and AS 386 have stronger Balmer decrements than predicted for low-density nebular gas, case-B recombination, and a temperature near 10^4 K. The spectral energy distributions of AS 386 and LB-1 both reveal infrared excess; Khokhlov et al. (2018) find that in AS 386, this stems primarily from dust emission in a circumbinary disk.

AS 386 exhibits periodic variability in the infrared. Khokhlov et al. (2018) attribute this to dust temperature variations in the Earth-facing inner rim of the circumbinary disk as the distance from the B star to the rim changes. We do not find evidence of such variability in the WISE light curves of LB-1, although the light curves are sparse. This could indicate that

the inclination is lower than in AS 386, or that the IR excess (which is weaker than that in AS 386) is emitted at larger radii. We also note that unlike LB-1, the spectrum of AS 386 contains several other emission lines, including forbidden lines emitted in the circumbinary disk. Further theoretical modeling of LB-1 will be required to explain these differences.

Although it likely does not contain a $70 M_{\odot}$ BH, LB-1 is an intriguing system which adds to the short but growing list of Milky Way stellar mass BH candidates that are not X-ray bright (e.g. [Casares et al. 2014](#); [Khokhlov et al. 2018](#); [Giesers et al. 2018](#); [Thompson et al. 2019](#)). If the unseen companion is confirmed to be a BH, LB-1's circular orbit will provide tantalizing evidence that some massive stars end their lives without an explosion (e.g. [Adams et al. 2017](#)). We are optimistic that further study of the system will shine additional light on its evolutionary history.

Acknowledgements

We thank the anonymous referee for a constructive report, and JJ Eldridge, Andrew Howard, Howard Isaacson, Andreas Irrgang, Chris Kochanek, Jifeng Liu, Ben Margalit, Todd Thompson, Dan Weisz, Yanqin Wu, and Wei Zhu for helpful discussions. K.E acknowledges support from an NSF graduate research fellowship.

Chapter 9

A stripped-companion origin for Be stars: clues from the putative black holes HR 6819 and LB-1

An earlier version of this article was previously published in El-Badry, K. and Quataert, E., 2020, MNRAS, 502, 3436.

9.1 Abstract

HR 6819 is a bright ($V = 5.36$), blue star recently proposed to be a triple containing a detached black hole (BH). We show that the system is a binary and does not contain a BH. Using spectral decomposition, we disentangle the observed composite spectra into two components: a rapidly rotating Be star and a slowly rotating B star with low surface gravity ($\log g \approx 2.75$). Both stars show periodic radial velocity (RV) variability, but the RV semi-amplitude of the B star's orbit is $K_B = (62.7 \pm 1) \text{ km s}^{-1}$, while that of the Be star is only $K_{\text{Be}} = (4.5 \pm 2) \text{ km s}^{-1}$. This implies that the B star is less massive by at least a factor of 10. The surface abundances of the B star bear imprints of CNO burning. We argue that the B star is a bloated, recently stripped helium star with mass $\approx 0.5 M_\odot$ that is currently contracting to become a hot subdwarf. The orbital motion of the Be star obviates the need for a BH to explain the B star's motion. A stripped-star model reproduces the observed luminosity of the system, while a normal star with the B star's temperature and gravity would be more than 10 times too luminous. HR 6819 and the binary LB-1 probably formed through similar channels. We use MESA models to investigate their evolutionary history, finding that they likely formed from intermediate-mass ($3 - 7 M_\odot$) primaries stripped by slightly lower-mass secondaries and are progenitors to Be + sdOB binaries such as ϕ Persei. The lifetime of their current evolutionary phase is on average 2×10^5 years, of order half a percent of the total lifetime of the Be phase. This implies that many Be stars have hot subdwarf and white dwarf companions, and that a substantial fraction (20 – 100%) of field

Be stars form through accretion of material from a binary companion.

9.2 Introduction

Large-scale multi-epoch radial velocity (RV) surveys have identified a number of binaries in recent years that are proposed to contain stellar mass black holes (BHs; e.g. Casares et al. 2014; Khokhlov et al. 2018; Giesers et al. 2018, 2019; Thompson et al. 2019). Confirmation of these BH candidates is challenging precisely because BHs are expected to be rare. That is, all plausible alternate explanations for observed data, even those which are rare, must be ruled out before a candidate BH can be considered reliable.

Recently, Rivinius et al. (2020) identified the object HR 6819 as a candidate host of a stellar-mass BH. Phase-resolved optical spectra of the object revealed two luminous components: a B star with relatively narrow, RV-variable absorption lines, which were observed to follow a nearly circular orbit with $P = 40.3$ days and velocity semi-amplitude $K_B \approx 61 \text{ km s}^{-1}$, and a classical Be star whose broad emission and absorption lines appeared to be stationary. No component was found to orbit in anti-phase with the B star. Assuming the mass of the B star to be at least $5 M_\odot$, as expected for a normal star of its spectral type, Rivinius et al. (2020) argued that any stellar companion massive enough to explain the B star’s orbit would also contribute to the spectrum at a detectable level. They thus concluded that the companion is a BH, with an estimated minimum mass of $4.2 M_\odot$. In this hierarchical triple scenario, the Be star must be at least a few AU from the B star and BH for the system to be dynamically stable, and its status as a Be star would likely be unrelated to the B star or BH.

HR 6819 is in many ways similar to the binary LB-1 (Liu et al. 2019), which was also proposed to contain a stellar-mass BH. Like HR 6819, LB-1 contains an RV-variable B star and apparently stationary emission lines. The emission lines were initially proposed to originate in an accretion disk, either around the BH (Liu et al. 2019) or around the binary (El-Badry & Quataert 2020; Abdul-Masih et al. 2020; Irrgang et al. 2020). Rivinius et al. (2020) proposed that LB-1 and HR 6819 are both hierarchical triples with a B star and a BH companion in the inner binary, and a distant Be star – the source of the emission lines – orbiting both components.

Shenar et al. (2020), however, recently used spectral disentangling to fit the multi-epoch spectra of LB-1 as a sum of two luminous components. They also found evidence for a Be star, including both emission and rotationally-broadened absorption lines, but found its spectrum to shift in anti-phase with the B star. This is not expected in the hierarchical triple scenario, so they argued that LB-1 is a binary containing two luminous stars. It was also recently noted by Liu et al. (2020) that the emission line shape in LB-1 varies coherently with the B star’s phase. This variation, which is consistent with expectations for an irradiated or tidally-perturbed disk, is expected if the system is a binary, but not if it is a hierarchical triple.

The RV variability amplitude found by Shenar et al. (2020) for the Be star in LB-1

suggests it is 5 times more massive than the B star. If the luminous binary scenario for LB-1 is correct, the B star thus would have a mass of $M_B \approx 1.5 M_\odot$, a factor of 4 lower than expected for a normal star of its spectral type. Shenar et al. (2020) propose that most of the B star’s envelope was recently stripped by its companion. In this case, the star is currently contracting and will likely soon become a core helium burning sdOB star (see also Irrgang et al. 2020).

In this paper, we use spectral disentangling to fit the multi-epoch spectra of HR 6819. We find that, like in LB-1, the Be star orbits in anti-phase with the B star, which is much less massive than expected for a normal star of its spectral type. The remainder of this paper is organized as follows. We describe the spectra in Section 9.3.1 and the disentangling method in Section 9.3.2. We constrain the atmospheric parameters of both components in Section 9.3.3. Section 9.3.4 presents the RV variability of both components, and Section 9.3.5 compares the system’s luminosity to models. The emission from the Be star’s disk is examined in Section 9.3.6, and the abundances of the B star in Section 9.3.7. We present possible evolutionary models for the system in Section 9.3.8. Finally, we discuss HR 6819 and LB-1 in the context of the broader Be star population in Section 9.4.

The appendices provide supporting information. Appendix 9.5 describes our estimate of both components’ masses, radii, and luminosities. Spectral disentangling is detailed in Appendix 9.6. Appendix 9.7 investigates the pulsation-driven variability of the B star. Spectroscopic constraints on the B star’s rotation velocity are presented in Appendix 9.8. Its surface helium abundance is investigated in Appendix 9.9.

9.3 Methods

9.3.1 Data

We analyze 51 optical spectra of HR 6819 with spectral resolution $R \approx 48,000$ that were taken in 2004 with the FEROS echelle spectrograph (Kaufer et al. 1999) on the ESO/MPG 2.2m telescope at La Silla Observatory. The data span 134 days (3.3 orbital periods) from MDJ 53138 to 53273 and are publicly available through the ESO archive; they are described in more detail by Rivinius et al. (2020). The spectra were reduced with the ESO-MIDAS pipeline, which performs bias-subtraction, flat fielding, and wavelength calibration, applies a heliocentric correction, and combines the spectra from individual orders. The pipeline failed to reliably measure the signal-to-noise ratio (SNR), so we estimated it empirically from the pixel-to-pixel scatter in regions without strong absorption or emission lines. The typical SNR per pixel is 300 at $\lambda = 6,000 \text{ \AA}$.¹ We verified the stability of the wavelength solution by checking that the ISM sodium absorption line at 5890 \AA is found at the same heliocenter-corrected wavelength at all epochs.

We also analyzed an additional 12 FEROS spectra of HR 6819 that were obtained in 1999 (dataset “A” in Rivinius et al. 2020). We found that the Be star’s emission line profiles varied

¹All wavelengths are quoted at their values in air.

substantially between the 1999 and 2004 datasets, with obvious changes in the shape, width, and amplitude of Balmer, Fe II, and O I emission lines. Such variation is common in Be stars due to changes in the structure of the disk (e.g. [Dachs et al. 1981](#); [Okazaki 1991](#)), but it complicates the disentangling of the composite spectra. We therefore focus our analysis on the 2004 dataset, within which the emission line profiles are relatively stable.

The spectra are not flux calibrated. Continuum normalization was performed by fitting a cubic spline to wavelength pixels without significant absorption or emission. The initial set of candidate continuum pixel was selected by taking all wavelengths from a TLUSTY model spectrum ([Hubeny & Lanz 1995](#)) with $T_{\text{eff}} = 18 \text{ kK}$ and $\log g = 3.5$ for which all pixels within $\pm 60 \text{ km s}^{-1}$ are within 0.5% of the theoretical continuum. We then refined the selection of continuum pixels by inspecting the individual spectra, removing pixels affected by telluric absorption and other features not present in the model spectrum. We tested application of this continuum normalization procedure on mock spectra with a range of $\log g$ to verify that it does not significantly overfit the continuum in the wings of broad lines. This is important, because improper continuum normalization can change the shape of the wings of broad lines, leading to biases in the inferred $\log g$. We focus our analysis on the part of the spectrum with $\lambda > 3900 \text{ \AA}$ because the narrow spacing of the Balmer lines at shorter wavelengths makes it challenging to estimate the continuum there.

9.3.2 Spectral disentangling

We used the code CRES ([Ilijic 2004](#)), which implements the wavelength-space spectral disentangling algorithm of [Simon & Sturm \(1994\)](#), to separate the spectra of the two components. The algorithm operates under the ansatz that all the observed composite spectra are produced by summing together the time-invariant spectra of two components, with known Doppler shifts between epochs and a known continuum luminosity ratio. The rest-frame spectra of the two components are modeled as arbitrary vectors, which under these assumptions can be solved for using singular value decomposition. No model spectra (or indeed, any physics beyond the Doppler shift) are use in this calculation.

The velocity of the B star in HR 6819 is reasonably well known at each epoch through the orbital solution from [Rivinius et al. \(2020\)](#), but the velocity of the Be star is not. Following [Shenar et al. \(2020\)](#), we therefore proceed under the ansatz that the Be star and B star are orbiting each other, and step through a grid of velocity semi-amplitudes for the Be star, K_{Be} , calculating the best-fit disentangled spectra for each value. The choice of $K_{\text{Be}} = 0 \text{ km s}^{-1}$ corresponds to a stationary Be star and would be expected in a triple scenario. We choose the optimal K_{Be} and spectral decomposition as the value that minimizes the total χ^2 over all spectra; we find $K_{\text{Be}} = (4.5 \pm 2) \text{ km s}^{-1}$. The determination of this value is described in detail in [Appendix 9.6](#).

When determining K_{Be} , we use the spectra at their native resolution and independently analyze narrow wavelength ranges centered on absorption and emission lines with clear contributions from both components. We then adopt $K_{\text{Be}} = 4.5 \text{ km s}^{-1}$ and use this value to disentangle the full spectrum. At this stage, we re-bin the spectra to a 0.2 \AA wavelength

resolution to reduce the computational cost of disentangling. This does not significantly decrease the information content of the spectra, because higher-frequency spectral features are smeared out by rotation and/or macroturbulence. The spectral disentangling and determination of K_{Be} are described in detail in Appendix 9.6.

Spectral decomposition only determines the spectra of the two components up to a multiplicative constant, which represents the continuum luminosity ratio of the two stars. Unless there are e.g. eclipses or well-characterized pulsations, the luminosity ratio must be determined based on external information, because the effects on the combined spectrum of increasing the continuum flux contribution of one component are identical to the effects of weakening all its spectral lines. We begin by assuming both components contribute 50% of the light, as estimated by Rivinius et al. (2020). Once a preliminary decomposition is found with this ratio, we adjust the luminosity ratio to match the depths of the B star’s Balmer lines to model spectra, while accounting for rotation, macroturbulence, instrumental broadening, and rebinning. The temperature of the model spectra is constrained primarily by ionization balance of metal lines (see Section 9.3.3; this temperature estimate is not sensitive to the luminosity ratio). We found the optimal flux ratio to be $f_{\text{B}}/f_{\text{tot}} = 0.47 \pm 0.05$ at $\lambda = 4000 \text{ \AA}$ (i.e., with the Be star the slightly brighter component). The uncertainty in the flux ratio accounts for uncertainty in T_{eff} and $\log g$, as all atmospheric parameters are fit simultaneously. Because the Be star is slightly hotter, its relative contribution to the total flux is expected to decrease toward redder wavelengths. For our fiducial temperature estimates of $T_{\text{eff,B}} = 16 \text{ kK}$ and $T_{\text{eff,Be}} = 18 \text{ kK}$ (Section 9.3.3), the expected flux ratio increases to 0.49/0.51 at $\lambda = 7000 \text{ \AA}$, still with the Be star the brighter component. This weak wavelength-dependence of the flux ratio is accounted for in our spectral disentangling. We normalize the component spectra following the method described by Ilijic et al. (2004).

Figure 9.1 shows a 600 \AA wide section of the disentangled spectra. The overplotted TLUSTY model spectra are taken from the BSTAR06 grid (Lanz & Hubeny 2007) and have the T_{eff} , $\log g$, and $v \sin i$ that we find to best fit the disentangled spectra. The chemical abundances of the models are derived in Section 9.3.7. The Be star spectrum primarily shows strong absorption lines from neutral H and He due to its large $v \sin i$, as well as excess due to double-peaked emission lines within the Balmer line cores and in Fe II lines (Section 9.3.6). The B star spectrum has many narrow metal lines (see Section 9.3.7).

Fitting the disentangled metal line profiles of both components with the rotational profile from Gray (1992a), we find $v \sin i = (180 \pm 20) \text{ km s}^{-1}$ for the Be star. For our best-fit Be star mass, radius, and inclination (Sections 9.3.4 and 9.3.5), this corresponds to near-critical rotation. For the B star, we derive an upper limit of $v \sin i < 20 \text{ km s}^{-1}$. We are not able to measure the actual value of $v \sin i$ from the B star’s line profiles because non-rotational broadening due to macroturbulence and/or non-radial pulsations dominates over rotation (see Appendix 9.8). If the B star is tidally synchronized, as might be expected following a period of mass transfer, its $v \sin i$ would be less than 5 km s^{-1} . The Be star is not expected to be tidally synchronized, as accretion of mass and angular momentum is expected to spin it up.

Overall, the consistency between the disentangled spectra and theoretical models, as




Figure 9.1: Disentangled spectra of the B star (top) and Be star (bottom) in HR 6819. Black lines show the reconstructed spectra. Red lines show TLUSTY model spectra, which are not used in the spectral disentangling, for comparison. The B star spectrum contains many narrow lines and is well-fit by a model with $T_{\text{eff}} = 16 \text{ kK}$ and $\log g = 2.75$ (see also Figure 9.2). The Be star spectrum is strongly broadened by rotation. It is reasonably well described by a model with $T_{\text{eff}} = 18 \text{ kK}$ and $\log g = 3.75$, but with clear excess emission inside the Balmer line cores (see also Figure 9.6).

manifest in properties such as the relative depth and width of absorption lines, is good. We emphasize that the spectral disentangling procedure does not use model spectra in any way – the disentangled spectra are free to take any shape they like – so the agreement between the observed and model spectra is encouraging.

9.3.3 Temperature and gravity

We fit the disentangled spectra of both components using 1D-NLTE model spectra computed with TLUSTY (Hubeny & Lanz 1995, 2017). We use the BSTAR06 grid, which has $\Delta T_{\text{eff}} = 1 \text{ kK}$ and $\Delta \log g = 0.25 \text{ dex}$, (Lanz & Hubeny 2007) in an initial grid search and then use the radiative transfer code SYNSPEC (Hubeny & Lanz 2011) to generate spectra with finer grid spacing, as well as spectra with different abundance patterns than assumed in the BSTAR06 grid (Section 9.3.7). We constructed grids of spectra with 4 values of the microturbulent velocity: $v_{\text{mic}} = (2, 5, 10, 15) \text{ km}^{-1}$. We find $v_{\text{mic}} = 10 \text{ km s}^{-1}$ to provide the best fit and adopt this value in the rest of our analysis (see also Bodensteiner et al. 2020), but we account for a $\pm 5 \text{ km s}^{-1}$ uncertainty in v_{mic} when calculating uncertainty in other spectral parameters and abundances.

We perform an iterative fit for the temperature and gravity of the B star, first measuring T_{eff} based on equivalent width ratios of lines in different ionization states (Figure 9.3), then measuring $\log g$ based on the wings of broad lines, then refining the T_{eff} estimate taking the estimated $\log g$ into account, and so forth. Detailed abundances are measured in a final step, once a converged solution for T_{eff} and $\log g$ is found. To estimate systematics due to model uncertainties, we also compared to 1D-LTE model spectra generated with SYNTHE (Kurucz 1993) from ATLAS-12 model atmospheres (Kurucz 1970, 1979, 1992). We find consistent solutions for the two sets of model spectra within 1 kK in T_{eff} and 0.2 dex in $\log g$.

The model spectra used in determining the temperature and gravity of the B star assume a solar abundance pattern. We show in Section 9.3.7 that the B star is enhanced in helium. We have checked that this does not significantly change our inferred atmospheric parameters, and the level of helium enhancement we find has a minimal effect on the shapes of the Balmer lines.

Our final best-fit parameters for both stars are listed in Table 9.1. The reported uncertainties in atmospheric parameters include formal fitting uncertainties due to noise in the spectra, uncertainty in the flux ratio, systematic differences between fits with TLUSTY/SYNSPEC and Kurucz spectra, and covariances between atmospheric parameters. Systematic uncertainty dominates the error in T_{eff} , and because T_{eff} and $\log g$ are strongly covariant, also contributes significantly to the uncertainty in $\log g$.

A critical result of our analysis is that the surface gravity of the B star is lower than that of a main-sequence star, $\log g \approx 2.75$. The main constraint on $\log g$ comes from the wings of the Balmer lines and some He I lines, which become broader with increasing $\log g$. This is illustrated in Figure 9.2, which shows a zoom-in on the Balmer lines H ϵ , H δ , and H γ in the disentangled spectrum of the B star. We compare the best-fit model, which has $\log g = 2.75$, to one with $\log g = 3.75$, the approximate value assumed by Rivinius et al. 2020, and the

Table 9.1: Physical parameters and 1σ (middle 68%) uncertainties for both components of HR 6819. Constraints on stellar parameters are based on the measured T_{eff} and $\log g$ of the Be star and the dynamical mass ratio. Table 9.3 lists constraints that also take the distance into account.

Parameters of the B star		
Effective temperature	T_{eff} [kK]	16 ± 1
Surface gravity	$\log(g/(\text{cm s}^{-2}))$	2.75 ± 0.35
Projected rotation velocity	$v \sin i$ [km s^{-1}]	< 20
Macroturbulent velocity	v_{macro} [km s^{-1}]	30 ± 5
Microturbulent velocity	v_{mic} [km s^{-1}]	10 ± 5
Continuum flux ratio at 4000 Å	$f_{\text{B}}/f_{\text{tot}}(4000 \text{ Å})$	0.47 ± 0.05
Radius	$R [R_{\odot}]$	$4.7^{+2.9}_{-1.9}$
Bolometric luminosity	$\log(L/L_{\odot})$	$3.11^{+0.42}_{-0.46}$
Mass	$M [M_{\odot}]$	$0.47^{+0.28}_{-0.22}$
Parameters of the Be star		
Effective temperature	T_{eff} [kK]	18^{+2}_{-3}
Surface gravity	$\log(g/(\text{cm s}^{-2}))$	$3.75^{+0.5}_{-0.25}$
Projected rotation velocity	$v \sin i$ [km s^{-1}]	180 ± 20
Macroturbulent velocity	v_{macro} [km s^{-1}]	60 ± 20
Continuum flux ratio at 4000 Å	$f_{\text{Be}}/f_{\text{tot}}(4000 \text{ Å})$	0.53 ± 0.05
Mass	$M [M_{\odot}]$	$6.7^{+1.9}_{-1.5}$
Radius	$R [R_{\odot}]$	$4.7^{+2.7}_{-1.5}$
Bolometric luminosity	$\log(L/L_{\odot})$	$3.35^{+0.47}_{-0.44}$
Fraction of critical rotation	$v_{\text{rot}}/v_{\text{crit}}$	$0.80^{+0.18}_{-0.14}$
Parameters of the binary		
Orbital period	P [day]	40.3 ± 0.3
B star RV semi-amplitude	K_{B} [km s^{-1}]	62.7 ± 1.0
Be star RV semi-amplitude	K_{Be} [km s^{-1}]	4.5 ± 2.0
Be star center-of-mass velocity	γ_{B} [km s^{-1}]	10.4 ± 0.6
Be star center-of-mass velocity	γ_{Be} [km s^{-1}]	10.5 ± 0.3
B star eccentricity	e_{B}	< 0.037
Be star eccentricity	e_{Be}	< 0.013
B star RV scatter	s_{B} [km s^{-1}]	4 ± 0.5
Be star RV scatter	s_{Be} [km s^{-1}]	< 1.0
Mass ratio	$q = M_{\text{B}}/M_{\text{Be}}$	0.071 ± 0.032
Orbital inclination	i [deg]	$32.1^{+3.0}_{-2.6}$
Separation	$a [R_{\odot}]$	96 ± 8
B star abundances relative to Solar		
Helium	[He/H]	0.55 ± 0.2
Carbon	[C/H]	-0.5 ± 0.2
Nitrogen	[N/H]	1.1 ± 0.3
Oxygen	[O/H]	0.3 ± 0.3
Neon	[Ne/H]	0.15 ± 0.2
Magnesium	[Mg/H]	0.2 ± 0.2
Aluminum	[Al/H]	0.4 ± 0.2
Silicon	[Si/H]	0.3 ± 0.2
Sulfur	[S/H]	-0.3 ± 0.2
Argon	[Ar/H]	0.4 ± 0.2
Iron	[Fe/H]	0.2 ± 0.2
Be star disk parameters		
Projected outer disk rotation velocity	$v_{\text{out}} \sin i$ [km s^{-1}]	105 ± 4
Ratio of inner to outer disk radii	$R_{\text{inner}}/R_{\text{outer}}$	0.24 ± 0.05
Fe II emissivity exponent	$\alpha = -d \ln j / d \ln r$	1.6 ± 0.3
Outer disk radius	$R_{\text{outer}} [R_{\odot}]$	33^{+5}_{-4}
Inner disk radius	$R_{\text{inner}} [R_{\odot}]$	$7.8^{+2.1}_{-1.8}$




Figure 9.2: Balmer lines as a surface gravity diagnostic. Black lines show the spectrum of the B star obtained through spectral disentangling. Red lines shows a TLUSTY model spectrum with $\log g = 2.75$, which reproduces the observed Balmer line profiles reasonably well. Cyan lines show model spectra with $\log g = 3.75$, as expected for a normal B star near the end of the main sequence. For this value of $\log g$, the wings of the Balmer and He I lines are much too wide to match the observed spectrum. The model T_{eff} values are chosen to match the observed ionization ratios at a given $\log g$ (Figure 9.3). For other choices of T_{eff} , the disagreement between the $\log g = 3.75$ model and the observed spectrum is even more severe.

best-fit T_{eff} for that $\log g$. The $\log g = 3.75$ model predicts much wider Balmer lines than are found in the observed spectrum. The poor fit of the higher $\log g$ model is also evident in the wings of the 4009 Å and 4388 Å He I lines. We note that while the superiority of the $\log g = 2.75$ model is clear in Figure 9.2, this is only the case because we have disentangled the spectra of the two components. In the composite spectra, the Balmer and He I lines appear much broader due to the broader absorption lines of the Be star.

The effective temperature of the B star is constrained from equivalent width ratios of lines of the same element in different ionization states. Higher ionization states become increasingly populated as temperature increases. The ratios of e.g. Si III to Si II equivalent widths are therefore sensitive to T_{eff} and relatively insensitive to the absolute Si abundance. Unlike absolute abundance measurements, equivalent width ratios are also insensitive to the assumed luminosity ratio of the two components, since light from the companion dilutes all lines by a similar factor.

Figure 9.3 compares the observed equivalent width ratios of Si III/II and Fe III/II lines to predictions from TLUSTY models for a range of T_{eff} and $\log g$. Our best-fit model with $\log g = 2.75$ is shown in black; colored lines show other values of $\log g$. All the line ratios we consider imply a temperature of $T_{\text{eff}} = (16 \pm 0.5)$ kK for $\log g = 2.75$, with values that are respectively higher and lower for higher or lower $\log g$.

Rapid rotation washes out most of the Be star’s lines, making it infeasible to measure T_{eff} from ionization state ratios. The constraint on T_{eff} comes primarily from the Balmer lines and the strength of Si II and He I lines. The star’s gravity can still be constrained from the shape of the Balmer lines, which are sufficiently broad that their wings are not strongly affected by rotation.

A natural concern when working with echelle spectra is that the profiles of broad lines such as the Balmer lines could be distorted by the blaze function and/or continuum normalization process, leading to a bias in $\log g$. To assess the reliability of our derived T_{eff} and $\log g$, we analyzed an archival FEROS spectrum of the slowly rotating standard B star HR 5285 that was presented and analyzed by Nieva & Przybilla (2007). The spectral resolution and SNR of this spectrum is similar to that of the HR 6819 data. We apply the same continuum normalization routine and procedure for estimating atmospheric parameters used for HR 6819 to this spectrum and obtain $T_{\text{eff}} = 21$ kK and $\log g = 4.25$, in good agreement with the literature values (Vrancken et al. 1996 found $T_{\text{eff}} = 22.6 \pm 0.9$ kK and $\log g = 4.2 \pm 0.1$, while Nieva & Przybilla 2012 found $T_{\text{eff}} = 20.8 \pm 0.3$ kK and $\log g = 4.22 \pm 0.05$). This allays any fear that our inferred $\log g$ values are significantly biased.

9.3.4 Radial velocities

We measure radial velocities for both components at each epoch by fitting the composite spectra as a sum of the two disentangled spectra obtained in Section 9.3.2, with the radial velocities of both components left free. We also leave the luminosity ratio free, because at least one component of the system is photometrically variable at the 5-10% level (Rivinius et al. 2020). We mask regions of the spectrum containing diffuse interstellar bands, telluric




Figure 9.3: Temperature of the B star from ionization equilibrium. Shaded regions show the observed ratios of the equivalent widths of Si III and Si II lines (left columns) and Fe III and Fe II lines (right columns). Lines show the ratios predicted for models of different $\log g$ as a function of T_{eff} . Given the best-fit $\log g \approx 2.75$ from the Balmer lines (Figure 9.2), the best-fit T_{eff} lies between 15500 and 16500 K for all line ratios.

absorption, and a few lines that are poorly fitted.

We divide the spectrum into four segments, fitting the RVs in each segment independently. This allows us to empirically estimate RV uncertainties from the scatter between different segments. It also allows us to test whether the RVs of the emission lines, which trace the Be star’s disk and circumstellar envelope, vary coherently with the absorption lines, which trace the photosphere. In particular, the region of the spectrum with $5100 < \lambda/\text{\AA} < 5400$ contains no strong absorption lines from the Be star but contains seven Fe II emission lines, which likely originate in a disk around the Be star.

As discussed by [El-Badry & Quataert \(2020\)](#) and [Abdul-Masih et al. \(2020\)](#), RVs measured from emission lines can be biased if contamination from absorption lines is not taken into account. Our procedure for measuring RVs is not susceptible to the bias pointed out in these works, because the spectra of both components are modeled simultaneously.

Figure 9.4 shows the measured epoch RVs for both components, phased to a period of 40.3 days. Error bars in the top two panels show the scatter between RVs fit to different spectral regions. The Be star displays sinusoidal RV variability with the same period as the B star, suggesting that they are orbiting each other. Emission and absorption lines show consistent variability, implying that the Be star and its disk move coherently. This implies that the disk is likely circumstellar, not circumbinary. For the B star, the RV uncertainties are smaller than the symbols for most points, suggesting that there is intrinsic RV scatter.

We fit the RVs of both components independently with Keplerian orbits, without the requirement that they have the same center-of-mass velocity or phase. Along with the 6 standard Keplerian orbital parameters, we fit for an intrinsic RV scatter or “jitter” term (see [El-Badry et al. 2018g](#), their Equation 9), which can represent either intrinsic RV variability not captured in the model or underestimated uncertainties. We initially search for the best-fit orbital solution using simulated annealing, and then sample from the posterior in the vicinity of this solution using a Markov chain Monte Carlo method, as described in [El-Badry et al. \(2018g\)](#). Given the large number of RV measurements (51), there is little danger of the sampler getting stuck in a local minimum. Our best-fit orbital parameters are listed in Table 9.1.

Our best-fit velocity amplitude for the B star is $K_B = 62.7 \pm 1.0 \text{ km s}^{-1}$, which is consistent with the value measured by [Rivinius et al. \(2020\)](#). Our constraints on P is also consistent, but our uncertainties are larger, because we did not include the additional RVs from 1999 in our fit. Unlike [Rivinius et al. \(2020\)](#), we do not measure any significant eccentricity: we find an upper limit of $e < 0.037$ for the B star, but the orbit is consistent with $e = 0$ once RV scatter is accounted for. For the Be star, we measure $e < 0.13$, again consistent with a circular orbit. The formal uncertainty in K_{Be} obtained from fitting the RVs shown in Figure 9.4 is only 0.3 km s^{-1} , but the true uncertainty is larger, since the points shown in Figure 9.4 were obtained after fixing $K_{Be} = 4.5 \text{ km s}^{-1}$ in spectral disentangling (see Appendix 9.6).

For the B star, we find an intrinsic RV scatter of $\approx 4 \text{ km s}^{-1}$. We did not find evidence of additional periodicity in the residuals. As we show in Appendix 9.7, the B star’s RV scatter is accompanied by modest variability in its spectral line profiles. HR 6819 also exhibits photometric variability with brightness variations of up to five percent in the total




Figure 9.4: Radial velocities for the B star (top panel) and Be star (middle and bottom panels), phased to a period of 40.3 days. RVs for both components are measured simultaneously by fitting the composite spectra with the single-star templates obtained from spectral disentangling. The Be star exhibits sinusoidal RV variability in anti-phase with the B star, suggesting the two components are orbiting each other with a period of 40.3 days. The much lower velocity amplitude of the Be star implies a mass ratio $M_B/M_{Be} = 0.071 \pm 0.032$. The bottom panel compares velocities for the Be component measured from absorption lines (presumably originating in the stellar photosphere) and emission lines (presumably originating in a circumstellar disk). Emission and absorption line RVs vary coherently.

light (implying variations of 10% if the light comes primarily from one component) on timescales down to half a day (Rivinius et al. 2020). The most likely explanation is that the B star is a slowly pulsating B star (SPB star; e.g. Waelkens 1991) undergoing non-radial g -mode pulsations that cause the RV scatter, line profile variability, and photometric variability. The B star’s atmospheric parameters, the amplitude of the RV and photometric variability, and the variability timescale are all within the normal range found for SPB stars (e.g. Dziembowski et al. 1993). We note, however, that photometric variability is also common for Be stars, due to both pulsations and variations in the structure of the disk (Rivinius et al. 2013), so some of the observed photometric variability could also come from the Be star.

If the masses of both components are known, the inclination of the binary’s orbital plane can be measured from the velocity semi-amplitudes. For a circular orbit,

$$\sin^3 i = \frac{(1 + K_{\text{Be}}/K_{\text{B}})^2 K_{\text{B}}^3 P}{2\pi G M_{\text{Be}}}. \quad (9.1)$$

Taking $M_{\text{Be}} = 6.7_{-1.5}^{+1.9} M_{\odot}$, a constraint which we obtain by comparing the Be star’s temperature and surface gravity to MIST evolutionary tracks (see Appendix 9.5; Dotter 2016; Choi et al. 2016), we obtain $i = 32.1_{-2.6}^{+3.0}$ deg. The mass ratio of the two components is constrained from the relative velocity amplitudes to $M_{\text{B}}/M_{\text{Be}} = 0.071 \pm 0.032$. This leads to $M_{\text{B}} = 0.47_{-0.22}^{+0.28} M_{\odot}$, where the dynamical mass ratio and the mass of the Be star contribute roughly equally to the uncertainty.

9.3.5 Luminosity

If the distance to HR 6819 is known, then its absolute magnitude can distinguish between models in which the B star is a normal star and models in which it was recently stripped. The bolometric luminosity of a star can be approximated as

$$\frac{L}{L_{\odot}} = \left(\frac{R}{R_{\odot}}\right)^2 \left(\frac{T_{\text{eff}}}{T_{\text{eff},\odot}}\right)^4 = \left(\frac{M}{M_{\odot}}\right) \left(\frac{g}{g_{\odot}}\right)^{-1} \left(\frac{T_{\text{eff}}}{T_{\text{eff},\odot}}\right)^4 \quad (9.2)$$

$$= 1400 \left(\frac{M}{0.5M_{\odot}}\right) \left(\frac{10^{\log g}}{10^{2.75}}\right)^{-1} \left(\frac{T_{\text{eff}}}{16 \text{ kK}}\right)^4. \quad (9.3)$$

As discussed by Rivinius et al. (2020), a normal star with the temperature and gravity of the B star in HR 6819 would have a minimum mass of at least $5 M_{\odot}$. It would thus be a minimum of 10 times brighter than a $0.5M_{\odot}$ stripped star with the same T_{eff} and $\log g$. And as we discuss below, the minimum mass for a normal B star matching HR 6918’s spectral type is even higher, of order $10 M_{\odot}$, once the lower $\log g$ measured from the disentangled spectrum is taken into account.

Distance

Gaia measured a parallax of $\varpi = 2.71 \pm 0.12$ mas for HR 6819 (*Gaia* DR3 `source_id` 6649357561810851328), implying a distance $d \approx 369 \pm 17$ pc (Gaia Collaboration et al. 2020a).² At this distance, the projected semimajor axis would be $\theta \approx 1.2$ mas. Since this is almost half the parallax, and the astrometric solution did not account for orbital motions, the parallax uncertainty is likely underestimated. The re-normalized unit weight error for the source is large, `ruwe` = 1.45, also indicating a problematic astrometric solution (e.g. Fabricius et al. 2020). Given that the *Gaia* data span many orbital periods, the *Gaia* distance is not expected to be catastrophically biased, but the parallax uncertainty should not be taken at face value.

Rivinius et al. (2020) derived an independent estimate of the distance to HR 6819 by comparing its flux-calibrated UV spectrum to models with temperature and radius matching their fit to the spectrum. They estimated $d = (310 \pm 60)$ pc, which is also consistent with the *Gaia* distance. The largest source of uncertainty was the unknown luminosity ratio between the two components. Although the spectral type we infer for the B star is significantly different from that estimated by Rivinius et al. (2020), only the temperature and radius are relevant for the distance estimate. The values they adopted ($T_{\text{eff}} = 16$ kK and $R = (5.5 \pm 0.5)R_{\odot}$) are consistent with our best-fit values, because our lower mass estimate is almost exactly compensated for by our lower estimate of $\log g$.

El-Badry et al. (2021) found that bright sources with `ruwe` > 1.4 typically have parallax uncertainties underestimated by up to a factor of 2. We adopt the *Gaia* distance estimate in our analysis, but we conservatively inflate the parallax uncertainty by a factor of ≈ 2 due to the large `ruwe` and expected orbital motion. We assume $d = 369 \pm 34$ pc. To account for the possibility of more severe problems with the *Gaia* distance, we also present stellar parameters that are inferred without taking the distance into account.

Synthetic photometry

We used synthetic photometry from MIST isochrones (Dotter 2016; Choi et al. 2016) to place models for HR 6819 on the color-magnitude diagram (CMD), assuming solar metallicity. We assume an extinction of $E(B - V) = 0.135$, as reported by Wegner (2002), and a Cardelli et al. (1989) extinction curve with $R_V = 3.1$. The SFD map estimates $E(B - V) = 0.13$ at the position HR 6819 (Schlegel et al. 1998).

We began by searching the MIST model grid for “normal” stellar models with atmospheric parameters similar to the B star ($15 < T_{\text{eff}}/\text{kK} < 17$ and $2.4 < \log g < 3.1$). These models have masses ranging from $9M_{\odot}$ to as much as $20M_{\odot}$ (see Figure 9.10). This is significantly higher than the mass of $(5 - 6)M_{\odot}$ that would be inferred if the B star had $\log g = 3.5 - 4$, as assumed by Rivinius et al. 2020. These models are shown in green in the right panel of Figure 9.5. Even without the contribution from the Be star, these models are far brighter

²The *Gaia* DR2 parallax was $\varpi = 2.91 \pm 0.18$ mas. The distance errors implied by the parallax are slightly asymmetric, $d = 369_{-16}^{+17}$ pc; we symmetrize them for simplicity.

than HR 6819.

We next calculate synthetic photometry for a stripped star with the same range of T_{eff} and $\log g$ as the models described above, accounting for the lower mass using Equation 9.3. These models are shown in magenta in Figure 9.5; we assume masses between 0.4 and 0.6 M_{\odot} for the stripped star. We also take synthetic photometry for the Be star from models with $15 < T_{\text{eff}}/\text{kK} < 20$ and $3.5 < \log g < 4.25$. These models are shown in cyan; we do not attempt to correct for gravity darkening due to the Be star’s rapid rotation. Finally, we calculate the combined magnitude and color of randomly paired Be star and stripped star models (orange points). These are generally consistent with the color and magnitude of HR 6819.

This analysis reveals that the B star in HR 6819 cannot be a normal B supergiant – its luminosity is too low. A consistent solution can, however, be obtained if the B star has unusually low mass, of order $0.5M_{\odot}$. The poor agreement between a “normal B star” model and the observed luminosity of HR 6819 cannot be alleviated by simply assuming a different distance. A distance of 3-5 times our fiducial value would make the observed magnitude of HR 6819 compatible with a normal B star (green points in Figure 9.5). But in this case, the B star would be a factor of ~ 20 brighter than the Be star, violating the observational constraint that the two components are roughly equally bright.

Luminosities, masses and radii

We estimate the mass and radius of the Be star by comparing its temperature and surface gravity to a grid of MIST evolutionary tracks (Figure 9.10). We then calculate the mass of the B star from the measured mass ratio, and its radius from this mass and its measured surface gravity. The bolometric luminosity of each component is then calculated through Equation 9.2. The resulting stellar parameters are reported in Table 9.1.

In Appendix 9.5, we calculate an estimate of the same stellar parameters that is anchored on the distance to HR 6819 rather than on evolutionary tracks. We obtain consistent solutions between the two methods.

9.3.6 Emission lines

As is often the case in classical Be stars, the disentangled Be star spectrum contains clearly double-peaked emission in the Balmer lines and in many Fe II lines. Weak double-peaked emission lines are also found in some C II and Si II lines. Figure 9.6 highlights the $H\alpha$ line (which is by far the strongest of the star’s emission lines) and several of the Fe II lines. We correct the emission lines for contamination from the Be star’s absorption lines by subtracting a TLUSTY model spectrum with $T_{\text{eff}} = 18\text{ kK}$, $\log g = 3.75$, and $v \sin i = 180\text{ km s}^{-1}$. This has little effect on the Fe II lines, but it increases the amplitude and slightly changes the shape of the $H\alpha$ line.

The integrated emission line profile of an optically thin, kinematically cold Keplerian disk with known emissivity profile can be calculated analytically (e.g. Smak 1969). Following Shen

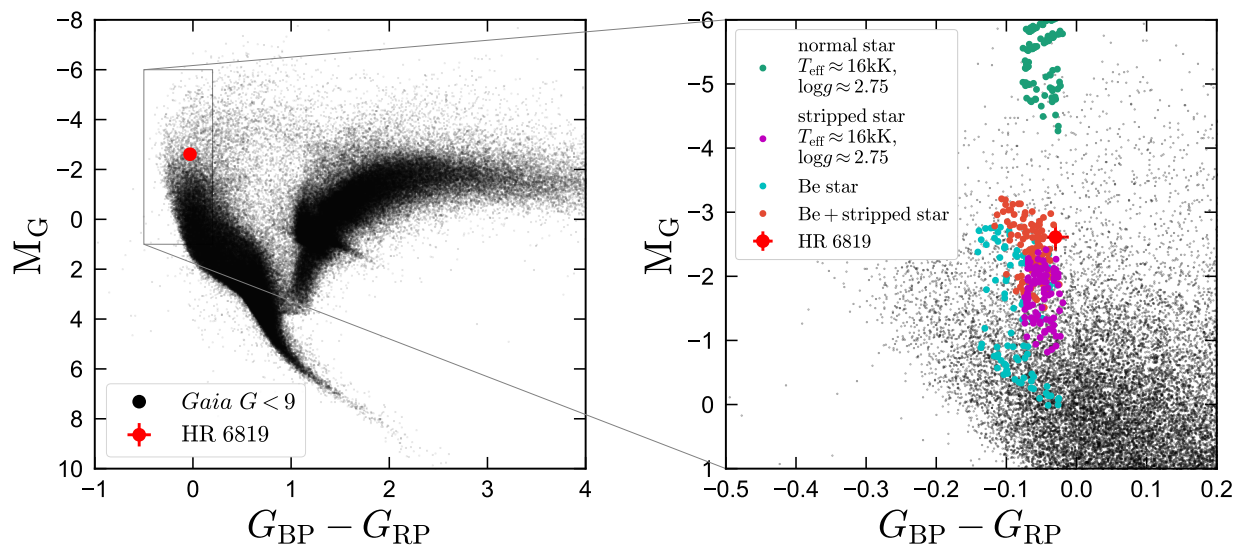


Figure 9.5: Color-magnitude diagram. Left panel compares HR 6819 to all *Gaia* sources brighter than $G = 9$. Right panel compares it to synthetic photometry from isochrones. Green points show predictions for a normal star, near the end of the main sequence, with temperature and surface gravity consistent with HR 6819. Such a star would be much brighter than HR 6819 and is ruled out. Magenta points show the CMD position of a stripped star with the same temperature and gravity. Due to its lower mass and smaller radius, a stripped star is much fainter. Cyan points show models with temperature and gravity consistent with our measurements for the Be star. Orange points show unresolved binaries composed of a Be star and stripped star model; these are consistent with the observed color and magnitude of HR 6819.

et al. (2019), we assume that the emissivity profile can be approximated by a power law, $j(r) \propto r^{-\alpha}$, and that the emissivity is sharply truncated at inner radius $r = R_{\text{inner}}$ and outer radius R_{outer} . This truncation could either represent the geometric edge of the disk, or a change in the ionization state. We assume an intrinsic thermal broadening with isothermal sound speed $c_s = 10 \text{ km s}^{-1}$, appropriate for ionized gas with temperature near 10^4 K .

Given these assumptions, we fit the emission line profile of the Fe II 5169 line. The free parameters of the fit are the projected rotation velocity at the outer edge of the disk, $v_{\text{out}} \sin i = (GM_{\text{Be}}/R_{\text{outer}})^{1/2} \sin i$, $r_{\text{inner}} = R_{\text{inner}}/R_{\text{outer}}$, the ratio of the inner disk edge to the outer disk edge, and the emissivity exponent α . We obtain $\alpha \approx 1.6$, $v_{\text{out}} \approx 105 \text{ km s}^{-1}$, and $r_{\text{inner}} \approx 0.24$ (see Table 9.1). This leads to $R_{\text{inner}} \approx 7.8R_{\odot}$ and $R_{\text{outer}} \approx 33R_{\odot}$. These parameters can be interpreted straightforwardly from the observed line profiles: R_{inner} is constrained by the maximum width of the emission lines (since the largest Keplerian velocities are found at the inner edge of the disk) and R_{outer} is constrained by the RV separation of the two peaks (since the total emission is maximized at the outer edge of the disk).

We note that applying the same fitting to the H α line yields significantly different results, including an unphysically small R_{inner} . This is because repeated absorption and re-emission of H α photons, as well as electron scattering, significantly broadens the wings of the H α line. This makes it an unreliable kinematic tracer (e.g. Hummel & Dachs 1992). Fe II lines have lower optical depth and are therefore more reliable optical tracers of the disk (Hanuschik 1988; Dachs et al. 1992), though they too many not be completely optically thin in the inner parts of the disk (Arias et al. 2006).

For our best-fit orbital parameters, the binary’s semimajor axis is $a \approx 0.44 \text{ AU} \approx 96 R_{\odot}$. That is, the disk emission extends to about $0.34a$. For a binary of mass ratio 0.07, the largest stable streamline at which a disk can be maintained is $r_{\text{max}} \approx 0.44a$ (Paczynski 1977), so the disk likely extends to near its maximum stable radius.

9.3.7 Chemical abundances

We use SYNSPEC (Hubeny & Lanz 2011) to generate grids of model spectra with $T_{\text{eff}} = 16 \text{ kK}$ and $\log g = 2.75$, varying each of the 11 elements listed in Table 9.1 with a grid spacing of 0.1 dex. All abundances are measured relative to the Solar abundance pattern, which is taken from Grevesse & Sauval (1998). The atmospheric structure is taken from the BSTAR06 atmosphere grid, which was generated with TLUSTY assuming $Z = 0.02$ and the Solar abundance pattern. We assume a microturbulent velocity $v_{\text{mic}} = 10 \text{ km s}^{-1}$ but also experimented with varying T_{eff} , $\log g$, and v_{mic} within their uncertainties.

Unlike the line ratios shown in Figure 9.3, the abundances are sensitive to the luminosity ratio of the two components. The smaller the ratio $L_{\text{B}}/L_{\text{Be}}$, the deeper the lines of the B star, and the larger the inferred abundances. Our reported abundance uncertainties account for a $\pm 5\%$ uncertainty in the luminosity ratio. Uncertainties in T_{eff} , v_{mic} , and the flux ratio, not the noise in the spectrum, dominate the abundance uncertainties.

The most striking result of our abundance analysis is that the surface of the B star is significantly enriched in nitrogen and helium, and significantly depleted of carbon. This




Figure 9.6: Line emission from the disk and circumstellar envelope of the Be star. Top panel shows a section of the disentangled Be star’s spectrum that displays several Fe II emission lines. Each line is double-peaked, as is expected for an optically thin Keplerian disk. Bottom left panel compares the line shapes of H α and one of the Fe II lines. The H α line has broader wings and a narrower peak, likely due to its higher optical depth. Bottom right panel compares the shape of the Fe II line to a model fit assuming a truncated, optically thin disk (see text). The best-fit model has an inner radius $R_{\text{inner}} \approx 8.6R_{\odot}$ and outer radius $R_{\text{outer}} \approx 35R_{\odot}$, reasonably consistent with the observed stellar radius and orbital semi-major axis.

suggests that material on the B star’s surface has been processed by the CNO cycle, as has also been reported for LB-1 (Irrgang et al. 2020; Shenar et al. 2020). An enhancement of nitrogen and helium at the expense of carbon and oxygen is expected in the stripped star scenario: the material at the surface of the star today was inside the convective core of the progenitor when it was on the main sequence and is thus contaminated with fusion products.

Unlike in LB-1, we do not find a significant depletion of oxygen for the B star. Most other elements are consistent with, or slightly enhanced relative to, the Solar value. We note that Irrgang et al. (2020) found most metals to be depleted in LB-1, but this is likely at least in part due to the unaccounted-for line dilution by the Be star, which makes all lines appear weaker by roughly a factor of 2.

We do not attempt to measure detailed abundance for the Be star, since most metal lines are washed out by rotation, uncertainties in continuum normalization will significantly affect the interpretation of weak and highly broadened lines, and contamination by disk emission complicates the interpretation of absorption line equivalent widths. Besides H and He, the strongest lines visible in the Be star spectrum come from C II and Si II. The strongest tension between the disentangled Be star spectrum and the TLUSTY model occurs for the 4075 Å CII line, which is much weaker than predicted. This suggests that the Be star’s surface is also depleted of C, which is consistent with our expectations if the Be star recently accreted much of the CNO-processed envelope of the B star.

We defer more detailed analysis and modeling of HR 6819’s abundance pattern to future work. Given the B star’s unusually large surface He abundance, it would be useful to recalculate the atmospheric structure while taking the abundance variations into account.

9.3.8 Evolutionary history

To better understand the possible formation pathways and future evolution of HR 6819, we began by searching the BPASS (v2.2; Eldridge et al. 2017) library of binary evolution models for models that go through a phase similar to the current state of HR 6819. We then investigated a few possible formation histories in more detail using Modules for Experiments in Stellar Astrophysics (MESA, version 12778; Paxton et al. 2011, 2013, 2015, 2018, 2019).

We search the $Z = 0.02$ BPASS grid and set the following constraints:

- $15 < T_{\text{eff},1}/\text{kK} < 17$
- $2.5 < \log g_1 < 3.25$
- $5 < M_2/M_\odot < 10$
- $20 < P/\text{days} < 60$
- $2.7 < \log(L_1/L_\odot) < 3.3$

We refrained from placing constraints on the temperature or luminosity of the secondary, because the evolutionary models in the BPASS do not follow the secondary’s evolution with




Figure 9.7: Initial masses and periods of models in the BPASS binary evolution library that at any point in their evolution pass through an evolutionary state similar to the current state of HR 6819.

detailed calculations following mass transfer. Because BPASS uses a grid of discrete models, we allow a relatively large range of periods. This search yielded 10 models, with initial primary masses ranging from 4.5 to 7 M_{\odot} , initial mass ratios ranging from 0.4 to 0.8, and initial periods of 1.5 to 6.5 days. These are shown in Figure 9.7.

We then used MESA to calculate binary evolution tracks for models in this region of parameter space, and to investigate how the systems’ evolution changes when their parameters are varied. Most of the single-star free parameters are set following the `inlist_7M_premsto_AGB` `inlist` in the MESA test suite. We use the wind mass loss prescription from Reimers (1975). The initial metallicity is set to 0.02. We use a mixing length parameter $\alpha_{\text{MLT}} = 1.73$ and an exponential overshoot scheme with overshoot efficiency parameter $f_{\text{ov}} = 0.014$. The overshooting scheme and motivation for this choice of f_{ov} are described in Herwig (2000).

The `MESAbinary` module is described in detail in Paxton et al. (2015). Both stars are evolved simultaneously, with tides and mass transfer taken into account. Roche lobe radii are computed using the fit of Eggleton (1983). Mass transfer rates in Roche lobe overflowing systems are determined following the prescription of Kolb & Ritter (1990). Tidal torques

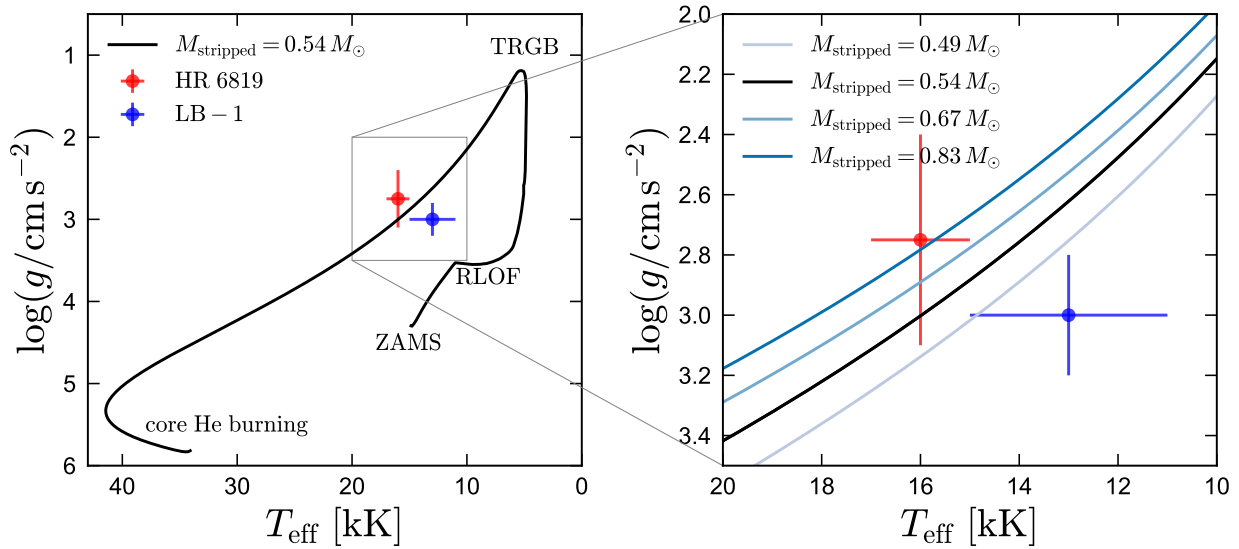


Figure 9.8: Evolutionary history of the primary for stripped star models similar to HR 6819 and LB-1. Left panel shows a model with initial masses $M_1 = 4M_\odot$ and $M_2 = 3M_\odot$, and initial period 2.3 days. This produces a stripped star with mass $0.54M_\odot$. We indicate the locations of the zero-age main sequence (ZAMS), first Roche lobe overflow (RLOF), tip of the red giant branch (TRGB), and stable core helium burning. Right panel compares models with a range of stripped star masses, corresponding to initial masses between 3.5 and $7M_\odot$.

are modeled following Hut (1981). The orbital separation evolves such that the total angular momentum is conserved when mass is lost through winds or transferred to a companion, as described in Paxton et al. (2015). We experimented with a range of mass transfer efficiencies. When mass transfer is not conservative, the mass that escapes the binary is modeled as being lost from the vicinity of the accretor through a fast wind, as described by Tauris & van den Heuvel (2006, their “ β ” parameter).

Our MESA models are summarized in Table 9.2. Figure 9.8 shows the evolution of 4 models in $T_{\text{eff}} - \log g$ space. The left panel shows a model with initial masses $M_1 = 4M_\odot$, $M_2 = 3M_\odot$, and an initial period of 2.3 days. The components of the binary evolve normally until the end of the primary’s main sequence lifetime. Shortly before the primary crosses the Hertzsprung gap, it overflows its Roche lobe and begins stable mass transfer to the companion. The primary’s envelope expands further as more mass is removed, and the star moves up the red giant branch. At the tip of the red giant branch (TRGB), the temperatures required to maintain hydrogen burning can only be maintained if the envelope contracts. At this point, when the primary’s luminosity is of order $1000L_\odot$, the primary leaves the RGB and passes through the region of $T_{\text{eff}} - \log g$ space where HR 6819 and LB-1 are currently observed. The star then continues to contract and heat until it reaches $T_{\text{eff}} \approx 30$ kK and $\log g \approx 5.5$, where it settles as a core helium burning hot subdwarf (sdOB star). The core helium burning phase would last of order 10^8 years, but for most of the models we consider,

Table 9.2: Summary of our MESA binary models.

$M_{1,\text{init}}$ [M_{\odot}]	$M_{2,\text{init}}$ [M_{\odot}]	P_{init} [days]	f_{mt}	P_{stripped} [days]	M_{stripped} [M_{\odot}]	$M_{\text{He core}}$ [M_{\odot}]	M_{env} [M_{\odot}]	M_{Be} [M_{\odot}]	t_{RLOF} [Myr]	t_{TRGB} [Myr]	t_{sdOB} [Myr]	bloated lifetime [10^5 yr]
3.5	2.8	2.3	1	90	0.49	0.43	0.06	5.8	232	269.9	272.7	2.07
4.0	3.0	2.3	1	91	0.54	0.47	0.07	6.4	162	203.8	205.5	1.29
5.5	4.2	2.3	0.3	152	0.67	0.11	0.56	5.6	71	97.1	98.7	2.20
7.0	4.0	2.3	0.2	79	0.83	0.15	0.68	5.2	41	55.8	57.2	2.27

$M_{1,\text{init}}$, $M_{2,\text{init}}$, and P_{init} are the initial masses and period at the zero-age main sequence. f_{mt} is the efficiency of mass transfer, defined such that a fraction f_{mt} of mass lost by the donor is gained by the accretor. P_{stripped} is the orbital period during the bloated stripped star phase, when $2 < \log g < 3.5$ (roughly when the star passes through the inset in Figure 9.8). M_{stripped} , $M_{\text{He core}}$, and M_{env} , are the total, core, and envelope mass of the initial primary (the B star) during this period, and M_{Be} is the mass of the companion during this period. t_{RLOF} , t_{TRGB} , and t_{sdOB} are the system age when the primary first overflows its Roche lobe, when it reaches the tip of the red giant branch, and when its luminosity is first dominated by He burning. The “bloated stripped phase lifetime” is the time it takes the primary to contract from $\log g = 2$ to $\log g = 3.5$ as it moves from the TRGB toward the extreme horizontal branch.

it is truncated prematurely when the secondary (the Be star) leaves the main sequence and overflows its Roche lobe.

The right panel compares stripped star models with masses between 0.49 and 0.83 M_{\odot} . These are produced from calculations with initial primary masses of 3.5, 4.5, 5.5, and 7 M_{\odot} . We reduce the efficiency of mass transfer (which is poorly constrained a priori) as the primary mass increases, so that the mass of the secondary during the stripped star phase is always about 6 M_{\odot} . As the total mass of the core increases, its radius and luminosity increase, and $\log g$ decreases. However, the sensitivity of these parameters to mass is relatively weak in this region of parameter space, so HR 6819’s atmospheric parameters are consistent with models with a broad range of core masses (and hence initial masses).

Figure 9.9 shows the evolution of the model with $M_1 = 4M_{\odot}$ and $M_2 = 3M_{\odot}$ after the primary begins to contract. We designate the “bloated stripped star” phase of the binary’s evolution as the phase after the tip of the red giant branch when the primary (the B star) has $\log g$ between 2 and 3.5. This is roughly the period when the primary has the temperature of a B star and is more inflated than a main sequence star of the same temperature. For the 4 models shown in Figure 9.8, the duration of this phase ranges from 1.3 to 2.3×10^5 years. The secondary (i.e., the Be star) does not evolve on the timescale shown in Figure 9.9: for all our models, the mass transfer timescale near the TRGB is significantly longer than the thermal timescale of the secondary, so the secondary remains in thermal equilibrium.

We note that, contrary to what is sometimes assumed in the literature (e.g. Eldridge et al. 2020), the contraction from the red giant branch to the extreme horizontal branch does not occur on a Kelvin-Helmholtz timescale of the donor star – the bottom panel of Figure 9.9 shows that it is longer by more than a factor of 10. The reason for this is that gravitational




Figure 9.9: Time-evolution of the MESA model with $M_{\text{stripped}} = 0.54M_{\odot}$ during the stripped star phase (black line in Figure 9.8). From top to bottom, panels show surface gravity, luminosity, mass, and Kelvin-Helmholtz timescale. Time is measured since the tip of the red giant branch (i.e. when $\log g$ reaches its minimum). The luminosity is dominated by H shell fusion, which burns through most of the remaining envelope as the star contracts. He fusion is ignited as contraction begins, but does not dominate over shell fusion until 1.5×10^6 years after the TRGB (beyond the time range of the plot). The lifetime of the contraction phase is a few $\times 10^5$ years. This is significantly longer than the Kelvin-Helmholtz timescale (bottom panel), because gravitational contraction contributes only a small fraction of the luminosity.

contraction is not the dominant source of energy during this phase; hydrogen shell burning and core helium burning both contribute luminosity to slow the star’s contraction.

The final period produced by these calculations is somewhat longer ($P = 79 - 152$ days) than the observed $P = 40.3$ days in HR 6819. This can be understood as a result of the orbit’s expansion during mass transfer. When mass transfer is conservative, the final period can be calculated analytically to be

$$P_f = P_i \left(\frac{M_{\text{donor},i}}{M_{\text{donor},f}} \times \frac{M_{\text{accretor},i}}{M_{\text{accretor},f}} \right)^3. \quad (9.4)$$

Here P_i and P_f represent the initial and final period, $M_{\text{donor},i}$ and $M_{\text{donor},f}$ the initial and final masses of the donor, and $M_{\text{accretor},i}$ and $M_{\text{accretor},f}$ the initial and final masses of the accretor. In our calculations $M_{\text{donor},i}/M_{\text{donor},f} = 7 - 8.5$, while $M_{\text{accretor},i}/M_{\text{accretor},f} = 1 - 2.1$. This causes the final period to be a factor of $\gtrsim 40$ larger than the initial period. The final period according to Equation 9.4 could be shorter if the initial mass ratio were more unequal ($M_1 \gg M_2$), but in this case mass transfer eventually becomes dynamically unstable, leading to an episode of common envelope evolution (CEE) and typically inspiral to a much shorter period (e.g. Heber 2016).

No analytic expression for the final period exists for non-conservative mass transfer. How the angular momentum of a binary changes following mass loss depends on how the mass is lost (see Soberman et al. 1997, and Tauris & van den Heuvel 2006, their Equation 16.18). In MESABinary, it is possible to specify what fraction of transferred mass is lost from the vicinity of the donor, from the vicinity of the accretor, and from a circumbinary toroid. Particularly near the end of the stripping process, when $M_{\text{donor}}/M_{\text{accretor}} \ll 1$, the orbital response to mass loss, which is parameterized through the Tauris & van den Heuvel (2006) subgrid model, depends quite sensitively on the assumed mass loss fractions for each channel. For our calculations with non-conservative mass transfer, the final period decreases by up to a factor of 10 if we assume mass is lost from a circumbinary toroid rather than from a wind in the vicinity of the accretor. Because we have limited physical intuition for what the efficiency of mass transfer is in this problem or how mass that leaves the binary escapes, the observed period cannot reliably constrain the binary’s initial period.

We note that while luminosity and radius increase monotonically with the total mass of the stripped star in our calculations (right panel of Figure 9.8), this is not necessarily expected to generically be the case. The radius and luminosity of the stripped star depend on the relative masses of the envelope and the He core, with lower envelope masses corresponding to larger radii and luminosities at fixed temperature. For example, Driebe et al. (1998) calculated evolutionary tracks for low-mass stripped pre-He WDs formed from a $1M_{\odot}$ progenitor by artificially enhanced mass loss on the RGB. Their $0.41 M_{\odot}$ track is nearly coincident with our $0.83M_{\odot}$ track, while their $0.33 M_{\odot}$ track falls below our $0.49M_{\odot}$ track and is consistent with the LB-1 atmospheric parameters (see also Irrgang et al. 2020). In HR 6819, both the dynamical mass and atmospheric parameters of the stripped star are consistent with a scenario in which it is a pre-He white dwarf, with a degenerate core and too low a mass to ever ignite helium. However, our binary evolution calculations do not produce

stripped stars with such low mass. Because the stripped star must have initially been the more massive component of the binary, and the total initial mass must have been at least $\gtrsim 6M_{\odot}$, the initial mass of the stripped star must have been at least $3M_{\odot}$. Stars with initial masses near $3M_{\odot}$ *can* produce stripped He stars with masses of order $0.4M_{\odot}$, but these ignite He burning on the RGB. Because their cores are not degenerate during the bloated stripped phase, they have thicker hydrogen envelopes, higher $\log g$, and lower luminosity than the [Driebe et al. \(1998\)](#) models.

We have not performed an exhaustive search of parameter space with detailed binary evolution calculations, so it is possible that another region of parameter space could produce systems similar to HR 6819 and LB-1, even though no other satisfactory models were found in the BPASS library. Below, we briefly summarize how the MESA models change when the initial conditions or model parameters are perturbed.

- *Longer periods:* For longer initial periods, mass transfer begins later, when the primary has expanded more. For $P_{\text{init}} \lesssim 20$ days, evolution is similar to our fiducial models, but stripping occurs over a shorter timescale, the final period is longer, and a somewhat larger fraction on the envelope remains when contraction begins, shifting models downward (toward larger $\log g$) by up to ≈ 0.2 dex in [Figure 9.8](#).

For $20 \lesssim P_{\text{init}}/\text{days} < 150$, mass transfer begins when the envelope is convective and the core mass is too small to support stable mass transfer, leading to dynamical instability and presumably an episode of CEE. MESA does not include robust treatment of CEE, but it is expected to result in shorter periods than the observed $P = 40.3$ days (e.g. [Ivanova et al. 2013](#)).

For $150 \lesssim P_{\text{init}}/\text{days} < 2000$, the primary begins core He burning before RLOF. Stripping then begins on the asymptotic giant branch (AGB), where the larger core mass allows mass transfer to remain dynamically stable (e.g. [Chen & Han 2008](#)). Eventually, the primary is reduced to a carbon-oxygen core with a thin helium envelope, whose contraction is similar to that of a single proto-white dwarf. These models pass through the same region of $T_{\text{eff}} - \log g$ space highlighted in [Figure 9.8](#), typically shifted upward (toward smaller $\log g$) somewhat from our fiducial models of the same total mass. However, they contract much more rapidly than our fiducial models, with lifetimes of $\lesssim 10^3$ years. The final periods are also much longer, $P \gtrsim 3000$ days, though we note that these depend sensitively on the assumptions made about how mass is lost from the binary.

- *Shorter periods :* For shorter periods, RLOF occurs earlier, when the primary is still on the main sequence. For $P_{\text{init}} \lesssim 2$ days, the primary loses about half its mass during the initial phase of RLOF, but then essentially continues its evolution as a lower-mass main sequence star. Meanwhile, the now-more-massive secondary evolves faster than the primary, reaching the end of the main sequence before the primary. Such models never resemble HR 6819.

- *More unequal mass ratios:* If the initial masses of the two components are too unequal ($M_2/M_1 \lesssim q_{\text{crit}}$, where q_{crit} depends on the structure of the donor; e.g. [Hjellming & Webbink 1987](#); [Han et al. 2002](#); [Ge et al. 2010](#)), mass transfer eventually becomes dynamically unstable, leading to an episode of CEE. When mass transfer begins, the envelope is radiative, and mass transfer is stable for a wide range of mass ratios. However, the envelope becomes convective soon after mass transfer begins, and mass transfer then becomes unstable. We note that the secondary is generally not expected to accrete much mass during CEE; this also speaks against an initially low secondary mass.
- *Higher primary masses :* As the mass of the primary increases, the efficiency of mass transfer must approach 0 to avoid making the Be star too massive. If this is achievable physically, then a higher primary mass produces leads to a more massive stripped He star. For our default parameters, $M_{1,\text{init}} = 10 M_\odot$ leads to a $1.3 M_\odot$ stripped He star. During the contraction phase, the $1.3 M_\odot$ model continues the monotonic trend seen in [Figure 9.8](#) of lower $\log g$ with increasing mass at fixed T_{eff} .
- *Lower primary masses:* For primary masses $M_{1,\text{init}} \lesssim 3 M_\odot$, the total initial mass of the binary is less than its current total mass, so the observed mass of the Be star cannot be reproduced.
- *More or less overshooting:* The treatment of overshooting changes the mapping between initial mass and stripped star mass. The larger the overshooting efficiency parameter, the larger the convective core on the main sequence, and the larger the He core. Increasing f_{ov} thus leads to higher M_{stripped} for fixed $M_{1,\text{init}}$, shifting models toward lower $\log g$ at fixed T_{eff} . For the models shown in [Figure 9.8](#), increasing f_{ov} from 0.014 to 0.028 increases M_{stripped} by $\approx 0.12 M_\odot$ on average, but does not significantly change the lifetime of the “bloated stripped star” phase.

We note that while the models shown in [Figure 9.8](#) allow for a consistent solution between the dynamically-inferred mass of the stripped star in HR 6819 and its position in $T_{\text{eff}} - \log g$ space, we have not found a consistent solution for LB-1. The T_{eff} and $\log g$ (taken from [Shenar et al. 2020](#)) imply a stripped star mass $M_{\text{stripped}} \lesssim 0.5 M_\odot$, which can be relaxed to $M_{\text{stripped}} \lesssim 0.6 M_\odot$ if we allow for longer initial periods. However, [Shenar et al. \(2020\)](#) report a dynamical mass of $(1.5 \pm 0.5) M_\odot$, which is significantly higher than this. A $1.5 M_\odot$ He star is generically difficult to produce in the context of our models (or the BPASS models); if stripping occurs near the Hertzsprung gap, it requires a primary mass of order $12 M_\odot$ for our adopted overshooting parameter. Unless mass transfer is highly inefficient, the Be star would end up too massive in such a scenario. Without binary interactions, a single star of solar metallicity must have initial mass $\gtrsim 9 M_\odot$ to reach a He core mass of $1.5 M_\odot$ prior to core He burning.

The dynamical mass constraint derived by [Shenar et al. \(2020\)](#) comes primarily from the $\text{H}\alpha$ line. Since spectral disentangling relies on the assumption that the individual spectra of

both components are time-invariant, and this assumption likely does not hold in detail for the $H\alpha$ line of LB-1 due to phase-dependent disk irradiation by the B star (Liu et al. 2019), it is possible that the dynamical mass measurement of the B star is biased and the true mass is lower. If the mass measurement is reliable, it would imply a significantly higher envelope mass ($M_{\text{env}}/M_{\text{He core}} \sim 1$) than produced by our models, since the luminosity during the stripped phase is set mainly by the He core mass. Three of the five confirmed Be + sdOB binaries have estimated masses above $1.0M_{\odot}$ (Wang et al. 2017), so a higher stripped-star mass is not beyond the realm of possibility. This is, however, at least partially a selection effect, since higher-mass sdOBs are hotter and easier to detect (e.g. de Mink et al. 2014; Wang et al. 2017; Schootemeijer et al. 2018). We also note that the secondaries in these systems have higher masses than in HR 6819 ($(9 - 12) M_{\odot}$), allowing for higher initial masses for the primaries.

9.4 Summary and Discussion

We have shown that HR 6819 is a binary containing a rapidly rotating Be star and an undermassive B star with mass of order $0.5 M_{\odot}$. Radial velocity measurements from double-lined spectra establish that the two luminous components of the binary are orbiting one another (Figure 9.4). There is thus no need to invoke an unseen 3rd object to explain the orbital motion of the B star, as proposed by Rivinius et al. (2020). We propose that the B star is the core of a star with initial mass of order $5 M_{\odot}$ whose envelope was recently stripped by its companion near the tip of the red giant branch. In this scenario, the B star is currently in a short-lived (few $\times 10^5$ years) phase of contraction toward the extreme horizontal branch, where it will become a core helium burning sdOB star (e.g. Heber 2016). In this scenario, HR 6819 is a progenitor system for Be + sdOB binaries, of which at least six are known (e.g. Gies et al. 1998; Koubský et al. 2012; Peters et al. 2013, 2016; Wang et al. 2017; Dulaney et al. 2017; Chojnowski et al. 2018) and another dozen candidates have been identified (Wang et al. 2018). Recent accretion of the B star’s envelope would then explain the near-critical rotation of the Be star.

HR 6819 is quite similar to the binary LB-1 (Liu et al. 2019), which has recently been the subject of much discussion. In El-Badry & Quataert (2020), we speculated that the emission lines in LB-1 might come from a circumbinary disk around a BH and a stellar companion. But spectral disentangling by Shenar et al. (2020) has since revealed the presence of a Be star in the spectrum, similar to the case in HR 6819. The B star in LB-1 would then also be a recently stripped He star, with no BH in the system.

Further modeling will be required to pin down the mass and current evolutionary state of these systems. Our MESA calculations produce stripped stars with masses, temperatures, and surface gravities comparable to HR 6819 and LB-1, but there is not yet full agreement between the dynamically- and evolutionarily-inferred masses for LB-1. For both systems, *Gaia* DR3 astrometry (including binary astrometric model fits) and ground-based interferometry are expected to prove useful in further constraining the system parameters.

9.4.1 Implications for the formation of Be stars

We now consider what the existence of the HR 6819 system implies for the broader Be star population.

HR 6819 is bright enough ($G = 5.22$, $V = 5.36$) to be seen with the naked eye. Estimating the rate of Be + stripped star binaries based on HR 6819 is challenging both because the search that led to its discovery does not have a well-defined selection function and because Poisson errors are large when only one object is known. We nonetheless make a crude estimate of the frequency of HR 6819-like binaries below.

Roughly 15% of all field B stars are Be stars (Jaschek & Jaschek 1983; Zorec & Briot 1997). There are about 1500 B stars brighter than $G = 6$, corresponding to about 225 Be stars.³ Most stars this bright have been studied in some detail, so it seems a reasonable assumption that HR 6819 is the brightest object in its class, and that there are at most a few stars brighter than $G = 6$ in its current evolutionary state.

Once the stripped star contracts, HR 6819 will likely continue to appear as a Be star for a few $\times 10^7$ years – a few hundred times the lifetime of the stripped star phase – depending on the mass of the Be star. This implies that for every system like HR 6819, there should be a few hundred binaries consisting of a Be star and a faint sdOB or WD companion. Such binaries certainly exist, but they are challenging to detect because the sdOB typically contributes at tiny fraction of the light in the optical. 5 of the 6 confirmed Be + sdOB binaries are brighter than $G = 6$ and were discovered via UV spectra. Recently, Wang et al. (2018) identified another dozen candidate Be + sdOB systems using IUE spectra, all of which are brighter than $G = 7$. Schootemeijer et al. (2018) estimated that there should be 30-50 yet undiscovered Be + sdOB binaries for each confirmed system of a given magnitude, which corresponds to 150-250 Be + sdOB binaries brighter than $G = 6$. Raguzova (2001) estimated through population synthesis simulations that about 70% of Be stars formed through binary evolution should have faint white dwarf companions, while 20% should have sdOB companions. No white dwarf companions have yet been unambiguously detected.

Given our estimate that there are only a few hundred Be stars brighter than $G = 6$, the existence of HR 6819 is consistent with a scenario in which a majority of Be stars form through this evolutionary channel. This is an appealing possibility, because the origin of the Be phenomenon is still not fully understood. It is certainly tied to rapid rotation, but there is considerable debate about how Be stars were spun up in the first place (e.g. Rivinius et al. 2013). Accretion of mass from a companion is one possibility. There is, however, considerable uncertainty in the frequency of HR 6819-like binaries. Given $N = 1$ objects discovered among 225 Be stars, the 1-sigma Poisson confidence interval for the fraction of Be stars in this evolutionary phase is $(0.08 - 1.5)\%$. If we take the mean lifetime of a Be star

³We estimate the number of B stars brighter than $G = 6$ based on *Gaia* photometry. There are 1275 sources in *Gaia* DR2 with $-0.3 < G_{BP} - G_{RP} < 0.1$, a color range corresponding to $T_{\text{eff}} = (11 - 30)$ kK without extinction. We inflate this to 1500 to account for extinction (which is modest for most naked-eye stars) and incompleteness in *Gaia* DR2 for bright stars. We obtain a similar estimate based on the Hipparcos catalog. The Be star catalog from Jaschek & Egret (1982) contains 141 sources brighter than $V = 6$ but does not claim to be complete.

to be 5×10^7 years and the mean lifetime of the stripped star to be 2×10^5 years, the lower limit of 0.08% implies that at least 20% of Be stars went through a phase like HR 6819.

HR 6819 and LB-1 are currently the only known objects in their class. Given that LB-1 is more than 6 magnitudes fainter than HR 6819 and was discovered in a multi-epoch RV survey of only 3000 objects (without selection on color or spectral type), there are probably significantly more similar systems with magnitudes between HR 6819 and LB-1 that have yet to be discovered. We estimate that there are of order 10^4 Be stars brighter than $G = 12$.⁴ Given the rate estimate from HR 6819, this implies between 8 and 150 Be + bloated stripped star binaries brighter than $G = 12$, a large fraction of which should be discoverable to wide-field spectroscopic surveys such as SDSS-V. While both HR 6819 and LB-1 were only identified as candidate stripped stars after significant spectroscopic study, additional candidates could likely be identified from a single medium-resolution spectrum as objects exhibiting both the emission lines of a classical Be star and narrow absorption lines from the companion (indicative of low $v \sin i$). Normal Be stars viewed pole-on would be a contaminant to such a search, but these can usually be identified from the shape of their emission lines (Hummel & Vrancken 2000; Rivinius et al. 2013), and they are rare due to geometric effects.

With LB-1 and HR 6819 demoted, the list of detached stellar mass BH candidates remains short. The most secure detached BH candidates to date are likely those discovered in globular clusters with MUSE (Giesers et al. 2018, 2019), but the evolutionary history of these objects is likely quite different from that of BH binaries in the field. No luminous companion has been discovered to the giant–BH binary candidate reported by Thompson et al. (2019), but its status as a BH is not uncontroversial (van den Heuvel & Tauris 2020; Thompson et al. 2020). We note that the other two BH candidates on the field also contain Be stars (Casares et al. 2014; Khokhlov et al. 2018). These systems are not likely to contain low-mass stripped stars, because in them it is the Be star that is clearly RV variable. Be stars do not constitute a dominant fraction of all massive stars, and there is no strong bias toward Be stars in the selection function of multi-epoch RV surveys. If these objects do contain BHs, then their status as Be stars is likely to be a result of accretion from the BH progenitor.

In the final stages of this manuscript’s preparation, a similar analysis of HR 6819 was completed by Bodensteiner et al. (2020). Overall, there is excellent agreement between their results and ours. The only point of disagreement, the helium abundance of the B star, is discussed in Appendix 9.9.

Data availability

All the spectroscopic data is publicly available through the ESO archive. The TLUSTY/SYN-SPEC spectral models are available at <http://tlusty.oca.eu/>. MESA is publicly available at

⁴This estimate is also based on *Gaia* DR2. Because reddening is more severe for the $G < 12$ sample, we select candidate B stars as objects with $-0.3 < G_{BP} - G_{RP} < 0.3$ and $M_G < 1.5$; this yields 9.5×10^4 B star candidates, of which we again assume 15% are Be stars.

<http://mesa.sourceforge.net/>. Other data is available upon reasonable request of the corresponding author.

Acknowledgements

We are grateful to the referee, Philip Dufton, for a constructive report. We thank Thomas Rivinius, Dan Weisz, Todd Thompson, and Andrea Antoni for helpful discussions. KE acknowledges support from an NSF graduate research fellowship and a Hellman fellowship from UC Berkeley. We thank Geoff Tabin and In-Hei Hahn for their hospitality during the preparation of this manuscript.

9.5 Stellar parameters anchored to evolutionary tracks and distance

Figure 9.10 compares the atmospheric parameters of both components of HR 6819 to MIST evolutionary tracks. As discussed in Section 9.3.5, the temperature and surface gravity of the B star would imply a mass of at least $9 M_{\odot}$ if it were a normal star, but this is inconsistent with the measured system luminosity (Figure 9.5). The atmospheric parameters of the Be star imply a mass in the range of 5 and $8 M_{\odot}$. For a single star, this would correspond to an age range of 30 to 80 Myr, but we caution that mass transfer between the two components is likely to make these age estimates unreliable.

To estimate the mass and radius of the Be star, we construct a grid of evolutionary tracks with a mass spacing of $0.05 M_{\odot}$. We sample each track with 500 points spaced uniformly in time. We thus make no attempt to account for the different IMF probability or lifetime of each track (which would downweight higher-mass tracks), or the observational selection function (which would downweight lower-mass tracks). We then retain samples with probability proportional to the likelihood of their T_{eff} and $\log g$ given the measured constraints on the Be star’s atmospheric parameters, approximating the likelihood function as a two-sided Gaussian. This effectively selects points from the evolutionary tracks that are close to the Be star in $T_{\text{eff}} - \log g$ space. We use the resulting samples to estimate the mass, radius, and luminosity of the Be star, and to propagate forward the uncertainty in other parameters in Table 9.1 that depend on the Be star’s mass.

The system parameters listed in Table 9.1 are anchored to the mass of the Be star estimated from its T_{eff} and $\log g$. The constraints can be tightened somewhat if the distance, $d = (310 \pm 60)$ pc, is taken into account. To this end, we add a term to the likelihood function that compares the predicted V -band apparent magnitude of the evolutionary tracks to the measured $V = 5.36$, assuming a V -band flux ratio of $f_{\text{Be}}/f_{\text{tot}} = 0.52 \pm 0.05$ and extinction $A_V = 0.42$. The resulting constraints are listed in Table 9.3. They are consistent with the constraints in Table 9.1, but the uncertainties are somewhat smaller. At fixed T_{eff} , a constraint on luminosity translates directly to a constraint on radius. When combined with




Figure 9.10: Atmospheric parameters for the two components of HR 6819 compared to MIST evolutionary tracks for normal (i.e. non-stripped) stars. If the B star were a normal star, its T_{eff} and $\log g$ would imply a mass of at least $9 M_{\odot}$, roughly 20 times the mass we infer dynamically. For the Be star, the atmospheric parameters imply a mass between 5 and $8 M_{\odot}$.

the measured $\log g$, this translates to an improved constraint on mass.

9.5.1 System parameters independent of evolutionary tracks

An estimate of the stellar parameters of both components that is independent of evolutionary tracks can be calculated from the apparent magnitude, distance, and flux ratio. The V -band apparent magnitude of HR 6819 is 5.36. Given extinction $A_V = 0.42$ mag, total bolometric correction $BC = -1.5 \pm 0.15$,⁵ and distance $d = 369 \pm 34$ pc, the absolute bolometric magnitude of HR 6819 is $M_{\text{bol}} = -4.39_{-0.24}^{+0.26}$ mag. This corresponds to a bolometric luminosity

$$L_{\text{bol,tot}} = 10^{0.4(M_{\text{bol},\odot} - M_{\text{bol}})} L_{\odot} \quad (9.5)$$

$$= 4500_{-1000}^{+1100} L_{\odot}, \quad (9.6)$$

where $M_{\text{bol},\odot} = 4.74$.

Given the effective temperatures of the two components, their respective uncertainties, and $(47 \pm 5)\%$ flux contribution from the B star at $\lambda = 4000$ Å, the fraction of the total light contributed by the B star is $f_{\text{B,bol}} = 0.43_{-0.07}^{+0.08}$. The individual luminosities are then:

$$\log(L_{\text{bol,B}}/L_{\odot}) = 3.28_{-0.13}^{+0.13} \quad (9.7)$$

$$\log(L_{\text{bol,Be}}/L_{\odot}) = 3.40_{-0.12}^{+0.12}. \quad (9.8)$$

The radius of each component can then be calculated as $R = \sqrt{L_{\text{bol}}/(4\pi\sigma_{\text{SB}}T_{\text{eff}}^4)}$, where σ_{SB} is the Stefan Boltzmann constant. This yields

$$R_{\text{B}} = 5.7_{-0.9}^{+1.1} R_{\odot} \quad (9.9)$$

$$R_{\text{Be}} = 5.3_{-1.1}^{+1.9} R_{\odot}. \quad (9.10)$$

Finally, the mass of each component can be calculated as $M = gR^2/G$, yielding

$$M_{\text{B}} = 0.65_{-0.38}^{+0.93} M_{\odot} \quad (9.11)$$

$$M_{\text{Be}} = 6.9_{-4.0}^{+15.7} M_{\odot}. \quad (9.12)$$

Equations 9.7-9.12 are fully consistent with the constraints in Tables 9.1 and 9.3, which are anchored on the fit of the temperature and gravity of the Be star with evolutionary tracks. Equations 9.7-9.12 are also consistent with, but independent of, the dynamically measured mass ratio. We adopt the parameters anchored to evolutionary tracks as our fiducial values due to their tighter constraint on the component masses.

⁵For solar-metallicity MIST models, T_{eff} values of 16 kK and 18 kK correspond to bolometric corrections of -1.35 and -1.64 , with weak dependence on $\log g$.

Table 9.3: System parameters anchored to both distance and evolutionary tracks. Uncertainties are 1σ (middle 68%).

Be star mass	$M_{\text{Be}} [M_{\odot}]$	$7.0^{+1.0}_{-0.8}$
Be star radius	$R_{\text{Be}} [R_{\odot}]$	$5.3^{+0.9}_{-0.7}$
Be star bolometric luminosity	$\log(L_{\text{Be}}/L_{\odot})$	$3.45^{+0.14}_{-0.16}$
Be star fraction of critical rotation	$v_{\text{rot}}/v_{\text{crit}}$	$0.83^{+0.13}_{-0.11}$
Orbital inclination	i [deg]	$31.7^{+1.7}_{-1.5}$
Semi-major axis	$a [R_{\odot}]$	96 ± 4
B star mass	$M_{\text{B}} [M_{\odot}]$	$0.49^{+0.24}_{-0.22}$
Be star outer disk radius	$R_{\text{outer}} [R_{\odot}]$	33^{+4}_{-3}
Be star inner disk radius	$R_{\text{inner}} [R_{\odot}]$	$7.9^{+1.9}_{-1.8}$

9.6 Spectral disentangling and determination of K_{Be}

Spectral disentangling solves for the two single-star spectra that can best reproduce all the observed composite spectra *for an assumed set of relative RVs at each epoch*. If the RVs are not known a priori, they can be solved for in tandem with the single-star spectra by searching for the RVs that allow the disentangled single-star spectra to best reproduce all the observed composite spectra. In the case of HR 6819, the problem is simplified by the fact that the narrow lines of the B star allow a preliminary estimate of its RVs to be made directly from the composite spectra.

We optimize for the RVs in two steps. In the first step, we assume that the velocity of the B star at each epoch is the velocity predicted by the orbital solution of Rivinius et al. (2020). We then step through a grid of RV semi-amplitudes for the Be star, K_{Be} , calculating the optimal disentangled spectra for each value. The choice of $K_{\text{Be}} = 0 \text{ km s}^{-1}$ corresponds to a stationary Be star, as would be expected in a triple scenario.⁶ We choose the optimal K_{Be} and corresponding spectral decomposition as the value that minimizes the total χ^2 summed over all observed spectra.

With the disentangled spectra from the first step in hand, we use them as templates to fit the composite spectra for the RVs at each epoch, as described in Section 9.3.4. This revealed intrinsic scatter of $s \approx 4 \text{ km s}^{-1}$ for the RVs of the B star (see Section 9.3.4).

In the second step, we repeat the disentangling, again stepping through a grid of K_{Be} values. Rather than setting the RVs of the B star to the values predicted by the Rivinius et al. (2020) orbital solution, we now set them to the values measured in the previous step. The best-fit values of K_{Be} for a given spectral region do not change much between the first and second step, but the quality of the fit improves somewhat in the second step.

We perform spectral disentangling and determine the optimal K_{Be} for 10 different regions of the spectrum, each centered on an emission or absorption line with clear contributions

⁶Note that there is no need to try different center-of-mass velocities: because we do not use any template spectra, the absolute RV scale is arbitrary.

from both components. Treating different lines independently allows us to test whether different ions trace different kinematic components, and to quantify the uncertainty in the K_{Be} values we infer.

Figure 9.11 shows the total χ^2 as a function of K_{Be} for the 10 lines. For easy comparison, we subtract the minimum χ^2 for each line. The Fe II lines in the top panel are dominated by emission in the Be star spectrum (e.g. Figure 9.6). The He I lines (middle panel) are dominated by absorption. The Balmer lines⁷ contain both absorption and emission, with a larger emission contribution toward redder wavelengths (Figure 9.1).

The 10 spectral regions we consider have optimal K_{Be} ranging from 2 to 10 km s^{-1} . Of these, 7 have optimal values between 3 and 6 km s^{-1} . None of the lines have minimal χ^2 values associated with $K_{\text{Be}} = 0 \text{ km s}^{-1}$, as would be expected for a stationary Be star. The scatter in K_{Be} values inferred from different lines, including lines tracing the same ion, is larger than would be expected if the formal uncertainties in K_{Be} , as quantified by $\chi^2(K_{\text{Be}})$ for a single line, were reliable. This is most likely a consequence of small errors in continuum normalization as well as time-variability in the flux ratio, which violate the underlying assumptions of the spectral disentangling algorithm (see Hensberge et al. 2008).

We do not find systematic difference between lines from different ions, or emission vs. absorption lines. This suggests that the emission lines, which are presumed to originate in a disk around the Be star, move coherently with the star. Considering the 10 spectral regions shown in Figure 9.11, as well as the RVs obtained when fitting a synthetic spectral template (Figure 9.13), we adopt $K_{\text{Be}} = (4.5 \pm 2) \text{ km s}^{-1}$ as our fiducial value. We then fix $K_{\text{Be}} = 4.5 \text{ km s}^{-1}$ in disentangling the full spectrum.

The shape of the disentangled spectra depends very weakly on the value of K_{Be} . Choosing $K_{\text{Be}} = 0 \text{ km s}^{-1}$ makes both the emission and absorption lines of the disentangled Be star spectrum slightly broader (by of order 0.1 Å) than in our fiducial solution. Broader lines suppress the RV variability of the Be star, but this leads to a worse overall fit to the composite spectra. The inferred T_{eff} and $\log g$ of the two components are not sensitive to the adopted value of K_{Be} , at least over the range $0 \leq K_{\text{Be}}/(\text{km s}^{-1}) < 10$.

In addition to our primary spectral decomposition using CRES, we also experimented with using the Fourier-space spectral disentangling code `fd3` (also referred to as `FDBinary` in the literature; Ilijic et al. 2004; Ilijic 2017), which uses the algorithm described by Hadrava (1995) to separate composite spectra. The disentangled spectra produced by the two codes are very similar. Figure 9.12 compares the two codes' constraints on K_{Be} from the $\text{H}\beta$ line. Under idealized circumstances, the velocity-space and Fourier-space spectra disentangling algorithms are mathematically identical (e.g. Hensberge et al. 2008). The most significant practical difference in the implementations for our purposes is that `fd3` requires the RVs of both components to follow Keplerian orbits, while CRES allows them to be specified manually at each epoch. The `fd3` solution thus does not model the scatter in the B star's RVs, which is accounted for in the CRES solution.

⁷We do not fit the $\text{H}\alpha$ line, because the shape of the emission varies between epochs. That is, there is no two-component spectral decomposition for which a good fit of all the observed spectra can be obtained. Such variability is common for Be stars (e.g. Rivinius et al. 2013); it may or may not be related to the companion.




Figure 9.11: Total χ^2 versus RV semi-amplitude of the Be star assumed during spectral disentangling. Top, middle, and bottom panels respectively show Fe II lines, He I lines, and H Balmer lines.




Figure 9.12: Constraints on K_{Be} from the H β line. We compare results obtained using the velocity-space spectral disentangling code CRES (solid line) and the Fourier-space code fd3 (dotted line).




Figure 9.13: RVs for the Be star measured by fitting the He I absorption line at 4388 \AA with TLUSTY/SYNSPEC model spectra. Dashed line shows the prediction for a circular orbit with $K_{\text{Be}} = 4.5 \text{ km s}^{-1}$. The typical RV uncertainty ($\sigma_{\text{RV,Be}} \approx 5 \text{ km s}^{-1}$) is substantial compared to the amplitude of the RV variability, but there is clear phase-dependent variability.

We also experimented with fitting RVs for both components using synthetic spectral templates. Figure 9.13 shows the best-fit RVs of the Be star obtained by jointly fitting the RVs of both components at each epoch, using the best-fit TLUSTY/SYNSPEC spectra for both components and fitting the He I absorption line at 4387.9 Å. Because the synthetic spectra do not perfectly reproduce the observed spectra, the narrow lines in the Be star’s spectrum are washed out by rotation, and many of the Be star’s lines are significantly contaminated by emission, we find that we obtain more stable RVs when fitting a narrow wavelength range than when fitting the full spectrum simultaneously. The RVs measured this way show larger scatter than those measured from the disentangled spectra (Figure 9.4), but they still show clear periodic modulation, the amplitude of which is consistent with the amplitudes we measure from spectral disentangling. The RVs we measure for the B star are in good agreement with those measured in Section 9.3.4.

9.7 B star variability

As discussed in Section 9.3.4, the measured RVs of the B star do not perfectly follow the predictions of a Keplerian orbit, but exhibit additional scatter with amplitude $\sigma_{\text{RV}} \approx 4 \text{ km s}^{-1}$. Here we show that this variability is accompanied by variability in the B star’s line profiles.

The top panel of Figure 9.14 shows the line profiles of two Si III lines from three single-epoch spectra of HR 6819, which are chosen to showcase the range of profiles found in the dataset. Weak and narrow metal lines such as these are ideal for characterizing the B star’s variability because contamination from the Be star – besides dilution by its continuum – is negligible. All three epochs are shifted to rest frame. The depth and width of the lines vary visibly between epochs. We quantify the line width with $v_{\text{macro}} = \sqrt{2}\sigma$, where σ is the dispersion of the best-fit Gaussian broadening profile (see Appendix 9.8).

The three lower panels of Figure 9.14 show v_{macro} , equivalent width, and RV deviation from a Keplerian orbit, measured from the Si III 4568 line across all 51 epochs. The epoch-to-epoch variation in v_{macro} is of order 10%, with a maximum-to-minimum range of almost a factor of 2. The typical variation in equivalent width is somewhat smaller, of order 5%. The RV scatter is 4 km s^{-1} , consistent with our finding in Section 9.3.4. There is no clear trend with orbital phase. Like the photometric variability (see Rivinius et al. 2020), the line profile variability exhibits power on timescales from ≈ 0.4 days to several days, and is not dominated by a single period.

The observed line profile variations cannot be understood as a result of photometric variability of the Be star, which would change the flux ratio – and thus, line depth and equivalent width – but not v_{macro} . It thus seems natural to connect the B star’s line profile variations with the observed photometric variability of HR 6819. The period and amplitude of the photometric variability are consistent with a scenario in which they are due primarily to the B star (see below). We note, however, that photometric variability is also common among Be stars, so we cannot rule out the possibility that some of the observed photometric




Figure 9.14: Top panel shows three single-epoch spectra of HR 6189, highlighting two Si III lines for which the absorption is dominated by the B star. The depth and width of the lines varies appreciably between epochs. Legend indicates the best-fit Gaussian v_{macro} , which quantifies the width of the lines. Bottom three panels show v_{macro} , equivalent width, and RV deviation from a Keplerian orbit, all measured for the Si III 4568 line, in all 51 epochs. These quantities vary between epochs, but do not vary systematically with orbital phase.

variability is due to the Be component.

The period, photometric variability amplitude, RV scatter, temperature, and luminosity of the B star in HR 6819 are all typical of slowly pulsating B-type (SPB) stars (e.g. [Waelkens 1991](#); [Waelkens et al. 1998](#)). Indeed, the prototypical SPB star, 53 Persei, has $T_{\text{eff}} = 17 \text{ kK}$, $\log(L/L_{\odot}) = 2.9$, and exhibits pulsation-driven RV variations with peak-to-peak amplitude of about 15 km s^{-1} , and photometric peak-to-peak amplitude of order 0.2 mag, quite similar to the B star in HR 6819 ([Smith et al. 1984](#); [Chapellier et al. 1998](#)). SPB stars undergo non-radial g -mode pulsations, which are excited by the κ mechanism acting on the metal opacity bump for $T_{\text{eff}} = 10 - 25 \text{ kK}$ (e.g. [Dziembowski et al. 1993](#)). Line profile variations in SPB stars are primarily due to variations in the velocity field at the stellar photosphere; variations due to changes in T_{eff} are usually negligible ([Buta & Smith 1979](#); [Townsend 1997](#); [Aerts & Eyer 2000](#)).

Photometric and spectral variability violate the assumptions of the spectral disentangling algorithm. When variability is purely photometric, it is straightforward to account for a time-dependent luminosity ratio (e.g. [Ilijic 2004](#); [Hensberge et al. 2008](#)). In the presence of line profile variations, however, there is no unique spectral decomposition, so the disentangled spectra represent a time-average of the component spectra. We carried out tests with mock data to assess the effects of unmodeled line profile variability similar to that found in HR 6819 on spectral disentangling. Because the variability is not systematically phase-dependent, we find that it does not lead to systematic biases in K_{Be} or in the width or depths of the disentangled line profiles, but it does reduce the total goodness-of-fit.

9.8 Rotation vs. macroturbulent broadening

The metal lines of the B star component of HR 6819 have a typical FWHM of 50 km s^{-1} . This is relatively narrow for a B star, but it is larger than the expected thermal width given the B star's spectral type in the absence of rotation or macroturbulent broadening (FWHM $\approx 10 \text{ km s}^{-1}$).

Figure 9.15 shows line profiles of the Si III 4568 and 4575 Å lines in a single-epoch ($\phi \approx 0.55$) spectrum of HR 6819. These lines are ideal for characterizing the B star's line broadening because they are intrinsically narrow and spectral disentangling suggests that the contribution of the Be star to them is negligible. We compare the observed line profiles to the predictions of a TLUSTY/SYNSPEC model ($T_{\text{eff}} = 16 \text{ kK}$ and $\log g = 2.75$) with no rotational or macroturbulent broadening, and to line profiles broadened only by macroturbulence or rotation. The rotational broadening profile is taken from [Gray \(1992a\)](#), using a linear limb darkening coefficient $\epsilon = 0.5$. The macroturbulent broadening profile is modeled as a Gaussian with $\sigma = v_{\text{macro}}/\sqrt{2}$. Instrumental broadening has negligible effects on the predicted line profiles. Rotational broadening alone reproduces the observed line profiles poorly. Disagreement between the observed and modeled profiles is most obvious in the line wings, which for the rotational profiles extend only to $\pm v \sin i$. The observed line wings are less steep and extend farther from the line center. This is not a result of




Figure 9.15: Black line shows an observed spectrum of HR 6819, highlighting two Si III lines for which the absorption is dominated by the B star. Blue line shows a model spectrum with no rotational or macroturbulent broadening and $v_{\text{mic}} = 10 \text{ km s}^{-1}$. Cyan and red lines show fits to the line profiles assuming purely rotational and Gaussian broadening, respectively. The observed line profiles have broader wings than predicted for pure rotational broadening and are reasonably well-fit by Gaussian profiles. This indicates that non-rotational broadening contributes significantly to the observed line widths.

contamination from the Be star: given its $v \sin i \approx 180 \text{ km s}^{-1}$, any absorption from it is spread over a significantly wider wavelength range, $\Delta\lambda > 5 \text{ \AA}$.

Purely Gaussian broadening does reproduce the observed line profiles reasonably well. An isotropic Gaussian broadening profile is not well-motivated physically (e.g. Gray 1992a; Simón-Díaz & Herrero 2014), but it is a serviceable simplification. The physical basis of non-rotational broadening in B stars is not well understood. As it is modeled here, macroturbulence represents any broadening mechanisms besides rotation, not necessarily turbulence. Given the observed photometric variability and line profile variations in HR 6819, it is likely that pulsation-driven fluctuations in the surface velocity field of the B star contribute to the broadening. Such fluctuations are not expected to produce strictly Gaussian broadening profiles (e.g. Aerts & Eyer 2000; Aerts et al. 2014), but in practice we find that Gaussians provide a reasonably good fit to the line profiles of uncontaminated metal lines in HR 6819 at all epochs, though the best-fit values of v_{macro} change somewhat from epoch to epoch (Figure 9.14). The conclusion that rotation alone produces line profiles with too-steep wings also holds at all epochs.

Because non-rotational broadening likely dominates over rotational broadening, we cannot reliably estimate $v \sin i$ of the B star from its line profiles. By fitting the observed Si III 4568 line profiles with a combination of rotational and macroturbulent broadening (e.g. Ryans et al. 2002; Simón-Díaz & Herrero 2014), we can limit the projected rotation velocity to $v \sin i < 20 \text{ km s}^{-1}$. If the B star is tidally synchronized, the limits on its radius and the binary’s inclination (Table 9.1) would imply $v \sin i = 3.1_{-1.2}^{+1.8} \text{ km s}^{-1}$.

9.9 Helium enrichment

As discussed in Section 9.3.7, we find that the surface of the B star is enriched in helium. Evidence for this is shown in Figure 9.16, which compares the strong He lines in the disentangled spectrum to TLUSTY/SYNSPEC models with solar helium abundance (blue) as well as models with $n_{\text{He}}/n_{\text{H}} = 0.35$ (red). The latter is 3.5 times the solar value (0.55 dex enhancement) and corresponds to a $\approx 55\%$ helium mass fraction. The helium-enriched models clearly provide a better fit. We note that while a 55% helium mass fraction is considerably higher than expected for a normal hydrogen-burning star, it is lower than predicted by the MESA models we analyze in Section 9.3.8, which predict mass fractions of $\approx 80\%$.

Bodensteiner et al. (2020) also analyzed the spectra of HR 6819. Their inferred atmospheric parameters and conclusions about the nature of the system are broadly consistent with ours. However, they did not find evidence of helium enrichment. The different conclusions are driven primarily by different choice of atmospheric models: their abundance constraints are based on PoWR models presented in Hainich et al. (2019), while we use TLUSTY/SYNSPEC models. Comparing the models directly, we find that at fixed T_{eff} , $\log g$, and v_{mic} , and He abundance, the PoWR models predict significantly stronger He lines, such that solar-abundance PoWR models indeed provide a good fit to the He lines in HR 6819. We defer a more detailed analysis of the differences between the two atmospheric and




Figure 9.16: Strong helium lines of the B star. Black line shows the disentangled spectrum of the B star. Blue shows a TLUSTY/SYNSPEC model with $T_{\text{eff}} = 16$ kK, $\log g = 2.75$, $v_{\text{mic}} = 10$ km s $^{-1}$, and solar helium abundance, $n_{\text{He}}/n_{\text{H}} = 0.1$. Red line shows a model with the same atmospheric parameters but 0.55 dex higher helium abundance; i.e., $n_{\text{He}}/n_{\text{H}} = 0.35$. The helium-enhanced model provides a much better fit.

spectral models to future work.

Chapter 10

A million binaries from Gaia eDR3: sample selection and validation of Gaia parallax uncertainties

An earlier version of this article was previously published in El-Badry, K., Rix, H.-W., Heintz, T. M., 2021, arXiv:2101.05282, MNRAS, in press.

10.1 Abstract

We construct from *Gaia* eDR3 an extensive catalog of spatially resolved binary stars within ≈ 1 kpc of the Sun, with projected separations ranging from a few au to 1 pc. We estimate the probability that each pair is a chance alignment empirically, using the *Gaia* catalog itself to calculate the rate of chance alignments as a function of observables. The catalog contains 1.3 (1.1) million binaries with $>90\%$ ($>99\%$) probability of being bound, including 16,000 white dwarf – main sequence (WD+MS) binaries and 1,400 WD+WD binaries. We make the full catalog publicly available, as well as the queries and code to produce it. We then use this sample to calibrate the published *Gaia* DR3 parallax *uncertainties*, making use of the binary components' near-identical parallaxes. We show that these uncertainties are generally reliable for faint stars ($G \gtrsim 18$), but are underestimated significantly for brighter stars. The underestimates are generally $\leq 30\%$ for isolated sources with well-behaved astrometry, but are larger (up to $\sim 80\%$) for apparently well-behaved sources with a companion within $\lesssim 4$ arcsec, and much larger for sources with poor astrometric fits. We provide an empirical fitting function to inflate published σ_{ϖ} values for isolated sources. The public catalog offers wide ranging follow-up opportunities: from calibrating spectroscopic surveys, to precisely constraining ages of field stars, to the masses and the initial-final mass relation of white dwarfs, to dynamically probing the Galactic tidal field.

10.2 Introduction

About half of all solar-type stars are members of binary systems, and a majority of these are so widely separated that the two components never interact (e.g. [Moe & Di Stefano 2017b](#)). With orbital periods ranging from ~ 10 to $\sim 10^8$ years, most of these binaries can in some sense be viewed as clusters of two: the components formed from the same gas cloud and have orbited one another ever since. They thus have essentially the same age, initial composition, and distance, but generally different masses and occasionally different evolutionary phases. This makes wide binaries useful for calibrating stellar models as well as spectroscopic and astrometric surveys.

At angular separations greater than about one arcsecond, wide binaries are easily resolvable as two point sources. Distinguishing physically bound binary stars from chance alignments (“optical doubles”) has been a long-standing challenge for binary star astronomy. Indeed, the first systematic binary star catalog was constructed under the assumption that all close pairs were chance alignments ([Herschel 1782](#)), an assumption that was only shown to be incorrect two decades later ([Herschel 1803](#)).

For bright binaries at close angular separations, chance alignments can be excluded probabilistically. However, the contamination rate from chance alignments increases at wider separations and fainter magnitudes. Inclusion of proper motion data can aid the selection of genuine binaries, in which the two stars have nearly identical proper motions (e.g. [Luyten 1971, 1979b](#); [Salim & Gould 2003](#); [Chanamé & Gould 2004](#); [Dhital et al. 2010](#)). Many wide binary searches have specifically targeted high-proper motion stars, which have fewer phase-space neighbors that can be mistaken for binary companions. If available, parallaxes and radial velocities are also useful for distinguishing binaries from chance alignments (e.g. [Close et al. 1990](#); [Andrews et al. 2017](#)).

Prior to the *Gaia* mission ([Gaia Collaboration et al. 2016](#)), useful parallaxes for this purpose were only available for (relatively) small samples of nearby and bright stars. *Gaia* DR2 ([Gaia Collaboration et al. 2018a](#)) dramatically expanded the sample of stars with well-measured parallaxes and proper motions, enabling the construction of unprecedentedly pure and extensive wide binary samples. [El-Badry & Rix \(2018c, hereafter ER18\)](#) searched *Gaia* DR2 for pairs of stars within 200 pc of the Sun with parallaxes and proper motions consistent with being gravitationally bound, and projected separations of up to 50,000 AU (0.24 pc). Their catalog prioritized purity over completeness, and thus imposed relatively strict cuts on astrometric and photometric quality and signal-to-noise ratio (SNR). This resulted in a catalog of $\sim 55,000$ binaries with an estimated contamination rate of $\sim 0.1\%$.

Using the same basic strategy but less stringent cuts on astrometric SNR, [Tian et al. \(2020\)](#) extended the [ER18](#) binary search to larger distances ($d < 4$ kpc) and wider separations ($s < 1$ pc). This produced a substantially larger sample of $\sim 800,000$ binary candidates, but with a higher contamination rate: chance-alignments dominate their catalog at $s \gtrsim 20,000$ au, though higher-purity subsamples can be selected by imposing stricter cuts on astrometric SNR. Their catalog contained 325,000 binaries with $s < 20,000$ au; it is expected to be reasonably pure in this separation regime.

Another binary catalog was produced by [Hartman & Lépine \(2020\)](#), who combined *Gaia* DR2 astrometry with a catalog of high-proper motion stars not contained in *Gaia* DR2. In order to reduce contamination from chance alignments, they limited their search to binaries with proper motions larger than 40 mas yr^{-1} ; this translates to a distance limit of order 200 pc for typical stars in the Galactic disk, but to a larger search volume for stars on halo-like orbits with large tangential velocities. Their primary binary catalog contains $\sim 100,000$ binary candidates, with projected separations as large as 10 pc.

In this paper, we use *Gaia* eDR3 data ([Gaia Collaboration et al. 2020a](#)) to further expand the sample of known wide binaries. Compared to DR2, eDR3 astrometry is based on a $\sim 1.5 \times$ longer time baseline. This yields significant improvements in parallax and especially proper motion uncertainties. For example, the median uncertainties in `parallax` and `pmra` (one-dimensional proper motion) at $G = 18$ have respectively improved from 0.165 mas to 0.120 mas, and from $0.280 \text{ mas yr}^{-1}$ to $0.123 \text{ mas yr}^{-1}$ ([Lindegren et al. 2018, 2020b](#)). Improvements at the bright end are more significant, due to better handling of systematics; i.e., at $G = 13$, the median parallax uncertainty decreased from 0.029 mas to 0.015 mas. This improved astrometric precision allows us to distinguish bound binaries from chance alignments to larger distances and wider separations than was possible with DR2. Our approach is a compromise between the strategies adopted by [Tian et al. \(2020\)](#) and [ER18](#). Like [Tian et al. \(2020\)](#), we search out to wide separations and relatively low astrometric SNR, so that at wide separations and faint magnitudes, the full catalog is dominated by chance alignments. However, we also empirically estimate and assign each binary a probability that it is a chance alignment, making it straightforward to select pure subsets of the catalog.

The remainder of the paper is organized as follows. In [Section 10.3](#), we describe how our binary candidate sample is selected and cleaned. [Section 10.4](#) describes how we quantify the contamination rate from chance alignments and estimate the probability that each binary candidate is bound. [Section 10.5](#) details basic properties of the catalog and a cross-match with the LAMOST survey. In [Section 10.6](#), we use the binary sample to validate the *Gaia* eDR3 parallax uncertainties. We summarize and discuss our results in [Section 10.7](#). The public catalog is described in [Section 10.8](#). Details about the calculation of chance-alignment probabilities are provided in the Appendix.

10.3 Sample selection

To reduce contamination from chance alignments, we limit our sample to pairs in which both components have moderately precise astrometry. We retrieved from the *Gaia* archive all sources with parallaxes greater than 1 mas (corresponding to a nominal distance limit of 1 kpc, which in reality is blurred due to parallax errors), fractional parallax uncertainties less than 20%, absolute parallax uncertainties less than 2 mas, and non-missing G -band magnitudes. This was achieved with the following ADQL query:

```
select *  
from gaiaedr3.gaia_source
```

```

where parallax > 1
and parallax_over_error > 5
and parallax_error < 2
and phot_g_mean_mag is not null

```

The query returns a total of $N = 64,407,853$ sources, corresponding to $N(N+1)/2 \approx 2 \times 10^{15}$ possible pairs. Of these, we consider as initial binary candidates all pairs that satisfy the following:

- *Projected separation less than 1 parsec*: the angular separation between the two stars, θ , must satisfy

$$\frac{\theta}{\text{arcsec}} \leq 206.265 \times \frac{\varpi_1}{\text{mas}}, \quad (10.1)$$

where ϖ_1 is the parallax of the star with the brighter G magnitude. The maximum search radius of 1 pc (corresponding to an orbital period of $\sim 10^8$ years) is chosen because a vanishingly small number of bound binaries are expected to exist at separations wider than this, where the Galactic tidal field becomes comparable to the gravitational attraction of the two stars (e.g. [Binney & Tremaine 2008b](#)). The separation beyond which the Galactic tidal field dominates a binary's internal acceleration is called the Jacobi radius. In the Solar neighborhood, it is given by $r_J \approx 1.35 \text{ pc} \times (M_{\text{tot}}/M_{\odot})$, where M_{tot} is the total mass of the binary ([Jiang & Tremaine 2010](#)). At separations slightly below r_J , binaries are efficiently disrupted by gravitational perturbations from objects such as other stars and molecular clouds (e.g. [Weinberg et al. 1987](#)).

- *Parallaxes consistent within 3 (or 6) sigma*: the parallaxes of the two components, ϖ_1 and ϖ_2 , must satisfy

$$|\varpi_1 - \varpi_2| < b \sqrt{\sigma_{\varpi,1}^2 + \sigma_{\varpi,2}^2}, \quad (10.2)$$

where $\sigma_{\varpi,i}$ is the parallax uncertainty of the i -th component, and $b = 3$ for pairs with $\theta > 4$ arcsec, or $b = 6$ for pairs with $\theta < 4$ arcsec. The less stringent cut at $\theta < 4$ arcsec is adopted because the chance alignment rate there is low and parallax uncertainties are significantly underestimated at close angular separations (Section 10.6).

- *Proper motions consistent with a Keplerian orbit*: The two stars in a wide binary will have proper motions that are similar, but, due to orbital motion, not identical. We require

$$\Delta\mu \leq \Delta\mu_{\text{orbit}} + 2\sigma_{\Delta\mu}, \quad (10.3)$$

where $\Delta\mu$ is the observed scalar proper motion difference, $\sigma_{\Delta\mu}$ its uncertainty, and $\Delta\mu_{\text{orbit}}$ the maximum proper motion difference expected due to orbital motion. The first two quantities are calculated as

$$\Delta\mu = [(\mu_{\alpha,1}^* - \mu_{\alpha,2}^*)^2 + (\mu_{\delta,1} - \mu_{\delta,2})^2]^{1/2}, \quad (10.4)$$

$$\sigma_{\Delta\mu} = \frac{1}{\Delta\mu} \left[\left(\sigma_{\mu_{\alpha,1}^*}^2 + \sigma_{\mu_{\alpha,2}^*}^2 \right) \Delta\mu_{\alpha}^2 + \left(\sigma_{\mu_{\delta,1}}^2 + \sigma_{\mu_{\delta,2}}^2 \right) \Delta\mu_{\delta}^2 \right]^{1/2}, \quad (10.5)$$

where $\Delta\mu_\alpha^2 = (\mu_{\alpha,1}^* - \mu_{\alpha,2}^*)^2$ and $\Delta\mu_\delta^2 = (\mu_{\delta,1} - \mu_{\delta,2})^2$. Here $\mu_{\alpha,i}^* \equiv \mu_{\alpha,i} \cos \delta_i$, α and δ denote right ascension and declination, and μ_α and μ_δ , the proper motion in the right ascension and declination directions. Following ER18, we take

$$\Delta\mu_{\text{orbit}} = 0.44 \text{ mas yr}^{-1} \times \left(\frac{\varpi}{\text{mas}}\right)^{3/2} \left(\frac{\theta}{\text{arcsec}}\right)^{-1/2}, \quad (10.6)$$

which is the maximum proper motion difference expected for a circular orbit of total mass $5M_\odot$; it corresponds to a projected physical velocity difference

$$\Delta V_{\text{orb}} = 2.1 \text{ km s}^{-1} \times \left(\frac{s}{1000 \text{ au}}\right)^{-1/2}, \quad (10.7)$$

where $s = 1000 \text{ au} \times (\theta/\text{arcsec})(\varpi/\text{mas})^{-1}$ is the projected separation. This quantity is of course not equal to the full 3D separation, or to the semimajor axis, a . For randomly oriented orbits with a realistic eccentricity distribution, s and a usually agree within a factor of 2 (see ER18, their Appendix B). For our purposes, it is generally an acceptable approximation to assume $a \sim s$. We note that the cut on proper motion difference removes from the sample a significant fraction of unresolved hierarchical triples and higher-order multiples (see Section 10.5.4).

We apply the cuts on projected separation, parallax difference, and proper motion difference to all possible pairs. The projected separation cut reduces the list of possible binaries to 10^{10} ; additionally requiring consistent parallaxes and proper motions reduces it to 10^8 initial candidate pairs. A random sample of 1% of these is plotted in the top panel of Figure 10.1, and their projected separation distribution is shown in gray in Figure 10.2. A large majority of initial candidate pairs are chance alignments, not genuine binaries. This is evident both from the fact that many candidate pairs are in the Galactic bulge, LMC, and SMC, and from the fact that their separation distribution peaks at the widest separations (Figure 10.2), where true binaries are rare.

Stars in the bulge, LMC, and SMC would ideally be excluded by the requirement of $\varpi > 1$ and $\varpi/\sigma_\varpi > 5$ in the initial query, but a significant fraction of sources in crowded fields have spurious parallaxes (e.g. Fabricius et al. 2020). Most of these spurious background pairs can be excluded by imposing astrometric quality cuts – for example, we find that adding the requirement of `astrometric_sigma5d_max` < 1 to our initial query reduces the number of initial candidates by a factor of 10 while only removing a minority of genuine binaries¹. We opted against applying such cuts because they do remove some real binaries, and we find that chance alignments with spurious parallaxes can be efficiently filtered out by the cleaning described in Section 10.3.1.

In addition to chance alignments of background sources with spurious parallaxes, our initial selection also efficiently selects members of star clusters and moving groups, about

¹`astrometric_sigma5d_max` is the longest principal axis in the 5-dimensional error ellipsoid, in mas. A large value indicates that at least one of the astrometric parameters is poorly determined in the 5-parameter solution. More information on this and other *Gaia* flags can be found in the [Gaia eDR3 data model](#).

100 of which can be seen in the top panel of Figure 10.1. These are not spurious, in the sense that they really do contain many pairs of stars within our search volume that are close in phase space and in some cases mutually bound (e.g. Oh et al. 2017b). However, most of them are not binaries and will become unbound when the clusters dissolve.

10.3.1 Cleaning clusters, background pairs, and triples

We clean the list of initial binary candidates in several passes. First, beginning with all the sources returned by our initial ADQL query, we count for each source the number of phase-space neighbors that are brighter than $G = 18$ and consistent with the size and velocity dispersion of a typical cluster. We define neighboring sources as those that satisfy the following:

- Projected separation less than 5 pc; i.e., $\theta \leq 17.19 \text{ arcmin} \times (\varpi/\text{mas})$.
- Proper motions within 5 km s^{-1} ; this translates to a proper motion difference $\Delta\mu \leq 1.05 \text{ mas yr}^{-1} \times (\varpi/\text{mas})$, with a 2σ tolerance.
- Parallaxes consistent within 2 sigma; i.e., $\Delta\varpi \leq 2\sqrt{\sigma_{\varpi,1}^2 + \sigma_{\varpi,2}^2}$.

We remove from our binary candidate list all pairs in which either component has more than 30 neighbors as defined above. Only 6.5 million of the 64 million sources in the search sample have more than 30 neighbors, and inspection reveals that a majority of these are not in the search volume at all, but rather are spurious sources in the Galactic bulge, LMC and SMC. Removing candidates containing these sources shrinks the candidate list from 112,473,599 pairs to 2,881,543 and removes most of the obvious structure seen in the top panel of Figure 10.1.

Next, we remove all overlapping pairs. That is, if either component of a binary candidate is a member of another binary candidate, we remove both pairs. This removes genuine resolved triples, which are efficiently identified by our initial search (e.g. Perpinyà-Vallès et al. 2019) and are not rare (Tokovinin 2014b). It also removes some additional chance alignments. This cut shrinks the candidate sample from 2,881,543 to 1,918,362.

Finally, we search for members of small clusters or moving groups not removed in the first pass. Adopting the phase-space coordinates of the brighter component of each pair as representing the pair, we count the number of neighboring pairs for each candidate, defined using the same three criteria used when counting neighboring sources (without any magnitude cut). We reject all candidates that have more than 1 neighboring pair, shrinking the sample from 1,918,362 to 1,817,594.

It is important to note that some real binaries will be removed during the filtering of resolved triples, clusters, and moving groups. In regions of high stellar density, a distant tertiary candidate that is really a chance alignment can cause a genuine binary to be rejected as a triple. Similarly, some bound binaries *do* exist within clusters, and these will all be rejected. An upper limit of $\sim 15\%$ can be set on the fraction of true binaries lost during

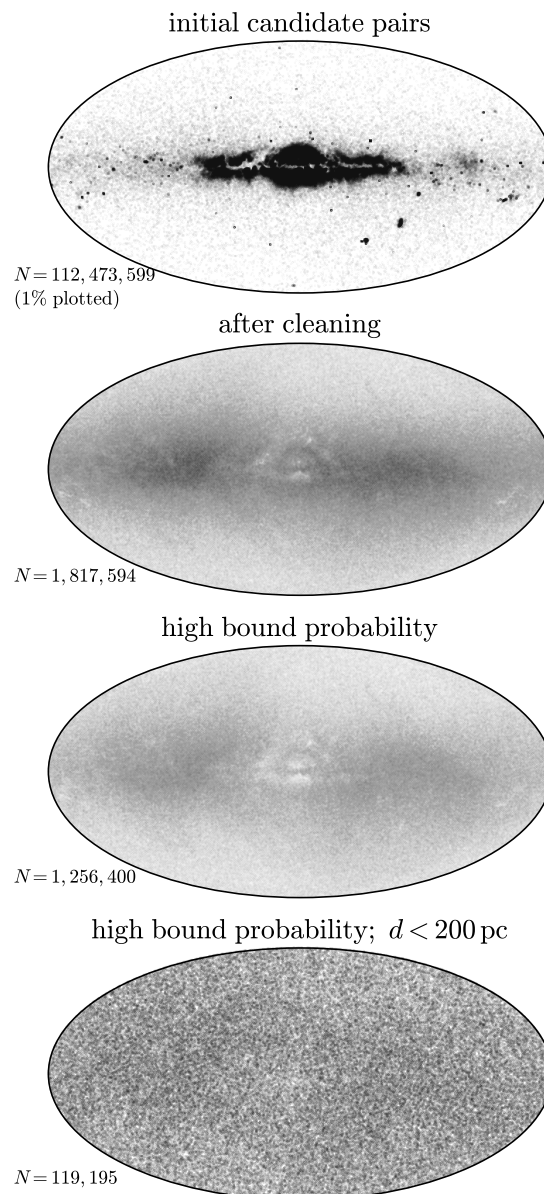


Figure 10.1: Sky distribution of binary candidates in Galactic coordinates. Top panel shows a random subset of all pairs with consistent parallaxes and proper motions. The vast majority of these are not genuine binaries. Many clusters are visible, as well as the Magellanic clouds and inner Galaxy. These are beyond the nominal 1 kpc search limit but enter the dataset due to spurious parallaxes. 2nd panel shows the sky distribution after clusters, moving groups, and resolved triples have been filtered (Section 10.3.1). This removes the majority of spurious background pairs. 3rd panel shows the pairs with $\mathcal{R} < 0.1$ (corresponding approximately to 90% bound probability; see Section 10.4.2). Some structure remains, primarily tracing dust. Bottom panel shows sources within 200 pc, which are distributed almost uniformly on the sky.




Figure 10.2: Separation distribution of binary candidates at three stages of the selection procedure, corresponding to the top three panels of Figure 10.1. Most initial candidates (gray) are chance alignments, which dominate at $s \gtrsim 3,000$ au. Cleaning resolved triples, clusters, and moving groups removes a large fraction of chance alignments (black), but they still dominate at $s \gtrsim 30,000$ au. The separation distribution of binaries with high bound probability ($\gtrsim 90\%$; red; see Section 10.4.2) falls off steeply at wide separations.

cleaning by comparing the gray and black histograms in Figure 10.2 at close separations; this is an upper limit because some of the pairs removed at close separations are genuinely members of resolved triples, moving groups, and clusters.

The binary candidate sample after removal of resolved triples and suspected cluster members is shown in the 2nd panel of Figure 10.1 and with the black histogram in Figure 10.2. This is the catalog published along with this work; it contains both high-confidence binaries and pairs that are very likely chance alignments. The contamination rate from chance alignments increases rapidly with separation (Section 10.4), while the true binary separation distribution decreases monotonically over the range of separations to which we are sensitive (e.g., ER18). The separation at which the number of binary candidates per dex of separation begins to increase ($\log(s/\text{au}) \approx 4.5$ for the black histogram in Figure 10.2) thus marks the separation at which chance alignments begin to dominate the full sample. That is, when considering the full catalog, a majority of binary candidates with $s \gtrsim 30,000$ au are chance alignments. It is, however, possible to select subsets of the catalog that are free of chance alignments out to much wider separations; see Section 10.4.

We define the “primary” and “secondary” components, denoted “1” and “2”, as the component with the brighter and fainter G magnitude, respectively. Both components are on the main sequence in a majority of binaries; in these cases, the “primary” is generally also the more massive component. For binaries containing white dwarfs, the secondary will often be more massive than the primary.

10.4 Chance alignments

The contamination rate from chance alignments depends on a variety of factors, including angular separation, parallax and proper motion, their respective uncertainties, and the local source density. We use two complementary approaches to constrain the chance alignment rate for different subsets of the catalog.

First, we repeat our binary selection procedure on the `GeDR3mock` catalog produced by Rybizki et al. (2020). This catalog is built on a realization of the Besançon model of the Milky Way (Robin et al. 2003) produced with `Galaxia` (Sharma et al. 2011). It contains no binaries, so by construction all binary candidates selected from the mock catalog are chance alignments. It does, however, contain a realistic population of open clusters, a variety of stellar populations, an approximation of the *Gaia* eDR3 selection function, and realistic astrometric uncertainties. We repeat the full binary selection process described in Section 10.3 on the mock catalog, including filtering of clusters and resolved triples. We remove pairs with angular separations $\theta < 0.5$ arcsec by hand, because the *Gaia* eDR3 sensitivity drops precipitously at closer separations (Fabricius et al. 2020), and this is not accounted for in the mock catalog. The separation distribution of candidates selected from the mock catalog is shown in the left panel of Figure 10.3.

Second, we produce an empirical chance alignment sample based on the actual *Gaia* eDR3 catalog, following the method introduced by Lépine & Bongiorno (2007b). Prior to selecting




Figure 10.3: Two methods for estimating the contamination rate from chance alignments. Left panel compares the separation distribution of binary candidates (black) to that of candidates selected from the eDR3 mock catalog (Rybizki et al. 2020; cyan), which contains no true binaries. Right panel compares the same binary candidates (black) to the separation distribution of candidates when stars are artificially shifted 0.5 degrees from their true positions when selecting candidate companions (orange; by construction, these are not true binaries). Both methods show that chance alignments dominate the full catalog at wide separations ($s \gtrsim 30,000$ au, where the separation distribution begins to increase), but true binaries dominate at closer separations.

potential binary companions to each star, we artificially shift it from its true position by ≈ 0.5 degrees, increasing its reported RA by $0.5 \sec(\delta)$ degrees. We then repeat the binary search, treating each star’s shifted coordinates as its true coordinates when searching for possible companions. This process avoids selecting real binaries, since stars are shifted away from their true companions, but preserves chance alignment statistics, because the source density within our 1 kpc search volume does not vary much on 0.5 degree scales. Copying and shifting the catalog effectively doubles the number of possible chance alignments, increasing the total number of pairs from $N(N+1)/2$ to $N(N+1)$, so we retain members of the shifted chance alignment catalog with 50% probability. We again remove pairs with $\theta < 0.5$ arcsec. The separation distribution of chance alignments produced in this way is shown in the right panel of Figure 10.3.

The two methods predict similar chance alignment rates. At wide separations ($\log(s/\text{au}) \gtrsim 4.5$ for the full catalog), the number of chance alignments is similar to the number of pairs in the binary candidates catalog. In this regime, most candidates are chance alignments. Not surprisingly, chance alignments begin to dominate at the separation where the binary candidate separation distribution begins to increase. The number of chance alignments per logarithmic separation interval increases steeply with separation: there are 100 times more chance alignments with $5 < \log(s/\text{au}) < 5.1$ than there are with $4 < \log(s/\text{au}) < 4.1$. This is a consequence of the larger available area for background stars to be found in at wider separations. At fixed distance, the area in which chance alignments can appear scales as $\sim 2\pi s ds$. At wide separations, the chance alignment distributions thus scale as $dN/d \log s \sim s^2$. This scaling does not hold exactly, in part due to small-scale clustering and in part due to the less stringent parallax consistency required at small angular separations, but it provides a good approximation to the chance alignment rate at wide separations ($\log(s/\text{au}) \gtrsim 4$).

We use the “shifted” chance alignment catalog in the rest of our analysis, because we find that it reproduces the separation distribution of binary candidates in the large-separation limit somewhat more reliably than the mock catalog.

10.4.1 Chance alignment rate for subsets of the catalog

Although chance alignments dominate the full sample at $s \gtrsim 30,000$ au, it is possible to select subsets of the catalog that are free of chance alignments out to wider separations. This is illustrated in Figure 10.4, which shows the separation distributions of various subsets of the binary candidate and the shifted chance alignment catalogs. As expected, the chance alignment rate at fixed separation is lower for samples with high galactic latitude, bright component stars, small fractional parallax errors, large space velocities, or large proper motions. Nevertheless, it is challenging to select any subset that remains pure beyond $s \sim 10^5$ au (~ 0.5 pc) without also dramatically reducing the sample size. We make the catalog of shifted chance alignments publicly available in order to facilitate estimation of the chance alignment rate in various subsamples of the catalog.

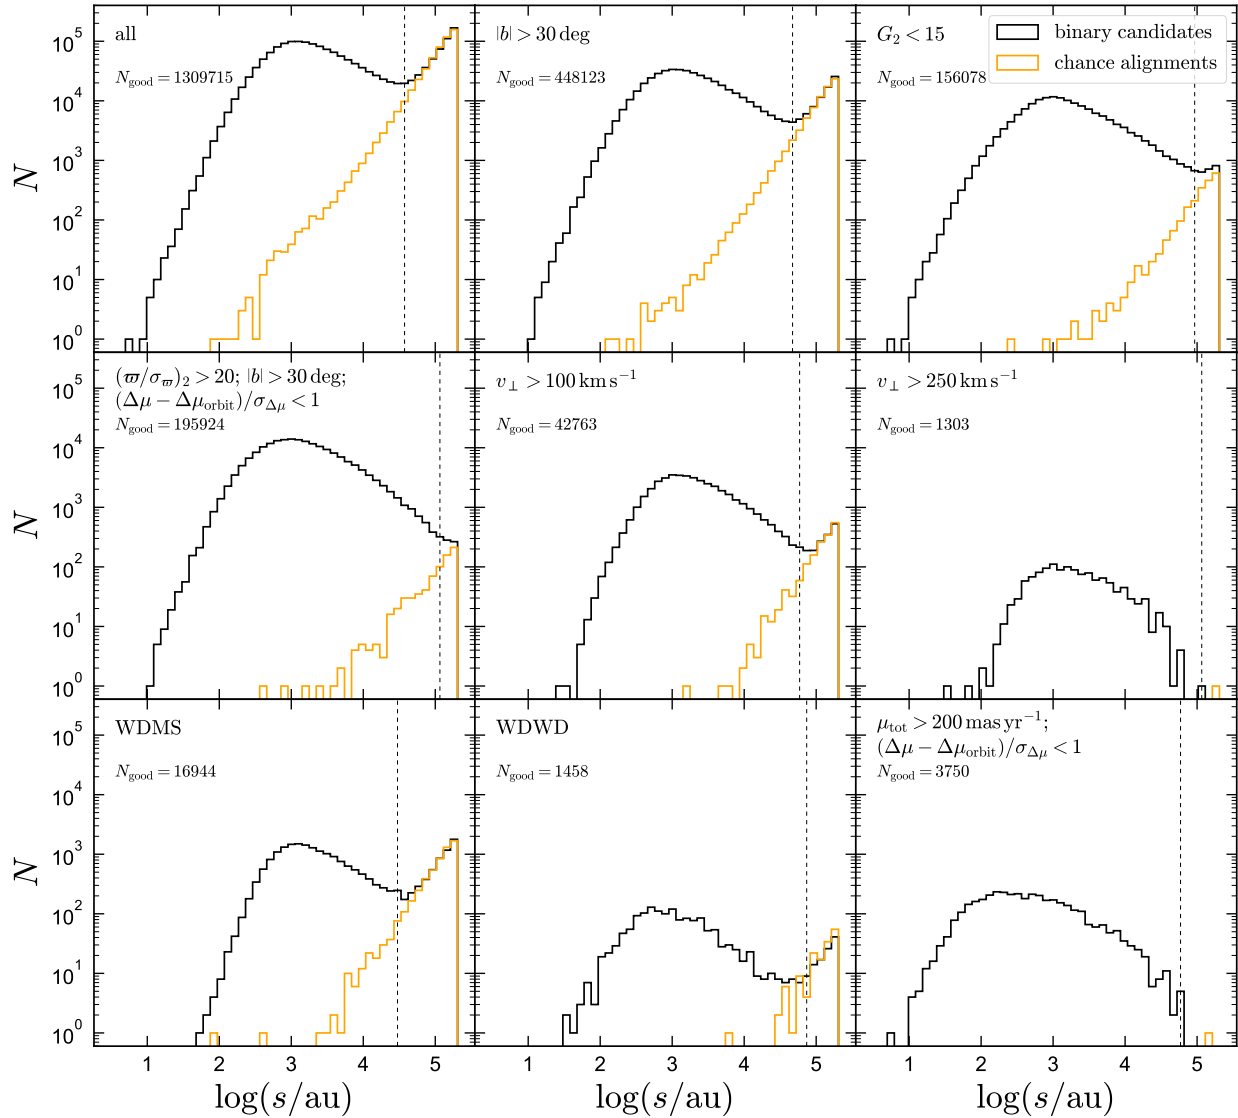


Figure 10.4: Contamination rate from chance alignments in various subsets of the catalog. Black histograms show binary candidates. Orange histograms show a catalog of chance alignments constructed by shifting stars 0.5 degrees when selecting candidate binaries. It contains no real binaries, but it has similar chance alignment statistics to the true binary candidates catalog. For the set of cuts illustrated in each panel, the vertical dashed line shows the widest separation at which there are more than $2\times$ the number of binary candidates as chance alignments. N_{good} is the number of binary candidates in that subset closer than this separation. Panels show the full catalog, high galactic latitudes, binaries with both components brighter than $G = 15$, precise parallaxes at high Galactic latitudes, high tangential velocities (e.g. halo-like orbits), WD+MS and WD+WD binaries, and high-proper motion pairs.

10.4.2 Estimating chance alignment probabilities

As illustrated in Figure 10.4, chance alignments and true binaries are found in different, but overlapping, regions of parameter space. We estimate the probability that a particular binary candidate is bound by comparing, at its location in parameter space, the local density of binary candidates and that of chance alignments from the shifted catalog. The process is described in detail in Appendix 10.9. The “densities” are evaluated in a seven-dimensional space using a Gaussian kernel density estimate (KDE). The dimensions are (1) angular separation, (2) distance, (3) parallax difference uncertainty, (4) local sky density, (5) tangential velocity, (6) parallax difference over error, and (7) proper motion difference over error. We re-scale these quantities to all have similar, order-unity dynamic range before fitting KDEs to both the binary candidate and the chance alignment distributions.

We denote the KDE-estimated density of chance alignments at a point \vec{x} in the 7-dimensional parameter space as $\mathcal{N}_{\text{chance align}}(\vec{x})$, and that of binary candidates as $\mathcal{N}_{\text{candidates}}(\vec{x})$. The latter quantity is expected to be the sum of the chance-alignment and true binary densities. We then calculate the ratio of these two quantities,

$$\mathcal{R}(\vec{x}) = \mathcal{N}_{\text{chance align}}(\vec{x}) / \mathcal{N}_{\text{candidates}}(\vec{x}). \quad (10.8)$$

This ratio approximately represents the probability that a binary candidate at position \vec{x} is a chance alignment, so selecting only candidates with small \mathcal{R} is an efficient method for eliminating chance alignments. \mathcal{R} is not strictly a probability – for example, it is not strictly less than one (Figure 10.5) – but it is a serviceable approximation for one. We calculate \mathcal{R} values for all members of the binary candidate and chance alignment catalogs.

Figure 10.5 (left) shows the distribution of \mathcal{R} values for both catalogs. There is a narrow population of binary candidates with \mathcal{R} near zero; these are objects that have a high probability of being bound. There is a second population of candidates with $\mathcal{R} \sim 1$; these objects are likely chance alignments. The separation distributions of binary candidates with \mathcal{R} below several thresholds are shown in the right panel of Figure 10.5. As expected, the vast majority of binary candidates with close separations ($\log(s/\text{au}) \lesssim 4$) have low \mathcal{R} values, indicating a high bound probability. At wider separations, the separation distribution of high-probability binaries falls off precipitously. We emphasize that this drop-off is steeper than that of the separation distribution of all binaries, because for a low \mathcal{R} threshold, more true binaries will be excluded from the high-confidence sample than chance alignments will be included in it.

We validate the use of \mathcal{R} as a proxy for the chance-alignment probability in Figure 10.6, which compares the *Gaia* DR2 radial velocities (Sartoretti et al. 2018; Katz et al. 2018) of the two components of candidates in which both stars have measured RVs. We only plot candidates in which both components have $\sigma_{\text{RV}} < 10 \text{ km s}^{-1}$ and the separation is wider than 50,000 au, where the full catalog is dominated by chance alignments. One expects the RVs of the two components to be similar for genuine wide binaries. For chance alignments, the RVs of the two components should be drawn from a broad distribution with width comparable to the local velocity dispersion of the Galactic disk, and thus they will generally be inconsistent.

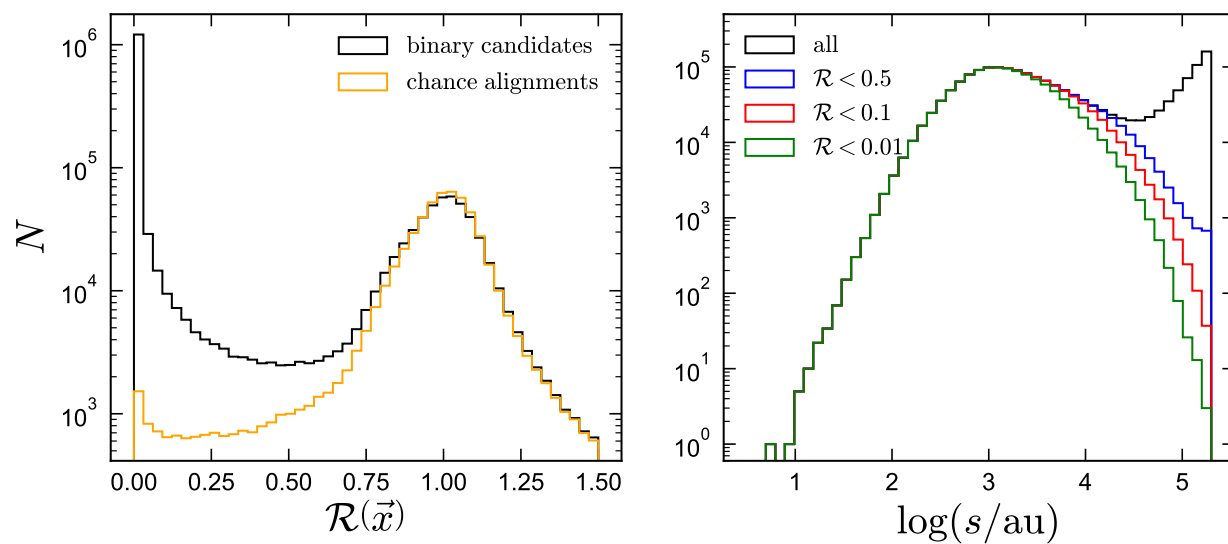


Figure 10.5: Left: ratio of the local density of chance alignments to binary candidates (Equation 10.8), for all pairs in the binary candidate sample (black) and the shifted chance alignment sample (orange). Likely chance alignments have $\mathcal{R} \sim 1$; objects with high probability of being bound have $\mathcal{R} \sim 0$. Right: separation distribution of binary candidates below different \mathcal{R} thresholds. The chance-alignment rate increases at wide separations, so the separation distribution of low- \mathcal{R} candidates falls off more steeply for progressively smaller \mathcal{R} thresholds.




Figure 10.6: *Gaia* RVS velocities of the components of binary candidates wider than 50,000 au. These are not used in constructing the catalog or calculating bound probabilities, but are useful for validation. Black points show all candidate binaries for which both components have $\sigma_{RV} < 10 \text{ km s}^{-1}$; red points show the subset of these with $\mathcal{R} < 0.1$, corresponding approximately to a 90% bound probability. These all fall close to the one-to-one line, suggesting that they are indeed bound.

For the full sample, there are indeed plenty of pairs with obviously inconsistent RVs, as chance alignments dominate at wide separations, even for bright pairs (e.g. Figure 10.4). Red points in Figure 10.6 show binaries with $\mathcal{R} < 0.1$. As expected, these pairs all have RVs close to the one-to-one line and are likely all bound. We stress that RVs are not used in creating the catalog or calculating \mathcal{R} values, so this result bolsters our confidence in the chance alignment ratios calculated from the shifted catalog.

We note that only bright stars ($G \lesssim 13.5$) had RVs published in *Gaia* DR2 (with no new RVs added in eDR3), so the fraction of chance alignments among binaries where both components have RVs is lower than in the full sample at the same separation. RVs for fainter stars are compared in Section 10.5.5, where we take RVs from the LAMOST survey.

In the rest of the paper, we define the “high bound probability” or “high confidence” subset of the catalog as the subset with $\mathcal{R} < 0.1$; this corresponds approximately to $> 90\%$ probability of being bound. This does *not* mean that 10% of the pairs in this subset are chance alignments: most candidates in it have $\mathcal{R} \ll 0.1$ (Figure 10.5). Interpreting \mathcal{R} as the probability that a given pair is a chance alignment, we estimate that 4,600 of the 1.26 million candidates with $\mathcal{R} < 0.1$ are chance alignments (0.4%). For $\mathcal{R} < 0.01$, the same fraction is 870 out of 1.15 million (0.08%). We make the full candidate catalog and \mathcal{R} values available, including pairs that are likely chance alignments.

10.5 Basic properties of the catalog

Basic properties of the binary catalog are shown in Figure 10.7 and listed in Table 10.1. Only the high-confidence pairs ($\mathcal{R} < 0.1$) are shown. Following ER18, we classify stars as white dwarfs (WDs) or main-sequence (MS) based on their position in the *Gaia* color-absolute magnitude diagram (CMD): defining $M_G = G + 5 \log(\varpi/100)$, we classify as WDs objects with $M_G > 3.25(G_{\text{BP}} - G_{\text{RP}}) + 9.625$; all other stars with measured $G_{\text{BP}} - G_{\text{RP}}$ colors are classified as MS stars. Under the ansatz that the two components have the same distance, we use the (usually more precise) parallax of the primary, (the brighter star) for both components when calculating M_G .

The adopted WD/MS boundary in the CMD is shown with a dashed line in Figures 10.7 and 10.8. The boundary is not entirely unambiguous – particularly for the WD+MS binaries, there are a few objects near the boundary that may be misclassified – but a majority of objects do fall clearly on the WD or MS sequences. We note that the “MS” classification serves only to exclude WDs. The CMDs in Figure 10.7 show that while most non-WD stars are indeed on the main sequence, the “MS” class also includes some giants, subgiants, pre-main sequence stars, and brown dwarfs. The number of binary candidates and high-confidence binaries in each class is summarized in Table 10.1.

The data acquisition window for BP/RP spectra is 2.1×3.5 arcsec wide, preventing colors from being measured for most close pairs (Arenou et al. 2018b). The majority of sources with a comparably bright companion within 2 arcsec thus do not have measured $G_{\text{BP}} - G_{\text{RP}}$ colors and cannot be classified as WD or MS stars based on *Gaia* data alone. Components

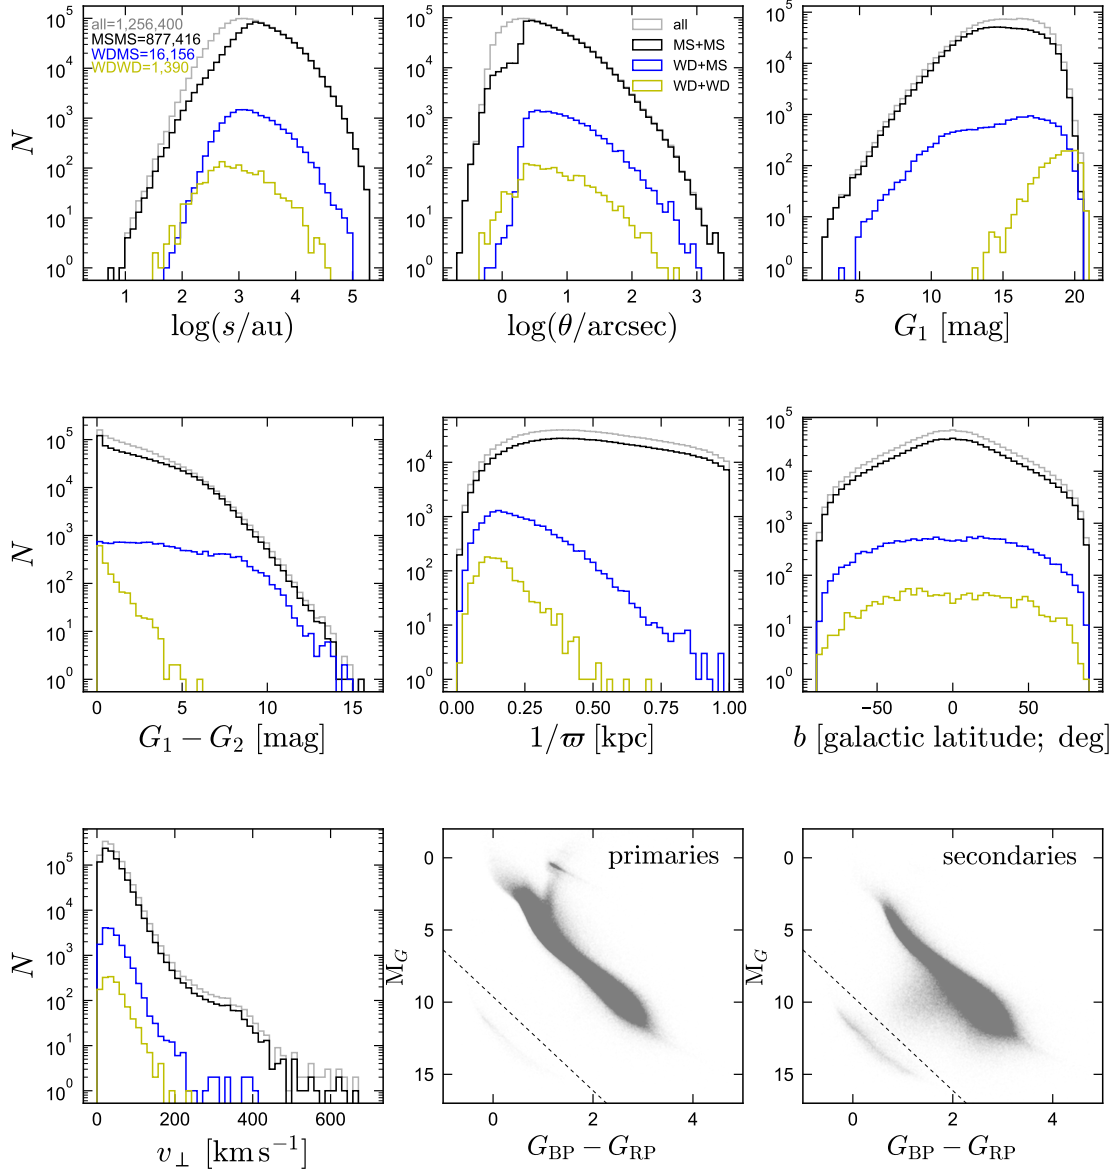


Figure 10.7: Basic properties of the high-bound probability binary sample: physical and angular separation, apparent magnitude of the primary, magnitude difference, distance, galactic latitude, plane-of-the-sky velocity, and color-magnitude diagrams. We designate stars as white dwarfs (WDs) or main sequence (MS; including giants) based on whether they fall below or above the dashed line in the bottom right panels. Black, blue, and yellow histograms show properties of MS+MS, WD+MS, and WD+WD binaries. We only classify binaries as MS+MS, WD+MS, or WD+WD if both components have a $G_{\text{BP}} - G_{\text{RP}}$ color. A majority of binaries with separation $\theta < 2$ arcsec do not have *Gaia* colors for a least one component and are therefore not classified; this accounts for most of the difference between the black and gray histograms (see also Table 10.1).

Classification	$N_{\text{candidates}}$	$N_{\mathcal{R}<0.1}$	Description
MSMS	1,412,903	877,416	both MS
WDMS	22,563	16,156	one WD, one MS
WDWD	1,565	1,390	both WD
MS??	378,877	360,180	one MS, one no colors
WD??	646	547	one WD, one no colors
????	1,040	711	both no colors
Total	1,817,594	1,256,400	

Table 10.1: Contents of the binary candidate catalog. All stars with measured $G_{\text{BP}} - G_{\text{RP}}$ colors are classified as “MS” or “WD” depending whether they fall above or below the dashed lines in Figure 10.8. $N_{\text{candidates}}$ is the number of candidate binaries with separations up to 1 pc (black histogram in Figure 10.2); $N_{\mathcal{R}<0.1}$ is the number with high bound probability (red histogram in Figure 10.2).

lacking a color measurement are denoted “??” in Table 10.1 and in the catalog. About 30% of all high-confidence binaries have one component with unknown color; 0.05% lack colors for both components. In 98% of cases where only one component has a color, it is the brighter component. In many cases where no *Gaia* color is available, colors from other surveys (e.g. Pan-STARRS) should be sufficient to distinguish WD and MS components.

The catalog contains 8.8×10^5 high-confidence MS+MS binaries, more than 16,000 high-confidence WD+MS binaries, and 1,390 high-confidence WD+WD binaries. Angular separations range from 0.2 arcsec to one degree. The peak of the angular separation distribution is at 1.2 arcsec. This is simply a result of the *Gaia* eDR3 angular resolution, since the intrinsic separation distribution falls off monotonically with increasing separation over all separations that are well-represented in the catalog (e.g. Duchêne & Kraus 2013). There are 271 pairs with separations between 0.2 and 0.4 arcsec, including 24 below 0.3 arcsec. Fabricius et al. (2020, their Figure 7) found some indication that the *Gaia* eDR3 catalog may contain spurious duplicated sources at separations below 0.4 arcsec (that is, single sources that were erroneously classified as two sources), so it is possible that a small fraction of the closest pairs in the catalog are spurious. We do not, however, find any increase in the angular separation distribution at close separations, as might be expected to arise from a population of duplicated sources.

The median magnitude of high-confidence primaries is $G = 15.2$, and that of secondaries is $G = 17.7$. Most WDs in the catalog are significantly fainter: the median magnitudes of primaries and secondaries in the WD+WD sample are 19.1 and 19.8. For WD+MS binaries, the same values are 15.5 for the primaries (in most cases, the MS star) and 19.4 for the secondaries. The median parallax of the full high-confidence sample is 2.05 mas ($1/\varpi \approx 485$ pc); the median distances for WD+WD and WD+MS binaries are 148 and 212 parsecs. Because of their closer distances, the WD+WD and WD+MS binaries are distributed roughly uniformly on the sky. The MS+MS sample, which extends to larger distances, bears clear imprint of the stratification of the Galactic disk. Most of the binaries

in the sample are part of the kinematic “disk” population, with a median tangential velocity $v_{\perp} = 4.74 \text{ km s}^{-1} \times (\mu_{\text{tot}}/\text{mas yr}^{-1})(\varpi/\text{mas})^{-1} \approx 35 \text{ km s}^{-1}$. There is also evidence of a kinematic “halo” population with $v_{\perp} \gtrsim 200 \text{ km s}^{-1}$ that contains a few thousand binaries.

The WD+MS and WD+WD binaries are shown separately on the CMD in Figure 10.8. On top of the WD+WD binaries, we plot WD cooling tracks for carbon-oxygen cores with hydrogen atmospheres (Holberg & Bergeron 2006; Kowalski & Saumon 2006; Tremblay et al. 2011; Bergeron et al. 2011; Bédard et al. 2020). WDs cool as they age, moving from the upper left to the lower right of the CMD. Cooling ages are indicated with triangular symbols along the tracks, which mark intervals of 1 Gyr. The faintest WDs in the catalog have implied cooling ages of about 10 Gyr. 31 WD+WD binaries have one component that falls below the $1.2M_{\odot}$ cooling track; i.e., with a photometrically-implied mass $M > 1.2M_{\odot}$. For WD+MS binaries, 75 WDs fall below the $1.2M_{\odot}$ cooling track. Hydrogen-atmosphere cooling tracks are not appropriate for WDs with non-DA spectral types, so spectroscopic classification must be obtained before masses and ages of individual WDs can be inferred with high fidelity.

The catalog also contains about 10,000 high-confidence binaries in which the primary is a giant (about half these giants are in the red clump), including about 130 giant-giant binaries. These are all quite bright (both components have $G < 11$). They, along with the WD+MS binaries and the $\approx 13,000$ binaries in which one component is a subgiant, can serve as useful calibrators for stellar ages.

Massive stars are not well-represented in the catalog; there are 351 high-confidence primaries with $G_{\text{BP}} - G_{\text{RP}} < 1$ and $M_G < 0$. This cut corresponds roughly to $M \gtrsim 3M_{\odot}$, though extinction complicates the mapping between M_G and mass. 75% of the high-confidence MS+MS binaries have primaries with M_G between 9.5 and 3.8, corresponding approximately to $0.4 < M_1/M_{\odot} < 1.3$. For secondaries, the 75% range is $M_G = 11.3 - 6.4$, corresponding to $0.2 < M_2/M_{\odot} < 0.8$. The catalog also contains about 80 binaries in which one component is likely a brown dwarf. We identify these on the CMD as objects with $M_G > 16.5$ and $G - G_{\text{RP}} > 1.3$; they are all within 80 pc of the Sun and therefore are also found in the binary catalog produced by Gaia Collaboration et al. (2020b).

A small but noticeable fraction of sources, particularly secondaries, are scattered below the main sequence in the CMD, between the WDs and MS stars (bottom right panel of Figure 10.7). The majority of these sources are separated from a brighter companion by only a few arcsec; the most likely explanation for their anomalous CMD position is thus that they are MS stars with contaminated BP/RP photometry. The majority of such sources can be filtered out using cuts on `bp_rp_excess_factor` and `phot_bp/rp_n_blended_transits` (e.g. Riello et al. 2020), but we refrain from employing such cuts since they also remove a significant fraction of sources with acceptable photometry and astrometry. Some sources below the main sequence may also be stars with spurious parallaxes or biased colors, and a few are likely real astrophysical sources, primarily cataclysmic variables and detached but unresolved WD+MS binaries (e.g. Abrahams et al. 2020; Belokurov et al. 2020).

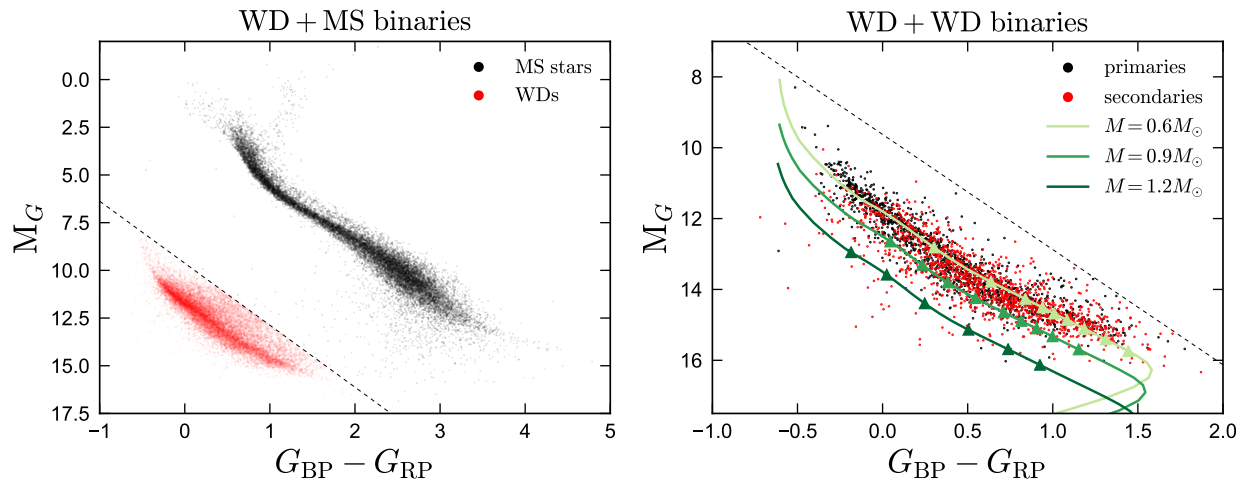


Figure 10.8: Color–absolute magnitude diagrams of high-confidence WD+MS and WD+WD binaries. Left: black points show MS components, which are usually the primaries; WD components are shown in red. Right: primary and secondary WDs are shown in black and red. We overplot cooling models for hydrogen-atmosphere WDs with masses between 0.6 and 1.2 M_{\odot} . Symbols mark 1 Gyr intervals of cooling age, with the first (leftmost) symbol at 1 Gyr.

10.5.1 Twin binaries and contrast sensitivity

An excess population of equal-brightness (and presumably, equal-mass) “twin” binaries is also found in the catalog. Its existence is most obvious in the distribution of magnitude difference, ΔG , as a function of separation, which is shown in Figure 10.9 and compared to chance alignments. Unlike the chance alignment catalogs constructed from the mock catalog and shifted catalogs (Section 10.4), these are selected in the same way as true binaries, but with the requirement that the parallaxes and proper motions of the two components be *inconsistent*. For easier visualization of the distribution of magnitude difference near $\Delta G = 0$, the sign of ΔG is randomized. The top panel illustrates the separation-dependent contrast sensitivity of the *Gaia* eDR3 catalog: at close angular separations, sources with significantly brighter companions are outshone. This leads to a contrast limit of $\Delta G \approx 4$ mag at $\theta = 1$ arcsec and $\Delta G \approx 7$ mag at $\theta = 2$ arcsec. The contrast limit at a given separation is not “sharp”, but is manifest as a smooth drop in sensitivity with increasing ΔG (e.g. Brandeker & Cataldi 2019b).² The contrast sensitivity is significantly improved in the binary catalog produced in this work compared to the one produced by ER18: that work required both components to have relatively uncontaminated BP/RP colors and thus contained basically no binaries closer than 2 arcsec.

The bottom panel of Figure 10.9 shows true binary candidates, in which the two compo-

²At very close separations ($\theta \lesssim 0.7$ arcsec; not shown in Figure 10.9), the *Gaia* eDR3 catalog contains only equal-brightness pairs (Lindgren et al. 2020b, their Figure 6). The contrast sensitivity is relatively smooth at $\theta > 1$ arcsec.




Figure 10.9: Magnitude difference between the two stars in chance alignments (top) and bound wide binaries (bottom) as a function of angular separation. The sign of ΔG is randomized. Chance alignments are selected following the same procedure used for real binaries, but with the requirement that the two stars’ parallaxes be *inconsistent* rather than consistent. The chance alignments illustrate the contrast sensitivity of the *Gaia* eDR3 catalog: at close separations, pairs with large magnitude difference are not detected. In the bottom panel, a narrow excess population of binaries with $\Delta G \approx 0$ is visible, which is absent in the top panel. This highlights the excess population of equal-mass “twin” binaries.




Figure 10.10: Comparison of the catalog produced in this work (black and red histograms) to the wide binary catalogs published by ER18, Tian et al. (2020), and Hartman & Lépine (2020), all of which are based on *Gaia* DR2 data. This catalog expands the one produced by Tian et al. (2020) by a factor of ~ 4 . The full catalog (black) has lower purity than the El-Badry & Rix (2018c) and Hartman & Lépine (2020) catalogs at wide separations. The subset with high bound probability (red) maintains high purity at wide separations, but does not purport to trace the intrinsic separation distribution there.

nents *do* have consistent parallaxes and proper motions. Unlike with the chance alignments, here there is an narrow excess population with $\Delta G \approx 0$. The extent and provenance of this population was studied by El-Badry et al. (2019); here we simply note that it is also clearly apparent in our catalog. Because the twin excess is most prominent at close physical separations, it is somewhat more obvious in our catalog, which extends to closer angular and physical separations.

10.5.2 Comparison to other catalogs

Figure 10.10 compares the distributions of projected physical separation, angular separation, and distance of the catalog produced here to other wide binary catalogs in the literature:

- The ER18 catalog (magenta) was produced from *Gaia* DR2 using a similar strategy to this work, but it was limited to binaries within 200 pc ($\varpi > 5$ mas) and used more stringent cuts on both photometric and astrometric quality and SNR. In order to maintain high purity, it only contains binary candidates with $s < 50,000$ au.
- Tian et al. (2020) expanded the ER18 search strategy to a larger volume (4 kpc) and used less stringent quality cuts, while still using *Gaia* DR2 data. Unlike this work, which uses a cut of $\varpi/\sigma_\varpi > 5$ for both components, they required $\varpi/\sigma_\varpi > 20$ for the primary and $\varpi/\sigma_\varpi > 2$ for the secondary. This results in somewhat different contamination properties and completeness. They also searched out to $s = 1$ pc. In

addition to their full catalog of all candidates, [Tian et al. \(2020\)](#) published 3 smaller catalogs with high purity, which are not shown in [Figure 10.10](#).

- [Hartman & Lépine \(2020\)](#) did not use a strict distance cut, but limited their search to high-proper-motion pairs with $\mu > 40 \text{ mas yr}^{-1}$. This preferentially selects nearby stars, since proper motion is inversely proportional to distance at fixed transverse velocity. For a “typical” tangential velocity of $v_{\perp} = 35 \text{ km s}^{-1}$, their proper motion cut corresponds to $d < 185 \text{ pc}$. However, stars on halo-like orbits with larger tangential velocities are included to larger distances; e.g., a binary with $v_{\perp} = 200 \text{ km s}^{-1}$ will have $\mu > 40 \text{ mas yr}^{-1}$ out to a distance of 1.05 kpc. Rather than employing strict cuts on parallax and proper motion consistency, [Hartman & Lépine \(2020\)](#) used empirical estimates of the chance alignment rate as a function of position and proper motion difference from a shifted catalog (similar to our approach in [Section 10.4](#)) to distinguish true binaries from chance alignments. Their approach has the advantage of not requiring specific cuts in parallax or proper motion difference, which are always somewhat arbitrary. A disadvantage is that it does not account for the heteroskedasticity of parallax and proper motion uncertainties – i.e., the chance alignment probability is higher for pairs with large astrometric uncertainties, and this is not accounted for in their analysis.

One difference between the catalog produced in this work and the other catalogs is obvious in [Figure 10.10](#): our sample extends to smaller angular separations, and thus also physical separations. This is partly a result of the improved angular resolution of *Gaia* eDR3 (e.g., [Fabricius et al. 2020](#)) but is primarily due to a change in search strategy. [ER18](#) and [Tian et al. \(2020\)](#) required both components of candidate binaries to have $G_{\text{BP}} - G_{\text{RP}}$ colors and to pass photometric quality cuts related to the `bp_rp_excess_factor` reported in *Gaia* DR2 ([Evans et al. 2018](#)). This set a soft resolution limit of $\sim 2 \text{ arcsec}$, with a wider effective limit for pairs with large brightness contrast. We do not require colors or employ a photometric quality cut in this work; this adds an additional $\sim 400,000$ binary candidates to the sample that would be excluded if we did ([Table 10.1](#)).

Both the catalog produced in this work and the one from [Tian et al. \(2020\)](#) become dominated by chance alignments at $s \gtrsim 30,000 \text{ au}$. The “high bound probability” subset of our catalog does not, but its separation distribution falls off steeply at wide separations. This decline at wide separations is steeper than that of the intrinsic separation distribution, since a decreasing fraction of binaries at wide separations can be identified as bound with high confidence. This can be seen in comparing the red separation distribution to the magenta one from [ER18](#), which tracks the intrinsic separation distribution over $5000 < s/\text{au} < 50,000$. The separation distribution from [Hartman & Lépine \(2020\)](#) has a similar logarithmic slope to the one from [ER18](#) in this separation range and likely tracks the intrinsic separation distribution out to wider separations. At the very widest separations represented in that catalog (5+ pc, exceeding the local Jacobi radius and corresponding to an orbital period of about a Gyr), it is unlikely that pairs are actually bound. This may reflect the fact that the search strategy employed by [Hartman & Lépine \(2020\)](#) is sensitive to any pairs that

are closer in phase space than chance alignments from the shifted catalog, without explicit consideration of the expected orbital velocities. That is, their search does not distinguish between bound binaries, moving groups, or stellar streams.

In terms of absolute numbers, the catalog represents a factor of ~ 4 increase in the number of high-confidence binaries over the one from [Tian et al. \(2020\)](#). To our knowledge, it is the largest published catalog of high-confidence binaries of any type. The sample could likely be expanded by a further factor of a few by loosening the distance and parallax uncertainty limits, or dropping the parallax cut entirely, while focusing on close angular separations (e.g. [Dhital et al. 2015](#)). However, the cuts we use in this paper provide a reasonable compromise between sample size, purity, and data quality.

A wide binary catalog based on *Gaia* eDR3 was also produced by [Gaia Collaboration et al. \(2020b\)](#), which contains pairs within 100 pc. It is not shown in [Figure 10.10](#), but we find that within 100 pc, it is almost identical to ours.

10.5.3 Space density

[Figure 10.11](#) compares the separation distributions of binaries in different (cumulative) distance bins. We scale these by the effective stellar volume corresponding to the distance cut, i.e., by a factor proportional to the total number of stars expected in each distance sample. Were it not for the finite scale height of the Galactic disk, this factor would simply be the search volume $V = 4\pi d_{\max}^3/3$, where d_{\max} is the distance limit. We approximate the total stellar density within our search volume as a plane-parallel exponential distribution with the Sun at the midplane and a scale height $h_z = 300$ pc ([Jurić et al. 2008b](#)). We then define an effective volume \tilde{V} , which represents the number of stars in a sphere of radius d_{\max} divided by the stellar density at the disk midplane:

$$\tilde{V} = 2\pi \int_0^{d_{\max}} e^{-z/h_z} (d_{\max}^2 - z^2) dz \quad (10.9)$$

$$= 2\pi \left[2h_z^3 (e^{-d_{\max}/h_z} - 1) + 2h_z^2 d_{\max} e^{-d_{\max}/h_z} + h_z d_{\max}^2 \right]. \quad (10.10)$$

As expected, this expression asymptotes to $4\pi d_{\max}^3/3$ in the limit of $d_{\max} \ll h_z$. In [Figure 10.11](#), the separation distributions for each value of d_{\max} are divided by the appropriate value of \tilde{V} .

Dashed vertical lines in [Figure 10.11](#) mark a separation of $s_{\text{res limit}} = (2 \text{ arcsec}) \times d_{\max}$. At separations $s < s_{\text{res limit}}$, incompleteness due to the *Gaia* eDR3 angular resolution starts to become severe. The figure shows that for $d_{\max} \lesssim 200$ pc, incompleteness is due primarily to the angular resolution limit: at $s > s_{\text{res limit}}$, the separation distributions in different distance bins overlap. However, for $d_{\max} = 500$ pc or 1000 pc, the catalog contains fewer binaries per effective volume than at closer distances, even at $s > s_{\text{res limit}}$. This reflects the fact that at sufficiently large distances, some binaries will have components that are too faint to pass the `parallax_over_error > 5` limit, or to be detected at all. At $d = 200$ pc, `parallax_over_error > 5` implies $\sigma_w < 1$ mas. This is satisfied by most sources with




Figure 10.11: Midplane number density of binaries in different cumulative distance bins. Vertical lines mark $\theta = 2$ arcsec at the distance limit of each bin, roughly the angular separation at which incompleteness due to blending becomes significant. At all distances, the turnover at close separations is driven primarily by the resolution limit, but the 15 and 25 pc samples show a relatively flat separation distribution below a few hundred au. Beyond 200 pc, there is significant incompleteness due to stars being too faint to pass our fractional parallax error threshold – i.e., the calculated number densities are lower than in the nearby samples even at wide separations, where blending should not be important.

$G \lesssim 20.5$ (Lindegren et al. 2020b), corresponding to $M_G = 14.0$, near the bottom of the main sequence (e.g. Figure 10.8). That is, the sample is expected to be almost complete for $d_{\max} = 200$ pc, except for crowding/blending effects at close separations. On the other hand, at $d = 1$ kpc, `parallax_over_error` > 5 implies $\sigma_{\varpi} < 0.2$ mas, $G \lesssim 18.8$, and $M_G \lesssim 8.8$, meaning that most of the lower main sequence will be excluded.

All of the separation distributions in Figure 10.11 increase toward smaller separations at $s > s_{\text{res limit}}$, but the distributions for $d_{\max} = 15$ pc and $d_{\max} = 25$ pc are similar and do appear to flatten above the resolution limit, at $s \sim 30$ au. For a typical binary in the catalog with total mass $1 M_{\odot}$, this corresponds to a period of order 200 years, which is indeed near the peak of the approximately lognormal separation distribution for solar-type binaries (Raghavan et al. 2010). We caution that effects of astrometric acceleration also become important in this regime (Section 10.7.3), potentially leading to spurious parallaxes and preventing pairs from being recognized as binaries by our search.

Integrating over separation, the distributions in Figure 10.11 imply a total space density of (0.006 ± 0.001) wide binaries with $s > 30$ au per cubic parsec in the solar neighborhood. For context, the space density of all unresolved *Gaia* eDR3 sources in the solar neighborhood is 0.07 pc^{-3} (Gaia Collaboration et al. 2020b), about 10 times higher. When all members of multiple systems are counted individually, the total stellar space density in the solar neighborhood is about 0.10 pc^{-3} (e.g. Winters et al. 2020).

10.5.4 Orbital velocities from proper motion differences

Precise parallaxes and proper motions make it possible to estimate orbital velocities (projected onto the plane of the sky) from the proper motion difference between the two stars. The plane-of-the-sky velocity difference ΔV can be calculated as

$$\Delta V = 4.74 \text{ km s}^{-1} \times \left(\frac{\Delta\mu}{\text{mas yr}^{-1}} \right) \left(\frac{\varpi}{\text{mas}} \right)^{-1}. \quad (10.11)$$

Here $\Delta\mu$ is the scalar proper motion difference (Equation 10.4), and ϖ is the parallax of the binary, for which we take the parallax of the brighter component. The corresponding uncertainty is

$$\sigma_{\Delta V} = 4.74 \text{ km s}^{-1} \sqrt{\frac{(\Delta\mu)^2}{\varpi^4} \sigma_{\varpi}^2 + \frac{\sigma_{\Delta\mu}^2}{\varpi^2}}, \quad (10.12)$$

with $\sigma_{\Delta\mu}$ calculated from Equation 10.5, $\Delta\mu$ and $\sigma_{\Delta\mu}$ in mas yr^{-1} , and ϖ and σ_{ϖ} in mas. We implicitly assume here that the two stars have the same parallax. Equation 10.12 is almost always dominated by the first term under the radical; i.e., parallax errors dominate over proper motion errors. The median value of $\sigma_{\Delta V}$ for all high-confidence binaries in the catalog is 0.33 km s^{-1} ; 195,601 have $\sigma_{\Delta V} < 0.1 \text{ km s}^{-1}$.

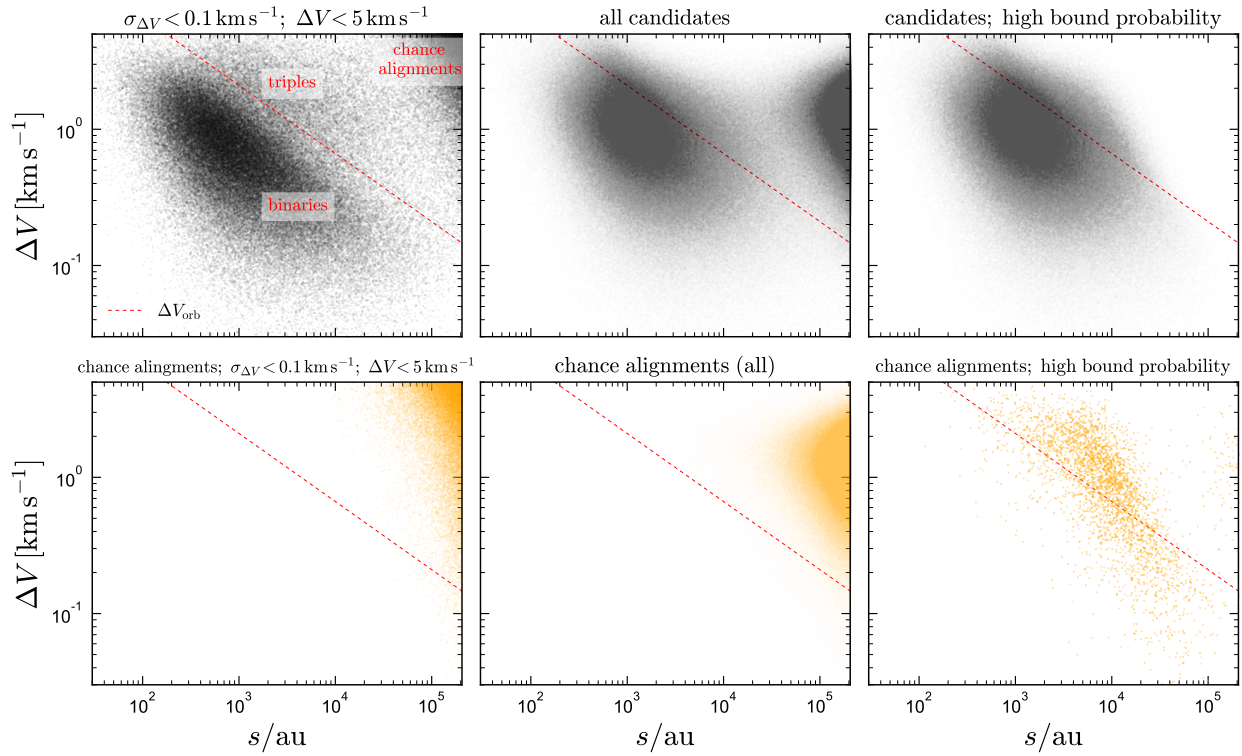


Figure 10.12: Upper left: plane-of-the sky velocity difference between the components of wide binary candidates selected with $\Delta V < 5 \text{ km s}^{-1}$ and $\sigma_{\Delta V} < 0.1 \text{ km s}^{-1}$. Dashed red line shows the maximum expected velocity difference for a binary with total mass $5M_{\odot}$ (Equation 10.7). Binary candidates in our catalog (top center) are required to be consistent, within 2σ , with falling below this line, but pairs with large proper motion uncertainties can scatter well above it. The clouds of points at large separation and ΔV are chance alignments; bottom panels show chance alignments from the shifted catalogs for the same selection. Points at large separation and ΔV are excluded by the requirement of high bound probability ($\mathcal{R} < 0.1$; upper right). In the upper left panel, there is a population of triples and higher-order multiples above the dashed line, with no corresponding population in the shifted catalog (lower left). These have increased ΔV due to the gravitational effects of an unresolved close companion. They are generally gravitationally bound, but are excluded from our primary catalog.

Figure 10.12 explores the ΔV values of binary candidates in the catalog, and the effects of the ΔV cuts we employ on its purity and completeness. The upper left panel shows all pairs from our initial query that have $s < 1$ pc, consistent parallaxes, $\Delta V < 5 \text{ km s}^{-1}$, and $\sigma_{\Delta V} < 0.1$. Note that the cut of $\Delta V < 5 \text{ km s}^{-1}$ is generally less strict than the one we adopt in constructing our primary catalog (Equation 10.3), which is equivalent to $\Delta V < \Delta V_{\text{orb}} + 2\sigma_{\Delta V}$, with ΔV_{orb} given by Equation 10.7.

In the upper left panel, there is a clear ridgeline of binaries with $\Delta V \sim s^{-1/2}$, as expected from Kepler’s laws. This population largely falls below ΔV_{orb} (dashed red line) because most of the binaries in the catalog have total masses less than $5M_{\odot}$. There is not, however, a sharp drop-off at $\Delta V > \Delta V_{\text{orb}}$. The population with $\Delta V > \Delta V_{\text{orb}}$ likely consists primarily of triples and higher-order multiples, in which the plane-of-the-sky velocity of one component is affected by a close, unresolved companion (e.g. Clarke 2020; Belokurov et al. 2020). It is also possible that some of these pairs are moving groups that are not actually bound but remain close in phase space (e.g. Pittordis & Sutherland 2019; Coronado et al. 2020); however, we find that the population exists even among binaries on halo-like orbits, favoring multiplicity as the primary explanation for it. We also find that most pairs with $\Delta V > \Delta V_{\text{orb}}$ have unusually large `ruwe`³ for at least one component (Figure 10.13), suggesting that these components are unresolved binaries (e.g. Belokurov et al. 2020).

A large fraction of these unresolved triples and higher-order multiples are excluded from our actual catalog by the requirement of $\Delta V < \Delta V_{\text{orb}} + 2\sigma_{\Delta V}$ rather than e.g., $\Delta V < 5 \text{ km s}^{-1}$; this should be kept in mind when using the catalog for applications involving higher-order multiplicity. We do not use a constant ΔV cut (e.g. $\Delta V < 5 \text{ km s}^{-1}$) in constructing the full catalog because this would result in a much higher contamination rate from chance alignments. For $5,000 < s/\text{au} < 10,000$, 22% of pairs with $\Delta V < 5 \text{ km s}^{-1}$ and $\sigma_{\Delta V} < 0.1 \text{ km s}^{-1}$ are excluded from the catalog by Equation 10.3. If we assume that large ΔV values are due mainly to subsystems, this implies that about 12% of wide binary components in this separation range have an unresolved subsystem that imparts a large enough photocenter perturbation to significantly increase ΔV . This corresponds to a subsystem separation range of $\sim 1 - 100$ au, since photocenter wobbles will average out over the 34-month *Gaia* eDR3 baseline at closer separations, and the subsystem-induced ΔV will be small at wider separations. About 25% of wide binary components have a subsystem in this separation range (Tokovinin & Smekhov 2002b; Tokovinin et al. 2010), so it is quite plausible that the large- ΔV pairs are mostly unresolved triples and higher-order multiples.

The clouds of points at large separation and ΔV in Figure 10.12 are chance alignments. These are primarily pairs with large $\sigma_{\Delta V}$; otherwise they would be excluded by the requirement of $\Delta V < \Delta V_{\text{orb}} + 2\sigma_{\Delta V}$. The bottom panels show the distribution of chance alignments from the shifted catalog (Section 10.4). These are distributed similarly to the large-separation cloud among binary candidates. No chance-alignment cloud is visible in the binary candidates with high bound probability (upper right), but a few chance alignments

³`ruwe`, the re-normalised unit-weight error, is a measure of astrometric goodness-of-fit that corrects for global trends in the other reported goodness-of-fit indicators with magnitude and color. Values above about 1.4 indicate potential problems.




Figure 10.13: Binary candidates from the upper left panel of Figure 10.12, now colored by the maximum **ruwe** of the two components. Most pairs with $\Delta V > \Delta V_{\text{orb}}$ and $s \ll 10^5$ au (those labeled “triples” in Figure 10.12) have **ruwe** $\gtrsim 1.4$ for at least one component, in most cases because that component is an unresolved binary. At $s \lesssim 500$ au, most pairs have large **ruwe** due to centroiding errors at close angular separations (see also Figure 10.18).

do scatter into the bound binary cloud, mostly at $\Delta V > \Delta V_{\text{orb}}$ (lower right).

We note that the interpretation of ΔV calculated from Equation 10.11 as a physical velocity difference between the two components of a binary breaks down at large angular separations, where projection effects become important (e.g. El-Badry 2019). Indeed, two stars in an ultra-wide binary can have *identical* space velocities but substantially different plane-of-the-sky proper motions. The magnitude of the apparent proper motion difference depends primarily on angular separation and is therefore largest for nearby binaries, which are also the binaries with the smallest $\sigma_{\Delta V}$. For the sample shown in the upper left panel of Figure 10.12, projection effects become important beyond about 20,000 au, which is – perhaps not coincidentally – the separation beyond which the trend of $\Delta V \sim s^{-1/2}$ appears to flatten. These projection effects can be corrected if the RVs of at least one component are known (El-Badry 2019).

10.5.5 Cross-match with LAMOST

We cross-matched the binary catalog with the LAMOST survey (Cui et al. 2012b, DR6 v2), the currently most extensive spectroscopic survey providing stellar parameters and abundances. We began with the LAMOST low-resolution “A, F, G and K Star” catalog, which contains atmospheric parameters and metallicities for 5,773,552 spectra (including some duplicate observations). We cross-matched the catalog with *Gaia* eDR3 using the CDS Xmatch service,⁴ which uses *Gaia* proper motions to propagate source positions to epoch J2000. We matched each LAMOST observation to the nearest *Gaia* source within 1 arcsec. For sources with more than one LAMOST observation, we retained only the observation with the highest *g*-band SNR. This left us with LAMOST data for 4,306,131 sources, which we then matched to our catalog using *Gaia* source ids.

This yielded 91,477 binaries in which at least one component has a LAMOST spectrum. This sample will be useful for a variety of applications, such as studying the dependence of the binary fraction on metallicity (e.g. El-Badry & Rix 2019a; Hwang et al. 2020). Here, we focus on a subset of the cross-match: those binaries in which *both* components have a LAMOST spectrum, and the angular separation is at least 3 arcseconds. The latter cut is to avoid cases where both stars fall inside a single fiber, leading to potentially biased stellar parameters and abundances (e.g. El-Badry et al. 2018f). The RVs of these binaries can be used to verify whether most binary candidates are bound. This test is similar to that shown in Figure 10.6 with *Gaia* RVs, but LAMOST spectra extend to fainter magnitudes than RVs from *Gaia*, which are currently only available at $G \lesssim 13$.

Figure 10.14 compares the LAMOSTs RVs and metallicities (i.e., [Fe/H]) of the components of these binaries. These are generally expected to be consistent for genuine wide binaries. The RVs for most binaries do indeed fall close to the one-to-one line, but there are some outliers: 91 of the 2840 binaries in the sample have RVs that are more than 3-sigma discrepant. A potential worry is that these pairs are not binaries at all, but chance

⁴<http://cdsxmatch.u-strasbg.fr/>

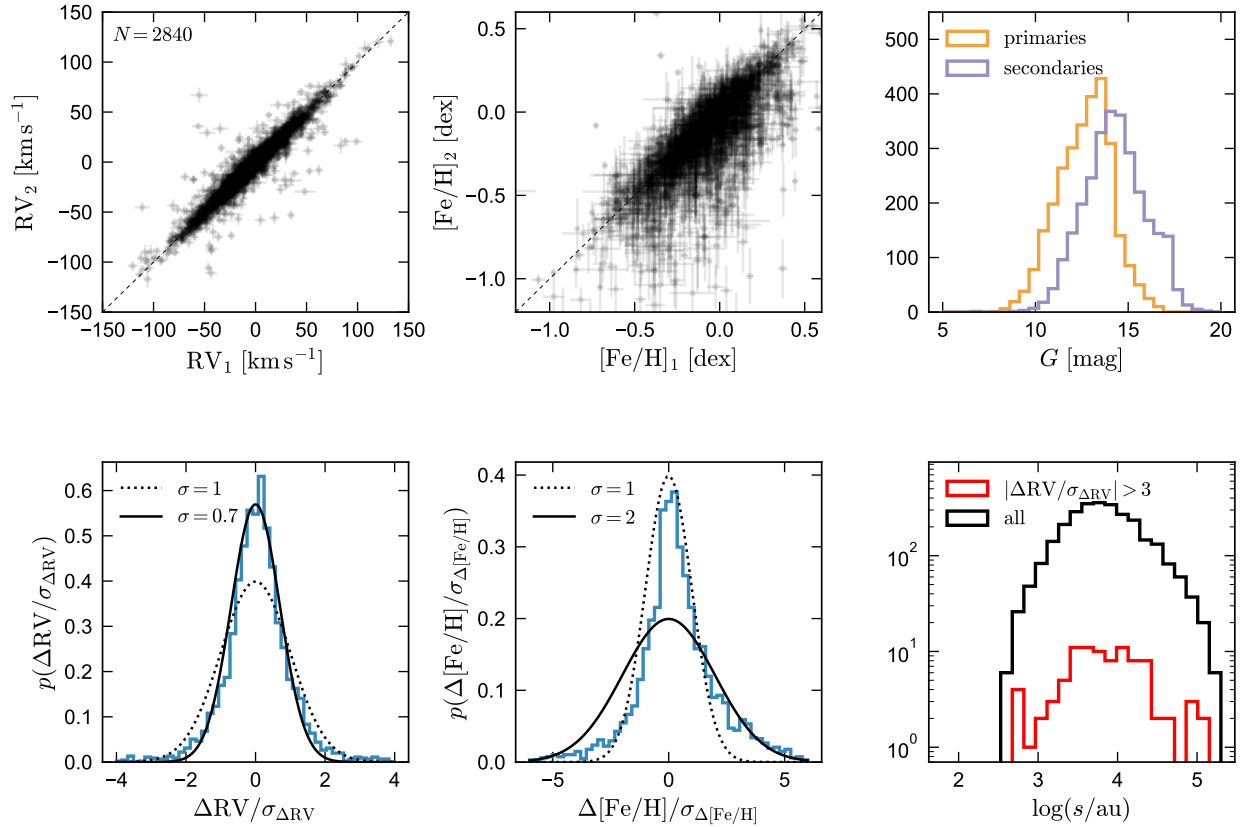


Figure 10.14: High-confidence binaries in which both components were observed by the LAMOST survey. We only include pairs with angular separation $\theta > 3$ arcsec to avoid blending. Top left and center: comparison of the RVs and metallicities ($[\text{Fe}/\text{H}]$) of the primary (x -axis) and secondary (y -axis). In most binaries, the two components have very similar RV and $[\text{Fe}/\text{H}]$, as expected. There is evidence of a systematic bias toward lower $[\text{Fe}/\text{H}]$ for fainter targets. Upper right: G magnitudes of both components. This sample includes fainter stars than the *Gaia* RVS sample. Bottom left and center: distributions of uncertainty-normalized RV and $[\text{Fe}/\text{H}]$ difference, compared to Gaussians with the listed σ . These suggest that the LAMOST RV uncertainties are overestimated by $\sim 30\%$ on average, while $[\text{Fe}/\text{H}]$ uncertainties are underestimated and subject to temperature systematics. Bottom right: separation distributions. Pairs in which RVs are not consistent within 3σ are shown in red. If these pairs were chance-alignments, one would expect them to be concentrated at the widest separations. They are not, suggesting that they are either higher-order multiples with the RV of one component biased by an unresolved close companion, or pairs in which one component has a catastrophically wrong RV.

alignments. To assess whether this is likely to be the case, we plot the separation distribution of the discrepant pairs in the bottom right panel of Figure 10.14. This shows that the separation distribution of pairs with large RV differences is similar to that of all pairs. Chance alignments are much more common at wide separations (Figure 10.4), so if chance alignments were the root of the discrepant RVs, one would expect these pairs to be clustered at large separations. The primary reason for the discrepant RVs is likely again higher-order multiplicity. This is likely to affect the LAMOST RVs more than it does those from *Gaia* (Figure 10.6), because the LAMOST RVs and their uncertainties are based on a single epoch. Unresolved short-period binaries are usually filtered out of the *Gaia* RV sample we consider, because σ_{RV} for that sample is calculated from the epoch-to-epoch RV dispersion, and we required $\sigma_{\text{RV}} < 10 \text{ km s}^{-1}$ for both components.

The lower left panel of Figure 10.14 shows the distribution of uncertainty-normalized RV difference between the two components of binaries; i.e., $\Delta\text{RV}/\sigma_{\Delta\text{RV}}$. The median $\sigma_{\Delta\text{RV}}$ for this sample is 7 km s^{-1} , which is larger than the typical $\sim 1 \text{ km s}^{-1}$ orbital velocity for these binaries, so we expect the width of the main distribution to be dominated by measurement uncertainties, with the tails dominated by higher-order multiples. If the reported σ_{RV} values are accurate, $\Delta\text{RV}/\sigma_{\Delta\text{RV}}$ should be distributed as a Gaussian with $\sigma \approx 1$, with some outliers at higher velocity difference. The distribution is indeed approximately Gaussian, but it is narrower than $\sigma = 1$; the bulk of the distribution is better approximated by $\sigma \approx 0.7$. This suggests that the LAMOST RV uncertainties are typically *overestimated* by $\sim 30\%$.

The distribution of uncertainty-normalized $[\text{Fe}/\text{H}]$ differences tells a different story. It is not well-described by a single Gaussian, but has a narrow component with $\sigma \lesssim 1$, and broad, asymmetric tails. This suggests that $\sigma_{[\text{Fe}/\text{H}]}$ values are considerably underestimated for a significant fraction of the catalog. We find (not shown in the figure) that the distribution becomes narrower and more Gaussian when we only consider binaries in which the two components have similar magnitude and effective temperature. This suggests that the larger-than-expected metallicity differences are due primarily to temperature systematics in the abundance pipeline. This is particularly evident in the upper middle panel of Figure 10.14, which shows that the secondaries (which have lower T_{eff}) systematically have lower reported $[\text{Fe}/\text{H}]$. It is of course possible that some binaries really have inconsistent surface abundances, but work with higher-quality spectra than those which underlie the LAMOST metallicities suggests such abundance anomalies are rare (Hawkins et al. 2020).

Many of the binaries in our catalog were also observed by other spectroscopic surveys. We defer analysis of these data to future work, but comment that the type of analysis shown in Figure 10.14 will be useful in calibrating the abundances derived by surveys (and their uncertainties). For example, an earlier version of our binary catalog constructed from *Gaia* DR2 was recently fruitfully used by Buder et al. (2020) to assess the reliability of abundances derived by the GALAH survey.




Figure 10.15: Distributions of uncertainty-normalized parallax difference between the two components of binaries, $\Delta\varpi/\sigma_{\Delta\varpi} = (\varpi_1 - \varpi_2)/\sqrt{\sigma_{\varpi,1}^2 + \sigma_{\varpi,2}^2}$. Each panel shows a different bin of angular separation and magnitude, with both components falling in the quoted magnitude range. Because the two components of a binary have essentially the same distance, this quantity would be expected to follow a Gaussian distribution with $\sigma = 1$ (dotted lines) if the formal parallax uncertainties were accurate. $\sigma > 1$ points toward underestimated parallax uncertainties. Solid black lines show Gaussian fits, with σ noted in the legends. At all magnitudes, the best-fit σ is larger at close separations, implying that σ_{ϖ} is more severely underestimated for sources with nearby companions. At fixed separations, the fractional underestimate of σ_{ϖ} is larger for bright stars.

10.6 Calibrating *Gaia* DR3 parallax uncertainties

Because the two stars in a wide binary have very nearly the same distance, our catalog provides a straightforward method of validating the *Gaia* eDR3 parallax uncertainties. We do this by calculating the uncertainty-normalized parallax difference between the two components of each binary, $\Delta\varpi/\sigma_{\Delta\varpi} = (\varpi_1 - \varpi_2) / \sqrt{\sigma_{\varpi,1}^2 + \sigma_{\varpi,2}^2}$, just as we did for the LAMOST RVs. In the limit of accurate parallax uncertainties (and small differences in the true distance to the two components), this quantity should be distributed as a Gaussian with $\sigma = 1$. If the reported parallax uncertainties are underestimated, one expects a wider distribution, and possibly deviations from Gaussianity.

Figure 10.15 shows distributions of $\Delta\varpi/\sigma_{\Delta\varpi}$ for binaries at a range of angular separations and magnitudes. To isolate underestimated random errors (as opposed to systematic shifts due to e.g. variations in the parallax zeropoint), we only show binaries in which both components have magnitudes in the quoted range (corresponding to a magnitude difference of $\Delta G < 1$). This figure only shows high-confidence binaries ($\mathcal{R} < 0.1$) in which both components have `ruwe` < 1.4 , indicating an apparently well-behaved astrometric solution. In this and all validation of the parallax errors, we exclude binaries that are wide and nearby enough that their physical size might measurably contribute to the parallax difference. Under the ansatz that the projected physical separation s is comparable to the line-of-sight distance difference and is much smaller than the distance, we expect the true parallax difference between the two stars to scale as $\Delta\varpi_{\text{true}} \sim 1/d - 1/(s+d) \approx s/d^2$, or

$$\Delta\varpi_{\text{true}} = \frac{1}{206265} \text{mas} \times \left(\frac{\theta}{\text{arcsec}} \right) \left(\frac{\varpi}{\text{mas}} \right). \quad (10.13)$$

This quantity is negligible compared to $\sigma_{\Delta\varpi}$ for most binaries in the catalog, but not for the nearest and widest binaries. We therefore exclude all binaries in which $\Delta\varpi_{\text{true}}/\sigma_{\Delta\varpi} > 0.05$.

Blue histograms in Figure 10.15 show the observed distributions. We show 5 bins of G magnitude, each 1 mag wide, and 3 bins of angular separation. At separations closer than 4 arcsec, the catalog contains binaries with $|\Delta\varpi|/\sigma_{\Delta\varpi} < 6$; at wider separations, $|\Delta\varpi|/\sigma_{\Delta\varpi} < 3$ (Section 10.3). Also shown in Figure 10.15 are Gaussian fits to the data. Because pairs with $|\Delta\varpi|/\sigma_{\Delta\varpi} < 3$ (or 6) do not enter the catalog, it is necessary to account for this truncation of the distribution. Particularly when parallax uncertainties are underestimated, simply calculating the sample standard deviation would underestimate the best-fit σ . We assume the observed values of $\Delta\varpi/\sigma_{\Delta\varpi}$ are drawn from a distribution defined as

$$p(\Delta\varpi/\sigma_{\Delta\varpi}) = \begin{cases} A \exp\left[-\frac{(\Delta\varpi/\sigma_{\Delta\varpi})^2}{2\sigma^2}\right], & |\Delta\varpi/\sigma_{\Delta\varpi}| < b \\ 0, & |\Delta\varpi/\sigma_{\Delta\varpi}| > b \end{cases} \quad (10.14)$$

where $A = \frac{1}{\sqrt{2\pi\sigma} \times \text{erf}(b/\sqrt{2\sigma})}$, erf is the error function, and $b = 3$ or 6 is the sigma-threshold above which binaries are rejected. The log-likelihood for a set of uncertainty-normalized




Figure 10.16: Fractional parallax uncertainty underestimate (i.e., the best-fit σ in Figure 10.15) as a function of magnitude and separation. Black symbols show binaries in which both components have `ruwe` $<$ 1.4; red symbols show those in which at least one component has `ruwe` $>$ 1.4, indicating a potentially problematic astrometric solution. From left to right, panels show increasing angular separation. As expected, sources with `ruwe` $>$ 1.4 have larger σ_{ϖ} underestimates at all magnitudes and separations. Sources with `ruwe` $<$ 1.4 have more severe underestimates of σ_{ϖ} at close separations and for bright stars, particularly at $11 < G < 13$.




Figure 10.17: Fractional parallax uncertainty underestimate (i.e., the best-fit σ reported in Figure 10.15) as a function of magnitude. All binaries considered have `ruwe` $<$ 1.4 for both components. Green points additionally have `ipd_gof_harmonic_amplitude` $<$ 0.1 and `ipd_frac_multi_peak` $<$ 10 for both components; blue points have at least one component that does not pass these cuts. Sources that do not pass the IPD cuts have more strongly underestimated uncertainties at all separations and magnitudes.




Figure 10.18: Fraction of pairs in which at least one component fails the indicated quality cut (larger values are indicative of a potentially problematic astrometric solution). All three indicators are significantly enhanced within $\theta \lesssim 2$ arcsec.

parallax differences is then

$$\ln L = \sum_i \ln p((\Delta\varpi/\sigma_{\Delta\varpi})_i), \quad (10.15)$$

where the sum is calculated over all binaries in the set. For each panel of Figure 10.15, we maximize Equation 10.15 to find the value of σ that best describes the truncated Gaussian. This is plotted with a solid line, and the value of σ is shown in the legend. For comparison, we also plot a Gaussian with $\sigma = 1$, the distribution expected in the limit of accurate parallax uncertainties. To avoid having a few outliers with strongly underestimated parallax uncertainties bias the fits, we exclude binaries with $|\Delta\varpi|/\sigma_{\Delta\varpi} > 3$ and set $b = 3$ in all separation bins. The figure shows that the distributions of $\Delta\varpi/\sigma_{\Delta\varpi}$ are indeed approximately Gaussian, and the fits are reasonably good representations of the data.

Two trends are clear in Figure 10.15: (a) at fixed apparent magnitude, the width of the observed distributions increases at closer separations, and (b) at fixed separation, their width

increases toward brighter magnitudes. That is, parallaxes are more strongly underestimated for bright sources with close companions. This is illustrated more clearly in Figure 10.16, in which black points with error bars show the best-fit Gaussian σ values for each bin of magnitude and angular separation, considering only sources with `ruwe` $<$ 1.4 for both components. At all separations, parallaxes are most severely underestimated at $G \approx 13$. This is likely related to the fact that the window class (i.e., the pixel sampling scheme around detected sources; see Rowell et al. 2020) changes at $G \approx 13$. The largest σ_{ϖ} underestimate factors are also accompanied by abrupt changes in the zeropoint (Lindgren et al. 2020a). For sources with $G < 13$, a 2D window is used and is fitted with a point spread function. For $G > 13$, only the collapsed 1D scan, which is fitted with a line spread function, is available. Sources near $G = 13$ have a mix of 2D and 1D windows and are likely more affected by any calibration issues.

Figure 10.16 also shows results for binaries in which at least one component has `ruwe` $>$ 1.4, indicative of a potentially problematic astrometric solution. As expected, the best-fit σ values are significantly larger for these binaries, at all separations and magnitudes.

A natural question is whether the broadened distributions of uncertainty-normalized parallax difference could be a result of contamination from chance alignments that are not actually bound, rather than underestimated parallax uncertainties. This hypothesis can be ruled out for two reasons. First, we can empirically estimate the chance-alignment rate for different subsets of the catalog (e.g. Figures 10.3 and 10.4), and we find it to be extremely low for the samples we use for parallax error validation. Second, our analysis suggests that σ_{ϖ} is most severely underestimated for bright binaries at close separations, and this is precisely the region of parameter space where the chance alignment rate is lowest (e.g. Figure 10.4).

Besides `ruwe`, *Gaia* eDR3 contains other diagnostics of potentially problematic astrometric fits. In particular, the parameter `ipd_gof_harmonic_amplitude` quantifies how much the image parameter determination goodness-of-fit varies with scan angle; a large value is likely indicative of a marginally-resolved binary. The related `ipd_frac_multi_peak` diagnostic quantifies in what fraction of scans multiple peaks are detected. Figure 10.17 separately plots the inferred σ_{ϖ} underestimate for sources in which both components have `ipd_gof_harmonic_amplitude` $<$ 0.1 and `ipd_frac_multi_peak` $<$ 10 (green), and those in which at least one component fails one of these cuts (blue). These thresholds are motivated by the experiments performed in Gaia Collaboration et al. (2020b). At all separations, the implied σ_{ϖ} underestimate is larger for sources that do not pass one of the IPD cuts; the difference is largest for bright pairs at close separations.

Figure 10.18 shows how the prevalence of problematic sources according to the `ruwe` and IPD diagnostics depends on separation. Considering all binaries in the catalog, we plot the fraction of pairs at a given separation in which at least one component does not pass the cut listed in the legend. At separations larger than a few arcseconds, this fraction is $\sim 10\%$ for the `ruwe` cut, and $\sim 3\%$ for both the IPD cuts. However, all three diagnostics of problematic solutions increase steeply at $\theta \lesssim 2$ arcsec, with the fraction approaching unity for the `ruwe` and `ipd_frac_multi_peak` cuts, and $\sim 40\%$ for the `ipd_gof_harmonic_amplitude` cut. This is not unexpected: for binaries at close separations, there will necessarily be two peaks




Figure 10.19: Same as Figures 10.16 and 10.17, but now comparing results obtained with *Gaia* DR2 vs DR3 parallaxes. All binaries considered have $\text{ruwe} < 1.4$ for both components in both DR2 and DR3. At $G \gtrsim 14$, the uncertainties are underestimated less in DR3 than in DR2. DR3 uncertainties are overestimated somewhat more severely at $11 < G < 13$.

in the image. This will unavoidably lead to biases in the image parameter determination, particularly for sources with $G > 13$, where the images are collapsed to 1D. There is also a danger of misattributing some scans to the wrong component in close pairs, leading to problems in the astrometric solution. In some cases, poor astrometric fits may also be due to astrometric acceleration (Section 10.7.3). Figures 10.16 and 10.17 show that the ruwe and IPD cuts are indeed useful for identifying sources with potentially problematic astrometric solutions. However, they likely do not catch all problematic sources: parallax uncertainties are underestimated somewhat even for sources that pass all cuts, and by a larger factor at close separations.

10.6.1 Comparison to *Gaia* DR2

Figure 10.19 compares the inferred underestimate factor of σ_w for DR2 and eDR3 astrometry. We attempt to match all the binary candidates in the catalog with *Gaia* DR2 using the `dr2_neighbourhood` catalog in the *Gaia* archive. For each component of each binary, we identify the likely corresponding DR2 source as the source within 100 mas that has the smallest magnitude difference compared to DR3. There are 1,894 primaries and 15,274 secondaries in the catalog for which no corresponding DR2 source could be identified. There are also 17,514 primaries and 62,424 secondaries for which there is a corresponding source in DR2 that only has a 2-parameter solution. Still, 96% of candidates have a corresponding DR2 source with a 5-parameter solution for both components. Figure 10.19 considers the subset of these binaries in which both components have $\text{ruwe} < 1.4$ in both DR2 and eDR3.

Overall trends with magnitude and separation are similar in DR2 and DR3. At $G > 13$, the inferred σ_w underestimates are smaller in DR3 at all separations, though they are




Figure 10.20: Median reported (dashed) and corrected (solid) parallax uncertainty as a function of G magnitude, for *Gaia* DR3 (black) and DR2 (red). The corrected values are obtained by multiplying the reported values by the factors shown in Figure 10.19, which are empirically determined from the reported parallax differences of the components of wide binaries. Panels show different bins of angular separation. Only stars with $\text{ruwe} < 1.4$ in both data releases are considered. The rightmost panel is appropriate for single stars outside of crowded fields.

still modest in DR2. Even at $G \sim 19$ and wide separations, where the inferred σ_w in DR3 are consistent with being accurate or very slightly overestimated, those in DR2 are underestimated by $\sim 5\%$ on average. However, at $11 < G < 13$, the DR3 uncertainties are overestimated more than those in DR2. Given that the reported values of σ_w decreased by a factor of two on average between DR2 and DR3 in this magnitude range, the DR3 parallaxes are still “better” on average than those from DR2, but Figure 10.19 implies that the true gains are somewhat more modest than those reported.

This is shown explicitly in Figure 10.20, which shows the reported and true (inferred) median σ_w for DR2 and DR3 as a function of apparent G – magnitude and angular separation. Dashed lines show the median value of σ_w in the sample, considering both primaries and secondaries. Solid lines show the result of multiplying these values by the appropriate factors from Figure 10.19. The corrected median σ_w values are at least 30% smaller in DR3 than in DR2 at all magnitudes and separations. For bright stars ($G \lesssim 13$), the gains are generally more than a factor of 2.

10.6.2 Parallax zeropoint corrections and 5-parameter vs 6-parameter solutions

The *Gaia* eDR3 parallax zeropoint is known to vary with apparent magnitude, color, and ecliptic latitude. We do not attempt to account for this variation when constructing the binary catalog. Because all the sources in the catalog have $\varpi > 1$ mas, the effects of zeropoint corrections, which are typically on the order of 0.02 mas, are modest. [Lindgren](#)

et al. (2020a) derived an empirical zeropoint for eDR3 using quasars, stars in the LMC, and binaries. Here we investigate whether “correcting” the parallaxes using the prescriptions they provide can reduce the inferred σ_{ϖ} underestimate factors.

Figure 10.21 shows the inferred σ_{ϖ} underestimate factors with and without the zeropoint correction. We also separately plot binaries in which both components have a 5-parameter astrometric solution and those in which both have a 6-parameter solution. Details about the differences between 5- and 6-parameter solutions are discussed in Lindegren et al. (2020b); 5-parameter solutions are generally more reliable. Binaries with one 5- and one 6-parameter solution are excluded. At $G \lesssim 18$, most sources have 5-parameter solutions. The exception is sources with a close companion ($\theta \lesssim 2$ arcsec), which usually lack both reliable colors and 5-parameter solutions.

At fixed separation and magnitude, the effects of applying the zeropoint correction are encouraging but modest: for widely-separated binaries with $G \approx 13$ and 5-parameter solutions, the inferred σ_{ϖ} underestimate factor decreases from ~ 1.30 to 1.25. Improvements are generally smaller at closer separations and for 6-parameter solutions. The small effect is not unexpected: because the binaries we consider all have small magnitude differences, the zeropoint corrections are similar for both components, and the parallax *difference* does not change much when the correction is applied. At close separations, where most stars have 6-parameter solutions, the underestimate factors are similar for 5- and 6-parameter solutions. This suggests that the increased uncertainties at close separations are not primarily due to the transition from 5- to 6-parameter solutions.

10.6.3 Color dependence

The parallax zeropoint is also known to depend somewhat on color, likely due to the color-dependence of the PSF (Lindegren et al. 2020a,b). We therefore investigate how the σ_{ϖ} underestimate factors depend on color in Figure 10.22, where we separately consider binaries in which both components are blue and those in which both components are red. The color boundary we use, $\nu_{\text{eff}} = 1.4 \mu\text{m}^{-1}$, corresponds to $G_{\text{BP}} - G_{\text{RP}} \approx 1.6$, which is approximately the red limit of the *Gaia* quasar sample. Figure 10.22 shows that for sources with $G > 16$, the inferred σ_{ϖ} underestimates are somewhat larger for red sources. We verified that there are not significant angular separation differences between the red and blue pairs at fixed magnitude, so color is the most likely driving variable.

10.6.4 Fitting function to inflate σ_{ϖ}

We fit a function to our inferred σ_{ϖ} inflation factors as a function of G magnitude, which can be used to empirically correct σ_{ϖ} values reported for isolated sources. To derive a correction appropriate for single sources with well-behaved astrometry, we consider binaries with $\theta > 5$ arcsec and $\text{ruwe} < 1.4$ for both θ components, and we apply the zeropoint correction from Lindegren et al. (2020a) to both components’ parallaxes. Because we observe a general

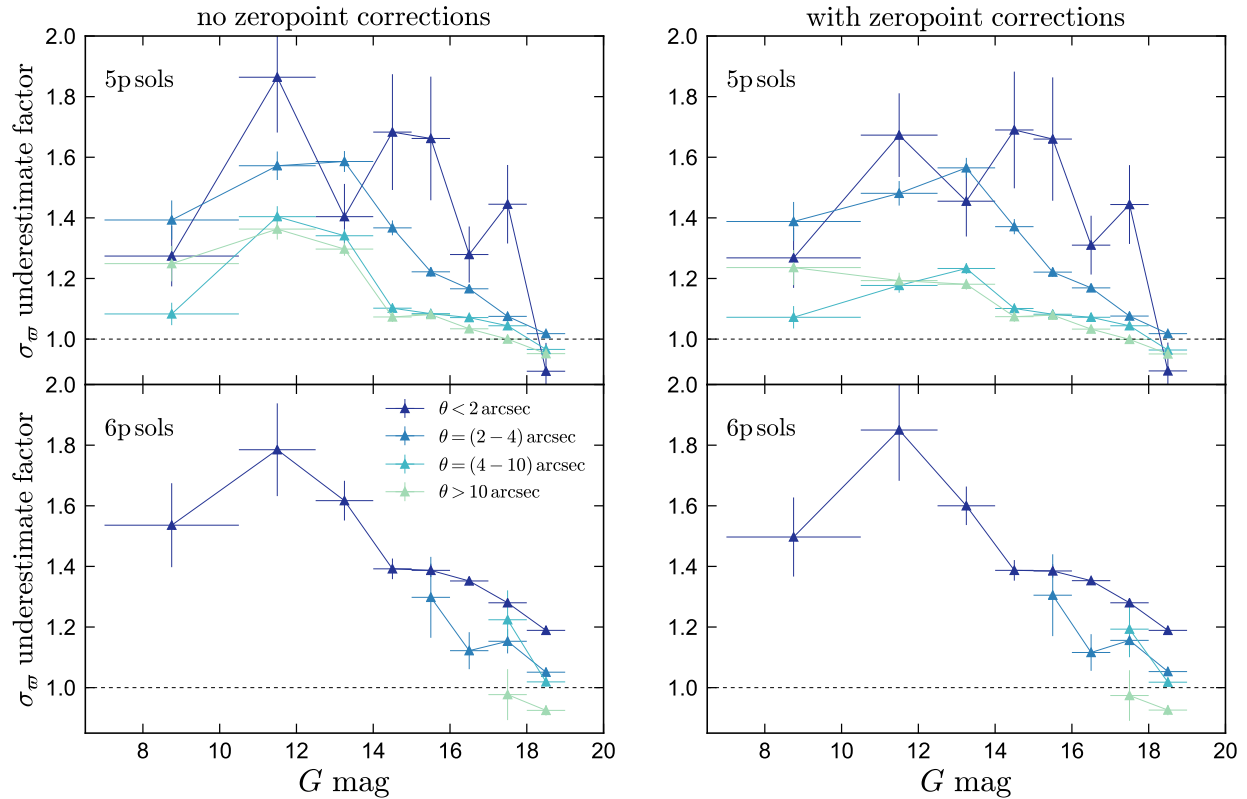


Figure 10.21: σ_w underestimate factors (similar to Figure 10.16) for binaries in which both components have five-parameters solutions (top) or both have six-parameter solutions (bottom). Left column uses the *Gaia* eDR3 parallaxes directly; right column applies the magnitude-, color-, and position-dependent parallax zeropoint correction from Lindegren et al. (2020a). For bright sources, only binaries with close angular separations have six-parameter solutions. The inferred σ_w underestimate factors depend mainly on separation and magnitude; they are similar for 5- vs 6-parameter solutions. Applying the zeropoint correction reduces the inferred overestimate factor slightly for bright sources.




Figure 10.22: Similar to Figure 10.16, but with binaries divided by color. Only binaries in which both components have `ruwe` < 1.4 and the two components fall in the same color bin are included. ν_{eff} refers to the parameters `nu_eff_used_in_astrometry` and `pseudocolour`, respectively, for 5- and 6-parameter solutions. The boundary of $\nu_{\text{eff}} = 1.4 \mu\text{m}^{-1}$ corresponds approximately to $G_{\text{BP}} - G_{\text{RP}} = 1.6$. Most quasars, white dwarfs, and MS stars with $M \gtrsim 0.6 M_{\odot}$ are bluer than this; lower-mass MS stars are redder. Left and middle panels show all separations; right panel shows wide separations. At the faint end, σ_{ϖ} overestimates are somewhat larger for red sources.

decline in the inflation factor with increasing G and a peak at $G \approx 13$ (Figure 10.23), we fit a polynomial plus a Gaussian bump:

$$f(G) = A \exp \left[-\frac{(G - G_0)^2}{b^2} \right] + p_0 + p_1 G + p_2 G^2 \quad (10.16)$$

We find $A = 0.21$, $G_0 = 12.65$, $b = 0.90$, $p_0 = 1.141$, $p_1 = 0.0040$, and $p_2 = -0.00062$. This is also plotted in Figure 10.23. Multiplying by f will – on average – correct reported σ_{ϖ} values for *Gaia* eDR3 and single-source solutions in DR3. The correction is appropriate for sources with $7 < G < 21$ that have `ruwe` < 1.4, have no comparably bright sources within a few arcsec, and have already had their parallaxes corrected by the zeropoint from [Lindgren et al. \(2020a\)](#). The effects of having a close companion on σ_{ϖ} likely depend on the brightness contrast and a variety of other factors; a rough estimate of the magnitude of the inflation can be obtained from Figure 10.16–10.17. The correction can be reasonably applied to both 5- and 6-parameter astrometric solutions. We do not fit separate corrections for red and blue sources, but we note that at the faint end, the inflation factors are generally somewhat smaller for blue sources.

10.6.5 Angular correlations in parallaxes

The σ_{ϖ} inflation factors inferred in this work and predicted by Equation 10.16 should be interpreted as lower limits. *Gaia* eDR3 parallaxes are subject to systematic trends on degree

scales (and larger) due to the scanning law (e.g. [Fabricius et al. 2020](#); [Lindegren et al. 2020b](#)). The angular separations of most of the binaries in the catalog are significantly smaller than this (Figure 10.7), so the “local” parallax zeropoint for the two stars is usually very similar. These local positional variations in the zeropoint are not accounted for in the correction from [Lindegren et al. \(2020a\)](#) and also inflate the effective parallax uncertainties. With a typical scale of about $10 \mu\text{as}$, these may contribute significantly to the uncertainties at $G \lesssim 13$.

A contemporaneous study by [Zinn \(2021\)](#) validated *Gaia* eDR3 parallaxes and their uncertainties using bright giants in the Kepler field ($9 \lesssim G \lesssim 13$) with independent distance estimates from asteroseismology. They tested our parallax uncertainty inflation function (Equation 10.16) and found it to perform well; i.e., no further uncertainty inflation was required after it was applied. The angular size of the Kepler field (~ 10 deg) is larger than the scale on which the strongest angular correlations in *Gaia* eDR3 are manifest (e.g. [Lindegren et al. 2020b](#), their Figure 14). This suggests that any additional uncertainty inflation required due to angular correlations in the zeropoint is modest.

10.6.6 Comparison to other work

The reliability of parallax uncertainties reported in *Gaia* eDR3 has been investigated by several other works. [Fabricius et al. \(2020, their Figure 19\)](#) used the dispersion in parallaxes reported for distant objects – quasars, stars in the LMC, and stars in dwarf spheroidal (dSph) galaxies, which should all have negligible true parallaxes – to estimate σ_ϖ inflation factors as a function of G magnitude. Their results from quasars and dSph stars, which are only available at the faint end ($G \gtrsim 16$) are broadly consistent with our results. In the LMC, they infer inflation factors that are larger at fixed magnitude than our results at wide separations, or their results for quasars and dSph at the same magnitude. This discrepancy is very likely a result of crowding: the source density in the LMC is large enough that a significant fraction of sources have another source within a few arcsec, and as we have shown (e.g. Figure 10.19), sources with companions within a few arcsec have more severely underestimated parallaxes. This likely also explains why [Maíz Apellániz et al. \(2021\)](#) found somewhat larger uncertainty inflation factors in globular clusters than we do with widely separated binaries. In the Kepler field, [Zinn \(2021\)](#) found a mean uncertainty estimate of $(22 \pm 6)\%$ for sources with $9 \lesssim G \lesssim 13$ (with most sources at the faint end of this range); this is consistent with our results.

10.7 Summary and discussion

We have constructed a catalog of 1.2 million high-confidence, spatially resolved wide binaries using *Gaia* eDR3 and have used them to empirically validate the reported parallax uncertainties in *Gaia* eDR3. Overall, the results are very encouraging: outside of crowded regions (i.e., for stars with no comparably bright companion within a few arcseconds), parallax uncertainties for sources with well-behaved astrometric fits are underestimated by *at*




Figure 10.23: Parallax uncertainty inflation factors inferred from widely-separated pairs with $\text{ruwe} < 1.4$ and parallaxes corrected by the zeropoint from [Lindgren et al. \(2020a\)](#). This is the σ_w underestimate factor for typical single sources with well-behaved astrometry. A fitting function (Equation 10.16) is provided.

most 30% (at $G \approx 13$), and by considerably less at fainter magnitudes. Our main results are as follows:

1. : *Catalog description:* The full catalog contains 1.2 million high-confidence binaries, including 15,982 WD+MS binaries and 1362 WD+WD binaries (Figures 10.7 and 10.8), and $\sim 20,000$ binaries containing giants and subgiants. The binaries span a projected separation range of a few au to 1 pc, have heliocentric distance up to 1 kpc, and include both (kinematic) disk and halo binaries. All binaries have reasonably precise astrometry, with $\varpi/\sigma_\varpi > 5$ for both components. The full catalog of 1.8 million binary candidates become dominated by chance alignments at $s \gtrsim 30,000$ au, but high-quality subsets can be selected that are relatively pure out to separations as large as 1 pc (Figure 10.4). The catalog builds on previous efforts to identify binaries using *Gaia* DR2, expanding the sample of known high-confidence binaries by a factor of 4 (Figure 10.10). This increase in sample size owes partly to the higher astrometric precision provided by *Gaia* eDR3, and partly to improvements in the binary identification and vetting strategy.
2. : *Quantifying and controlling chance alignments:* We estimate the contamination rate from chance alignments using two approaches: a mock catalog that does not contain any true binaries, and a version of the *Gaia* catalog in which stars have been artificially shifted from their true positions, removing real binaries but preserving chance-alignment statistics (Figure 10.3). Both approaches show that the full catalog has high purity at $s < 10,000$ but becomes dominated by chance alignments at $s \gtrsim 30,000$ au. Using the shifted chance alignment catalog, we show how one can select subsamples that have lower contamination rates, including some that are pure out to 1 pc (Figure 10.4). We also include in the catalog an estimate of the probability that each binary candidate is a chance alignment; this is constructed empirically from the distribution of known chance alignments in a seven-dimensional space of observables (Figure 10.5 and Appendix 10.9). We use radial velocities from *Gaia* (Figure 10.6) and LAMOST (Figure 10.14) to validate these probabilities.
3. : *Orbital velocities:* The high-precision of *Gaia* astrometry makes obvious the plane-of-the-sky velocity difference of the components of binaries due to orbital motion (Figure 10.12). About 200,000 binaries in the catalog have sufficiently accurate astrometry that the plane-of-the-sky velocity difference between the components can be measured with accuracy $\sigma_{\Delta V} < 100 \text{ m s}^{-1}$. A Keplerian decline in the velocity difference, $\Delta V \propto s^{-1/2}$, is visible out to $s \sim 20,000$ au, where projection effects become important. The requirement of proper motions consistent with Keplerian orbits excludes a significant fraction of hierarchical triples and higher-order multiples from the catalog (Figure 10.13).
4. : *Validation of *Gaia* DR3 parallax uncertainties:* We use the sample of high-confidence binaries to validate the published parallax uncertainties included in *Gaia* eDR3. This

analysis makes use of the fact that the two stars in a binary have essentially the same distance and thus should generally have reported parallaxes that are consistent within their uncertainties (Figure 10.15). We find that the published uncertainties are accurate for faint stars ($G \gtrsim 18$) that have well-behaved astrometric solutions and do not have a companion within a few arcseconds (Figure 10.16). They are underestimated somewhat for brighter stars, particularly in the range of $11 < G < 13$, where the published uncertainties should be multiplied by a factor of 1.3 on average. The degree to which uncertainties are underestimated is larger for sources with large `ruwe`, `ipd_gof_harmonic_amplitude`, and `ipd_frac_multi_peak` (Figure 10.16 and 10.17), and is larger for red sources than blue sources (Figure 10.22). The reported parallax uncertainties are generally more reliable in *Gaia* eDR3 than they were in DR2, except at $11 < G < 13$ (Figure 10.19). In an absolute sense, the σ_ϖ values improved by at least 30% from DR2 to eDR3, at all magnitudes and separations (Figure 10.20). We provide an empirical fitting function to correct reported σ_ϖ values (Figure 10.23).

Parallax uncertainties are underestimated more for binaries with angular separations less than a few arcseconds. Sources with resolved close companions are more likely to have high `ruwe` and IPD diagnostics related to binarity (Figure 10.18), but the underestimates of σ_ϖ is enhanced at close separations even for pairs in which both components have low `ruwe` and IPD flags. This is true both for sources with 5- and 6-parameter solutions (Figure 10.21).

10.7.1 Scientific uses for the catalog

This paper was primarily concerned with assembling the wide binary catalog. Here we note a few possible uses for the sample, which will be pursued in future work.

- *Calibrating stellar ages*: A useful property of wide binaries is that the two stars have basically the same total age, but can have different evolutionary states. If the age of one component can be constrained (e.g., because it is a WD, a subgiant, or a giant with asteroseismic mass constraints), that age constraint can be transferred to the companion (e.g. Chanamé & Ramírez 2012; Fouesneau et al. 2019; Qiu et al. 2020). This in turn can be used to calibrate more poorly-understood age indicators for MS stars, such as gyrochronology, stellar activity, and X-ray luminosity (e.g. Soderblom 2010).
- *The initial-final mass relation*: The WD+WD sample will be useful for constraining the initial-final mass relation (IFMR) for WDs: the masses and cooling ages of both WDs can (often) be well-constrained from photometry, and the IFMR can be constrained by the fact that both WDs have the same total age (and presumably, follow the same IFMR; e.g., Andrews et al. 2015). The subset of the WD+MS sample in which the age of the non-WD component can be constrained independently will also be useful for constraining the IFMR (e.g. Catalán et al. 2008).

- *WD masses from gravitational redshift:* Because the WD and MS star in a wide binary have essentially the same RV, the difference in their apparent RVs is due primarily to the WD’s gravitational redshift (which is typically $20 - 100 \text{ km s}^{-1}$). This provides a useful way of measuring WD mass that are essentially model-independent if the distance is well-constrained (e.g. [Koester 1987](#)). Masses from gravitational redshift will be particularly useful for measuring the mass distributions of WDs with rare spectral types and poorly-understood formation histories, such as the “Q-branch” WDs revealed by the *Gaia* CMD ([Bergeron et al. 2019](#); [Cheng et al. 2020](#)).
- *Abundances for WD progenitors:* a MS companion provides a window into the chemical abundances of a WD’s progenitor, which are otherwise inaccessible. Among other applications, this provides an avenue to compare the primordial and final abundances of disintegrating planets around polluted WDs, which represent a significant fraction of the WD population ([Koester et al. 2014](#); [Farihi 2016](#)).
- *Wide binary spin alignment:* An open question in binary star formation is how aligned the spins vectors of binaries are, and how this varies with separation (e.g. [Justesen & Albrecht 2020](#)). Of particular interest is whether the spins of excess “twin” binaries, which may have formed in circumbinary disks, are more aligned than those of non-twins at the same separation ([El-Badry et al. 2019](#)). Spin inclinations can be measured by combining a spectroscopic measurement of $v \sin i$ with a rotation period measured from spots (which requires a light curve) and a radius from parallax and temperature. About 6,000 of the high-confidence binaries in the catalog have both components with $G < 14$ and angular separations $\theta > 30$ arcseconds. These are ideal for follow-up study with TESS, because they are separated widely enough that high-quality light curves are available for both components separately. About 39,000 binaries in the catalog are in the K2 fields, and 5,000 are in the Kepler field.
- *Calibration of spectroscopic surveys:* The surface abundances of stars in wide binaries are generally very similar (e.g. [Hawkins et al. 2020](#)). This enables diagnosis of systematics in the abundances reported by surveys (e.g. Figure 10.14). Similar analyses can be done with stars in clusters, but binaries are more abundant and populate abundance space more densely than clusters.
- *Dynamical probes:* At separations wider than about 10,000 au, wide binaries are susceptible to dynamical disruption through gravitational encounters with other stars, compact objects, or molecular clouds (e.g. [Weinberg et al. 1987](#)). This makes the wide binary separation distribution a sensitive probe of the population of possible perturbers (e.g. [Yoo et al. 2004](#); [Tian et al. 2020](#)).

10.7.2 Constraining the parallax zeropoint with binaries

A further application of the *Gaia* binary sample, which we have not explored in this work, is calibration of the parallax zeropoint. We have intentionally limited our analysis

to binaries in which both stars have almost the same magnitude, (and, because almost all the stars are on the main sequence, the same color). This avoids complication arising from the magnitude- and color-dependence of the parallax zeropoint, which should be nearly the same for both stars. If we were to consider binaries with substantially different primary and secondary magnitudes, magnitude-dependence of the zeropoint would be manifest as a shift in the *mean* signed parallax difference: that is, the distributions in Figure 10.15 would no longer be centered on zero. If the absolute zeropoint at a particular magnitude and color can be pinned down from external data (e.g. quasars at the faint/blue end), binaries then allow for determination of the zeropoint at all other magnitudes and colors. Some analysis along these lines was carried out by Lindegren et al. (2020a) and Fabricius et al. (2020).

A challenge to carrying out this type of calibration with our current catalog is Lutz-Kelker bias (Lutz & Kelker 1973): because the fainter secondaries have larger parallax errors, their parallaxes in our sample will on average be overestimated more than those of the primaries. That is, in the absence of any magnitude-dependence of the zeropoint, the mean value of $\varpi_1 - \varpi_2$ would be negative. This bias must be eliminated or accounted for in order for reliable determination of the zeropoint from binaries to be feasible. The most straightforward path forward is likely to select binaries without explicit cuts on parallax.

10.7.3 Astrometric acceleration

We have shown that parallax uncertainties are more severely underestimated at close angular separations (e.g. Figure 10.16), and that a larger fraction of close binaries have $\text{ruwe} > 1.4$ for at least one component (Figure 10.18). Here we consider whether this is likely due to actual astrometric acceleration, or other issues.

Whether orbital acceleration of a binary is detectable depends on a variety of factors, including the eccentricity, orientation, and phase of the orbit. Here we derive a crude estimate. We consider a face-on circular orbit with $M_1 + M_2 = 1M_\odot$ and $M_1 \gg M_2$, with an angular separation θ viewed at a distance d . The semi-major axis is $a = 1 \text{ au} \times (d/\text{pc}) (\theta/\text{arcsec})$, and the orbital period is $P = 1 \text{ yr} \times (d/\text{pc})^{3/2} (\theta/\text{arcsec})^{3/2}$. During the 34-month baseline of *Gaia* DR3, the azimuthal angle swept out by the secondary is

$$\phi = 2\pi \times ((34 \text{ months}) / P) \quad (10.17)$$

$$= 0.018 \text{ radians} \times (d/(100 \text{ pc}))^{-3/2} (\theta/\text{arcsec})^{-3/2}, \quad (10.18)$$

where $\phi = 2\pi$ would signify a full orbit.

In the limit of small ϕ , the total orbital motion that is perpendicular to the instantaneous proper motion vector at the first observation is $\vartheta_\perp \approx \frac{1}{2}\theta\phi^2$, from Taylor expanding $x = \theta \cos \phi$. The total perpendicular deviation from linear motion is thus

$$\vartheta_\perp \approx 0.16 \text{ mas} \times \left(\frac{d}{100 \text{ pc}}\right)^{-3} \left(\frac{\theta}{\text{arcsec}}\right)^{-2} \quad (10.19)$$

$$\approx 0.16 \text{ mas} \times \left(\frac{d}{100 \text{ pc}}\right) \left(\frac{s}{100 \text{ au}}\right)^{-2}. \quad (10.20)$$

The deviation from the *best-fit* single-star orbit will likely be a factor of a few smaller than this.

To determine whether orbital acceleration is plausibly detectable, this quantity can be compared to the typical astrometric precision (e.g. Figure 10.20). At the typical distance of binaries in the catalog, $d \approx 500$ pc, the predicted deviation for $\theta = 1$ arcsec is on the order of 0.001 mas, well below the sensitivity of *Gaia* eDR3. Astrometric acceleration due to orbital motion is thus not expected to be detectable for the large majority of binaries in our catalog, and it is therefore likely that the more strongly underestimated σ_{ϖ} at close separations is primarily due to other issues, such as centroiding errors or some scans being attributed to the wrong component.

Astrometric acceleration should, in principle, be non-negligible for the nearest and closest binaries in the catalog. Considering only binaries with $s < 100$ au, the median deviation predicted by Equation 10.20 is 0.3 mas, which is larger than σ_{ϖ} for the majority of that sample. To investigate whether there is evidence of acceleration in our sample at close separations, we compared the DR2 and eDR3 proper motions of both components, under the assumption that acceleration should manifest as a change in mean proper motion from epoch 2015.5 to 2016.0 (e.g. Kervella et al. 2019). It is important to note that the coordinate systems of DR2 and eDR3 are not identical. An ad-hoc correction was applied to the eDR3 coordinate frame to remove a $\sim 0.1 \text{ mas yr}^{-1}$ rotation that was present in the coordinate system for bright stars in DR2 (Lindgren et al. 2020b). Properly aligning the coordinate systems between the two releases is nontrivial (e.g. Brandt 2018). For this reason, and because proper motion uncertainties, like parallax uncertainties, are likely underestimated somewhat, it is beyond the scope of our investigation to determine which proper motion differences are significant. Instead, we simply consider how the fraction of sources with proper motion differences above a particular threshold depends on separation. We find that the fraction of binaries with inconsistent proper motions for one or both components is strongly enhanced at close angular separations, as would be expected in the presence of accelerations. For the full binary catalog, the fraction of pairs that have at least one component with DR2 and eDR3 proper motions inconsistent within 3σ is 18% at $\theta > 4$ arcsec, but 50% at $\theta < 1$ arcsec, and 85% at $\theta < 0.5$ arcsec.

However, the fraction of sources with inconsistent proper motions is significantly enhanced even at $d > 500$ pc, where Equation 10.20 suggests that any perpendicular acceleration should be negligible. Indeed, at fixed magnitude, the fraction of sources with inconsistent proper motions depends primarily on angular, not physical, separation. This suggests spurious astrometry (due to bias in the image parameter determination or source misidentification) for sources with close companions is the primary cause for the apparent acceleration. Although Equation 10.20 suggests that acceleration should often be detectable in our sample at $s < 50$ au, the expected sensitivity is not yet realized there due to problematic astrometry for barely-resolved sources. The detectability of accelerations with *Gaia* was also investigated by Belokurov et al. (2020). They found that while *ruwe* is often enhanced in close binaries, the enhancement can be reliably tied to orbital motion only in the regime where a significant (order unity) fraction of the orbit is covered by the *Gaia* time baseline.

10.8 Catalog description

The full binary catalog will be hosted at CDS. It can also be accessed [online](#). All columns in the `gaiaedr3.gaia_source` catalog are copied over for both components. We also include the columns `source_id`, `parallax`, `parallax_error`, `pmra`, `pmdec`, `pmra_error`, `pmdec_error`, and `ruwe` from *Gaia* DR2 for both components; these have the prefix `dr2_`.

Columns ending in “1” and “2” refer to the primary and secondary component, respectively. The primary is always the component with the brighter *G* magnitude. We also include columns `pairdistance` (angular separation θ , in degrees), `sep_AU` (projected separation s , in au), `R_chance_align` (\mathcal{R} ; Equation 10.8), and `binary_type` (e.g. MSMS, WDMS, etc.; see Table 10.1). The ordering in `binary_type` does not account for primary/secondary designations; i.e., all binaries containing a WD and a MS star are designated WDMS, irrespective of whether it is the WD or MS component that is brighter.

The shifted chance alignment catalog is also available. It contains the same columns as the binary candidate catalog, except the *Gaia* DR2 columns. Because one component of each pair has been shifted from its true position in the `gaia_source` catalog, the `ra` and `dec` columns in it do not match those reported in the `gaia_source` catalog.

Acknowledgements

We thank the anonymous referee for a constructive report, Jackie Faherty for help creating visualizations of the catalog, and Eliot Quataert, Dan Weisz, Jan Rybizki, and Anthony Brown for helpful comments. We acknowledge earlier discussions with Tim Brandt that proved seminal for this paper. We are grateful to Geoff Tabin and In-Hei Hahn for their continued hospitality during the writing of this manuscript. KE was supported by an NSF graduate research fellowship and a Hellman fellowship from UC Berkeley. T.M.H. acknowledges support from the National Science Foundation under Grant No. AST-1908119.

This project was developed in part during the 2020 virtual eDR3 Unboxing *Gaia* Sprint. This work has made use of data from the European Space Agency (ESA) mission *Gaia* (<https://www.cosmos.esa.int/gaia>), processed by the *Gaia* Data Processing and Analysis Consortium (DPAC, <https://www.cosmos.esa.int/web/gaia/dpac/consortium>). Funding for the DPAC has been provided by national institutions, in particular the institutions participating in the *Gaia* Multilateral Agreement. Guoshoujing Telescope (the Large Sky Area Multi-Object Fiber Spectroscopic Telescope LAMOST) is a National Major Scientific Project built by the Chinese Academy of Sciences. Funding for the project has been provided by the National Development and Reform Commission. LAMOST is operated and managed by the National Astronomical Observatories, Chinese Academy of Sciences. This research made use of Astropy,⁵ a community-developed core Python package for Astronomy ([Astropy Collaboration et al. 2013b, 2018b](#)). This research made use of the cross-match service provided by CDS, Strasbourg.

⁵<http://www.astropy.org>

Data Availability

All the data used in this paper is publicly available. The *Gaia* data can be retrieved through the *Gaia* archive (<https://gea.esac.esa.int/archive>), and the LAMOST data are available at <http://dr6.lamost.org>. The binary catalog and code to produce it can be found at <https://zenodo.org/record/4435257>.

10.9 Chance alignment probabilities

We estimate the local density in parameter space of binary candidates and known chance alignments from the shifted catalog using a Gaussian kernel density estimate (KDE). The parameters (“features”) we use are listed in Table 10.2. Most of them are described in Section 10.3.

We also add a measure of the local sky density, Σ_{18} . This represents the number of sources per square degree that (a) pass the cuts of our initial query (Section 10.3) and (b) are brighter than $G = 18$. We calculate the value of Σ_{18} around every binary candidate, counting the number of sources within 1 degree of the primary and dividing by π . Values of Σ_{18} range from 280 toward the Galactic poles to 8700 toward the Galactic center. A significant fraction of the sources toward the Galactic center are likely background stars that are not actually within the 1 kpc search volume.

We rescale the features so that they have similar dynamic range. This is accomplished by applying a few simple functions, which are listed in the “scaled parameter” column of Table 10.2. The rescaled parameters all have a dynamic range of about 4. We then calculate a 7-dimensional Gaussian KDE using a bandwidth $\sigma = 0.2$.

The distance metric in this parameter space is somewhat ill-defined due to the different units and distributions of the features. We nevertheless proceed boldly, making no claim that the set of features, rescalings, or the choice of kernel are optimal. Our choices are designed to make the kernel (a) narrow enough that it does not smooth over the sharpest features in the data, and (b) wide enough to prevent density peaks around individual, discrete binaries (overfitting).

When calculating the KDE for the binary candidates, we use a leave-10% out method wherein the density at the positions of 10% of the binary candidates is evaluated using a KDE constructed from the other 90%. To minimize discreteness noise in the chance-alignment KDE, we produce 30 different realizations of the shifted chance alignment catalog, shifting the declination of each star by a random variable $\mathcal{U}(-0.5, 0.5)$ degrees for each realization. We combine the realizations when calculating the KDE, and then divide the calculated density by 30 to reflect the number of pairs in a single realization.

Figure 10.24 compares the distributions of shifted chance alignments, all binary candidates, and candidates with $\mathcal{R} < 0.1$, in the space of features used for the KDE. It is clear that there are two modes in the binary candidate distribution, only one of which has a corresponding population in the chance alignment catalog. The clearest divisions between

Parameter	units	scaled parameter	(1,99)% range	scaled range	Description
θ	arcsec	$\log \theta$	(0.7, 607)	(-0.16, 2.78)	angular separation
ϖ_1	mas	$4/\varpi_1$	(1.02, 10.7)	(0.37, 3.93)	parallax (primary)
$\sigma_{\Delta\varpi}$	mas	$4\sigma_{\Delta\varpi}$	(0.02, 0.74)	(0.08, 2.96)	parallax difference error
Σ_{18}	deg^{-2}	$4 \log(\Sigma_{18})$	(345, 2790)	(10.15, 13.78)	$G < 18$ local source density
$v_{\perp,1}$	km s^{-1}	$v_{\perp,1}/50$	(3.4, 121)	(0.07, 2.43)	tangential velocity (primary)
$\Delta\varpi/\sigma_{\Delta\varpi}$	–	$ \Delta\varpi /\sigma_{\Delta\varpi}$	(-3.24, 3.30)	(0.02, 3.95)	normalized parallax difference
$(\Delta\mu - \Delta\mu_{\text{orbit}})/\sigma_{\Delta\mu}$	–	$2\text{erf}[(\Delta\mu - \Delta\mu_{\text{orbit}})/\sigma_{\Delta\mu}]$	(-88, 1.97)	(-2, 1.99)	scaled proper motion difference

Table 10.2: Features used by our Gaussian KDE in computing the local density of binaries and chance alignments. The feature vector is the “scaled parameter” column, in which all variables have been rescaled to have comparable dynamic range. We list the middle-98% ranges of both the raw and scaled parameters.

chance-alignments and binaries are in the dimensions of angular separation and proper motion difference (see also Figure 10.12), but their distributions also differ in other features. For example, a binary candidate is more likely to be genuine if the parallax errors are small (low $\sigma_{\Delta\varpi}$), or if it is found in a region of low stellar density (low Σ_{18}).

Figure 10.25 shows the ratio of the number of chance alignments (from the shifted catalog) that have a given \mathcal{R} value to the number of binary candidates with similar \mathcal{R} . If \mathcal{R} is interpreted as the probability that a candidate is a chance alignment, one would expect this ratio to follow the one-to-one line. It does indeed fall close to the one-to-one line (dashed) but with some deviations, likely due to over-smoothing of the KDE. This figure suggests, for example, that about 6% of binary candidates with $\mathcal{R} \sim 0.1$ are chance alignments, implying that chance alignment probabilities inferred when \mathcal{R} is interpreted as a probability are conservative.

A small fraction of binary candidates in the catalog have extremely small \mathcal{R} values; i.e., 0.5% have $\log(\mathcal{R}) < -10$, and 0.09% have $\log(\mathcal{R}) < -20$. These are primarily at close separations, where the chance alignment probability is indeed very low, but in this regime the \mathcal{R} values should not be interpreted as probabilities due to the finite size of the shifted chance alignment catalog.

10.9.1 Sources with spurious astrometry

A non-negligible fraction of sources in *Gaia* eDR3 have spurious astrometric solutions, meaning that they have large reported parallaxes and small reported uncertainties, but the parallaxes are significantly in error. The types of problems that can cause spurious solutions – typically crowding and marginally resolved sources – are generally equally likely to produce positive and negative parallaxes. The impact of spurious sources on our sample can thus be assessed by considering sources with significant negative parallaxes.

To this end, we repeat our initial ADQL query (Section 10.3) but require `parallax < -1` and `parallax_over_error < -5`. This yields 2,877,625 sources, implying that about 4.5% of the sources returned by the initial query have spurious solutions. We add the sources to our initial sample, treating their parallaxes as if they were positive. We then repeat the neighbor-counting procedure described in Section 10.3.1 for these sources, again removing

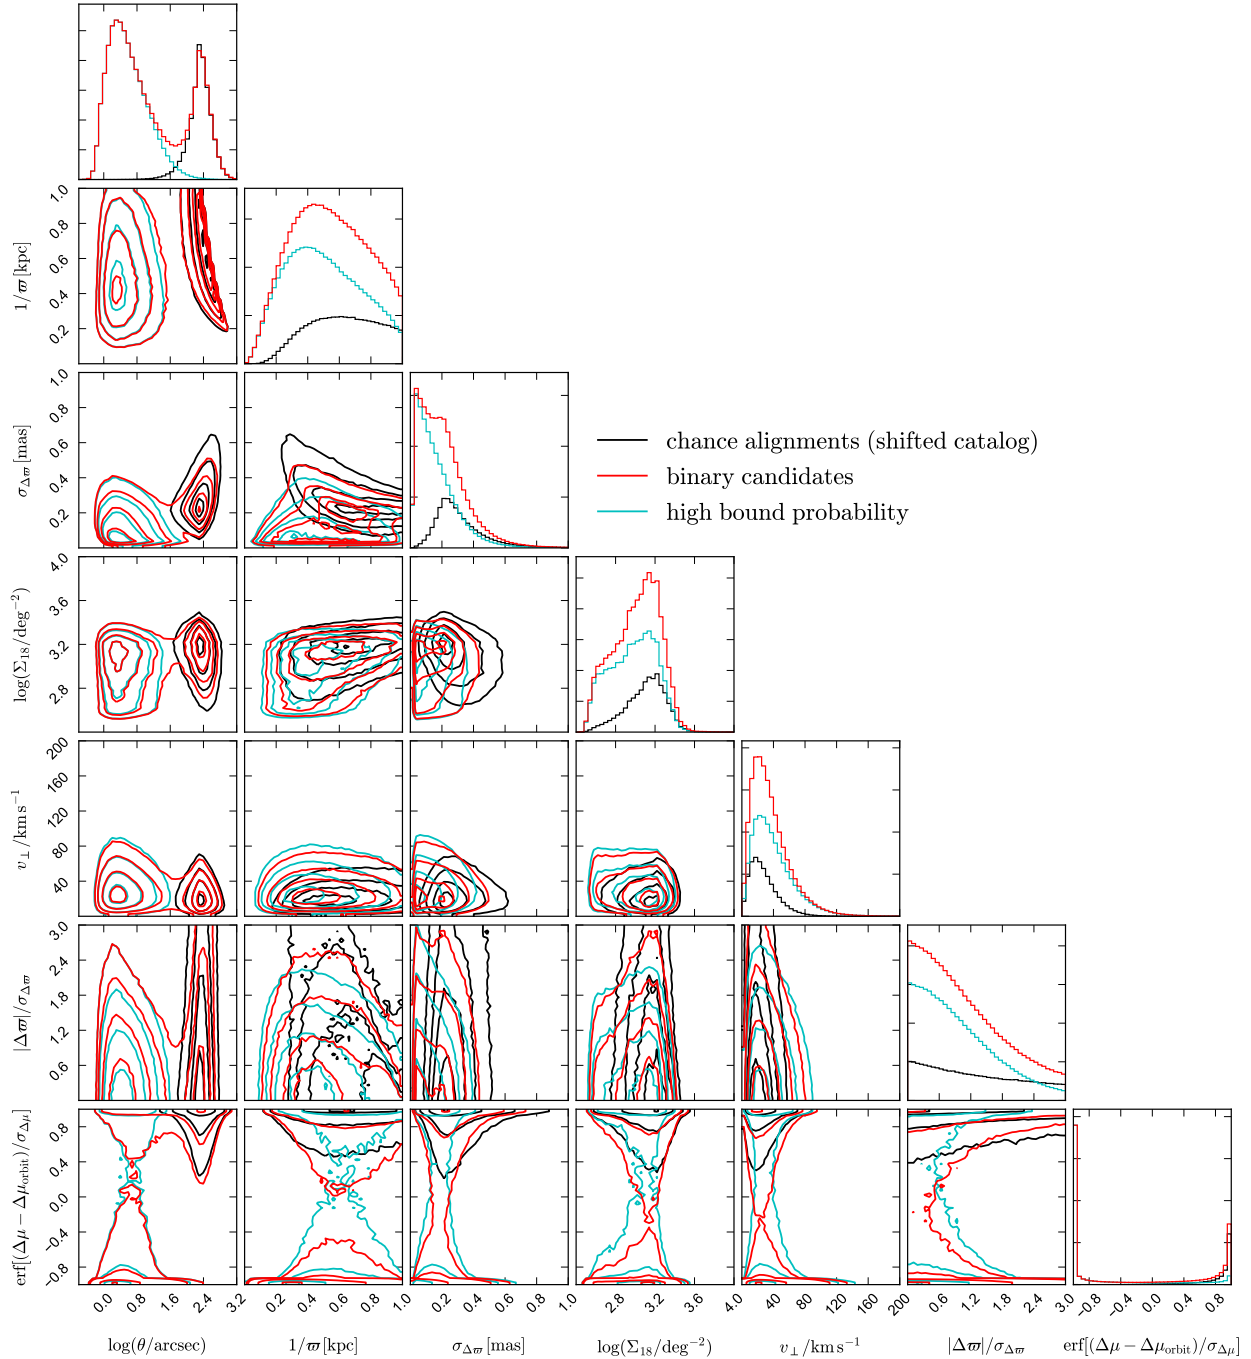


Figure 10.24: Parameter distribution of chance alignments from the shifted catalog (black), binary candidates (red), and high-confidence binary candidates with $\mathcal{R} < 0.1$ (section 10.4.2; cyan). Compared to chance alignments, the high-probability binaries have smaller angular separations, closer distances, larger parallax uncertainties, higher local source densities, larger tangential velocities, and more consistent parallaxes and proper motions.




Figure 10.25: \mathcal{R} is the ratio of the local “density” of chance-alignments from the shifted catalog to that of binary candidates (Equation 10.8). We compute \mathcal{R} for binary candidates and for chance alignments from a realization of the shifted catalog. We then plot the ratio of the number of pairs in the chance alignment catalog that have a given \mathcal{R} value to pairs in the binary candidate catalog with the same \mathcal{R} value. Dashed line shows a one-to-one relation for comparison.

objects with more than 30 neighbors. Of the 2,877,625 known spurious sources, only 380,379 (13%) survive this cut. That is, spurious sources are overwhelmingly found in regions of high source density, and a majority of them are removed by the first pass of cleaning.

We then carried out the full catalog construction procedure, now operating on an input sample that includes the initially selected sources as well as the known spurious sources with negative parallaxes, where the sign of the parallax is inverted for the known spurious sources. This yielded 15,852 candidate pairs in which one component is from the spurious sample. As expected, these pairs are concentrated at large separations; only 187 (593) have projected separations $s < 10,000$ au ($s < 30,000$ au). Finally, we repeat the calculation of \mathcal{R} on the candidates in which at least one component is known to be spurious, yielding 133 pairs with $\mathcal{R} < 0.1$. This implies that about 1 in 10,000 binary candidates with $\mathcal{R} < 0.1$ contains a source with a spurious parallax as defined here.

Chapter 11

LAMOST J0140355+392651: An evolved cataclysmic variable donor transitioning to become an extremely low mass white dwarf

An earlier version of this article has been submitted for publication as El-Badry, K., Quataert, E., Rix, H.-W., Weisz, D. R., Kupfer, T., Shen, K., Xiang M., Yang Y., Liu, X., 2021, MNRAS, submitted.

11.1 Abstract

We present LAMOST J0140355+392651 (hereafter J0140), a close ($P_{\text{orb}} = 3.81$ hours) binary containing a bloated, low-mass ($M \approx 0.15M_{\odot}$) proto-white dwarf (WD) and a massive ($M \approx 0.95M_{\odot}$) WD companion. The system's optical light curve is dominated by large-amplitude ellipsoidal variability but also exhibits additional scatter, likely driven by pulsations. The proto-WD is cooler ($T_{\text{eff}} = 6800 \pm 100$ K) and more puffy ($\log [g/(\text{cm s}^{-2})] = 4.74 \pm 0.07$) than any known extremely low mass (ELM) WD, but hotter than any known cataclysmic variable (CV) donor. It either completely or very nearly fills its Roche lobe ($R/R_{\text{Roche lobe}} = 0.99 \pm 0.01$), suggesting ongoing or recently terminated mass transfer. No dwarf nova-like outbursts have been observed. The spectrum is dominated by the proto-WD but shows tentative hints of H α emission, perhaps due to accretion onto the massive WD. The properties of the system are well-matched by MESA binary evolution models of CVs with donors that underwent significant nuclear evolution before the onset of mass transfer. In these models, the bloated proto-WD is either still losing mass via stable Roche lobe overflow or was doing so until very recently. In either case, it is evolving toward higher temperatures at near-constant luminosity to become an ELM WD. If the system is detached, mass transfer likely ended when the donor became too hot for magnetic braking to remain efficient. Evo-

lutionary models predict that the binary will shrink to $P_{\text{orb}} \lesssim 10$ minutes within a few Gyr, when it will either merge or become an AM CVn binary. J0140 provides an observational link between the formation channels of CVs, ELM WDs, detached ultracompact WD binaries, and AM CVn systems.

11.2 Introduction

Extremely low mass white dwarfs (ELM WDs, with masses $M \lesssim 0.25M_{\odot}$) are degenerate and semi-degenerate helium stars that never ignited core helium burning. The Universe is too young to produce ELM WDs by single-star evolution. It is thus thought that all ELM WDs are products of binary evolution: they are the stripped cores of stars that were initially more massive ($M_{\text{init}} \gtrsim 1M_{\odot}$) but lost most of their envelope to a companion (e.g. [Iben & Tutukov 1986](#)). This envelope stripping can occur through common envelope evolution or through stable mass transfer (e.g. [Li et al. 2019](#)). In the stable mass transfer channel, short-period ELM WDs are descendants of cataclysmic variables (CVs) with evolved donor stars (e.g. [Kalomeni et al. 2016](#)). When mass transfer ceases, the donor shrinks inside its Roche lobe and appears as a detached ELM WD. When and why mass transfer ceases is not fully understood, but is presumably linked to the end of efficient angular momentum loss via magnetic braking (e.g. [Sun & Arras 2018](#)).

Consistent with predictions of the mass transfer scenario, essentially all known ELM WDs are in binaries. Their companions are normal WDs and neutron stars ([Brown et al. 2020](#)). ELM WDs are rare: their local space density is about 5×10^{-5} that of ordinary WDs ([Brown et al. 2016](#)). Most of the ~ 100 spectroscopically confirmed ELM WDs were discovered in the last decade by the dedicated ELM survey ([Brown et al. 2010](#)), which performed a comprehensive magnitude-limited search beginning with SDSS color cuts. Known ELM WDs are relatively hot ($T_{\text{eff}} \gtrsim 8,000$ K) and compact ($\log g \gtrsim 5.5$). This is largely a consequence of the color cuts adopted by the ELM survey: cooler and more bloated proto-ELM WDs are expected to exist but are not easily distinguishable from main-sequence A and F stars based on their colors alone (e.g. [Pelisoli et al. 2018a,b, 2019](#)).

The mass-transferring progenitors of short-period ELM WDs differ from ordinary CVs in that the donor stars underwent significant nuclear evolution before the onset of mass transfer. In most short-period ($P_{\text{orb}} \lesssim 6$ hours) CVs, the mass-losing stars fall on a tight “donor sequence” of mass, radius, spectral type, and luminosity as a function of period ([Patterson 1984](#); [Beuermann et al. 1998](#); [Smith & Dhillon 1998](#); [Knigge 2006](#); [Knigge et al. 2011](#)). The regularity of the CV donor population is a consequence of the fact the donor stars fill their Roche lobes (so that their mean density is a nearly deterministic function of orbital period) and are only mildly out of thermal equilibrium. They thus have radii and effective temperatures similar to main-sequence stars of the same mass.

CVs donors that were significantly evolved at the onset of Roche lobe overflow (RLOF) are predicted to deviate from the donor sequence (e.g. [Tutukov et al. 1985](#); [Podsiadlowski et al. 2003](#); [van der Sluys et al. 2005](#); [Kalomeni et al. 2016](#)) because they have already

formed a helium core when mass transfer commences. By the time the orbital period reaches $P_{\text{orb}} \lesssim 6$ hours, such donors will consist of a low-mass ($M \lesssim 0.2M_{\odot}$) helium core with a thin hydrogen-burning envelope. The structure of such a donor is quite different from that of a main-sequence star of the same mean density. The most obvious observable difference is that the donor stars in CVs formed from evolved donors are expected to be hotter and smaller than normal CV donors with the same period. Several CVs with potentially evolved donors have been identified (e.g. [Augusteijn et al. 1996](#); [Thorstensen et al. 2002b,a](#); [Littlefair et al. 2006](#); [Thorstensen 2013](#); [Kato & Osaki 2014](#); [Rebassa-Mansergas et al. 2014](#); [Ashley et al. 2020](#)). Most of these systems have K-type donor stars, sometimes with abundances bearing imprints of CNO burning, at periods where normal CVs have M-type donors. If the donors are sufficiently evolved, they can become detached ELM WDs.

CVs with evolved donors are of interest for several reasons. Shortly after initial RLOF, their mass transfer rate is often high enough to allow for steady hydrogen burning on the surface of the accreting WD. This can form a bright supersoft X-ray source (e.g. [van den Heuvel et al. 1992](#)) and, if a high mass transfer rate is sustained long enough, trigger a type Ia supernova (e.g. [Han & Podsiadlowski 2004](#)). If magnetic braking remains efficient, evolved CVs can become mass-transferring ultra short-period ($5 \lesssim P_{\text{orb}}/\text{minute} \lesssim 30$) ‘‘AM-CVn’’ binaries containing a helium WD donor (e.g. [Podsiadlowski et al. 2003](#)). If magnetic braking becomes inefficient (e.g., because the donor loses its convective envelope), evolved CVs become ELM WDs, which can eventually become ultracompact via gravitational wave radiation (e.g. [Kilic et al. 2014](#); [Burdge et al. 2020](#)).

This paper presents a newly discovered close binary containing a normal white dwarf and a low-mass companion that is nearly or completely Roche lobe filling. The lower-mass component appears to be in a rarely observed transitional state between evolved CV donors and ELM WDs. It is hotter than any known CV donors, but cooler than any known ELM WDs. The optical spectrum is dominated by the low-mass component, which has an early F spectral type, and contains only hints of emission lines. No dwarf novae outbursts are detected; if they occur, their recurrence timescale would be unusually long. Because there is no strong evidence of ongoing mass transfer, the system either is detached or has an unusually low mass transfer rate. These unusual properties warrant detailed study, as the binary serves as a unique observational link between populations of CVs, ELM WDs, and ultracompact binaries.

The remainder of this paper is organized as follows. Section [11.3](#) summarizes the observed data, including archival spectra, light curves from several time-domain photometric surveys, astrometry from *Gaia*, the UV-to-near infrared spectral energy distribution (SED), and follow-up radial velocity observations. We constrain the physical parameters of the binary using these data in Section [11.4](#) and derive an upper limit on the outburst frequency in Section [11.5](#). In Section [11.6](#), we present binary evolution models for J0140 and discuss the future evolution of the system. We summarize our results and discuss future observations in Section [11.7](#). Supporting data is provided in the appendices.

11.3 Data

11.3.1 LAMOST discovery spectra

J0140 first came to our attention because LAMOST observations (Cui et al. 2012a) of the object suggested large epoch-to-epoch radial velocity (RV) variability. The object was visited by LAMOST three times, yielding two usable spectra (see Table 11.4). Between these two visits, the reported RV varied by $\approx 400 \text{ km s}^{-1}$.

Most LAMOST “visits” consist of 2-4 consecutive “subexposures”, with individual exposure times of 10-45 minutes. In the default LAMOST data reduction pipeline, these subexposures are coadded to improve SNR, without any velocity shift. The stellar parameters and RVs reported in most LAMOST catalogs are inferred from the coadded spectra. For close binaries, RVs may change significantly between subexposures, and the coadded spectra may have artificially broadened lines due to orbital smearing. We therefore investigated the RVs for J0140 across subexposures of each visit (Appendix 11.8 and Figure 11.15), finding RV variations of up to 600 km s^{-1} within a single visit. This prompted us to obtain additional follow-up spectra, as described in Section 11.3.5.

11.3.2 ZTF light curve

J0140 was observed regularly by the public survey of the Zwicky Transient Facility (ZTF; Bellm et al. 2019; Graham et al. 2019; Masci et al. 2019). We queried the public DR5 light curves for all clean photometry (`catflags = 0`). This yielded 524 *g*-band exposures, 576 *r*-band exposures, and 42 *i*-band exposures. The exposure time is 30 seconds in all bands. The ZTF data span 964 days between 2018 and 2021, sometimes including several observations on a single night.

To estimate the orbital period of J0140, we calculated *g*- and *r*- band periodograms from the ZTF light curves. Both were dominated by a peak at 1.904 hours. In binaries whose light curves are dominated by ellipsoidal variation, the dominant light curve period is typically half the orbital period, so we folded the light curves on twice this period (Figure 11.1). This revealed a light curve shape characteristic of ellipsoidal modulation: quasi-sinusoidal variability with alternation in the depth of adjacent minima. This alternation stems from the fact that a tidally distorted star is fainter when its cooler end faces the observer (phase 0 or 1.0) than when its hotter end does (phase 0.5).

11.3.3 Other photometry

J0140 was also observed by several other time-domain photometric surveys. Figure 11.2 shows all the photometric data we analyze, including both detections and upper limits between 2005 and 2021. The highest-quality photometry is from ZTF, but ZTF photometry only covers a small fraction of the total time baseline. Phased light curves from three surveys are shown in Figure 11.1.




Figure 11.1: ZTF, PTF, and CSS light curves of J0140, phased to a period of 3.81 hours. Each datapoint appears twice. The y -axis spans 0.55 mag in all panels. Variability is dominated by ellipsoidal modulation. Constraints on the peak-to-peak flux variability amplitude are listed in the upper left of each panel.




Figure 11.2: Light curve of J0140, showing data from several time-domain surveys. Points with errorbars are detections; inverted triangles are upper limits. CSS magnitudes are unfiltered but are scaled to roughly approximate the V band. The two apparent outbursts in the CSS light curve (in 2007 and 2011) are spurious; no reliable dwarf-novae like outbursts were detected between 2005 and 2021.

CSS

The Catalina Sky Survey (CSS; Drake et al. 2009) observed J0140 regularly between 2005 and 2013. CSS magnitudes are unfiltered but are transformed to V using an empirically calibrated, color-dependent relation. We queried all clean CSS photometry for J0140 with photometric uncertainty $\sigma_V < 0.12$ mag. This yielded 330 photometric points with a typical uncertainty of 0.09 mag.

The phased CSS light curve exhibits similar ellipsoidal modulation to the ZTF light curves (Figure 11.1). The full light curve (Figure 11.2) also contains two episodes, in 2007 and 2011, in which the object appears to brighten by 2-3 mag. These episodes at first appear similar to outbursts found in dwarf novae due to the disk instability (e.g. Osaki 1996). However, closer inspection revealed that within each episode, the unusually bright light curve points (with magnitude 15-16) were flanked on either side by observations with normal brightness (magnitude 17-18), with temporal separations of only a few minutes. Since dwarf nova outbursts typically last several days, this aroused suspicion that the apparent brightening might instead be due to problems with the photometry. We retrieved the raw CSS images for the 2011 episode from the CRTS web portal¹ and found that the brightening is indeed spurious. Images for the 2007 brightening episode are not available via the CRTS portal, but given their erratic time-evolution, we suspect this episode is also due to data problems and is not a dwarf nova outburst.

¹<https://crts.iucaa.in/CRTS/>

PTF

J0140 was also observed by the Palomar Transient Factor (PTF; [Law et al. 2009](#)). The light curve contains 86 good datapoints with a typical uncertainty of 0.04 mag, including 71 measurements in the g band and 15 in the R band. The exposure time is 60 seconds. The phased PTF light curve exhibits similar ellipsoidal modulation to the ZTF and CSS light curves (Figure 11.1). We limit our analysis to the g band data.

ASAS-SN

J0140 is regularly observed by the The All-Sky Automated Survey for Supernovae (ASAS-SN; [Kochanek et al. 2017](#)), but the object is generally too faint to be detected. We retrieved V - and g -band light curves from the sky-patrol web portal² on MJD 59232. The light curve contains 992 V band datapoints and 1030 g band datapoints, but the majority are upper limits (Figure 11.2). There are 12 V band detections and 43 g band detections, both with a median magnitude of 17.1. The brightest nominal detection has a g band magnitude of 16.41 ± 0.18 .

Most ASAS-SN exposures are collected in sets of three consecutive exposures, each with an exposure time of 90 seconds. In all the ASAS-SN detections, there was a detection in only one of the three exposures in the set, usually with a magnitude similar to the upper limits reported for the other two exposures. It is thus unlikely that any of the anomalously bright ASAS-SN detections are dwarf novae outbursts.

11.3.4 Light curve analysis

Variability amplitude

To constrain the photometric variability amplitude, we fit the photometry from each survey shown in Figure 11.1 with a 3-term Fourier model; i.e., modeling the light curve in normalized flux space as a sum of sines and cosines with periods P , $P/2$, and $P/3$, and the phases held fixed. In fitting these amplitudes, we also model a photometric scatter term that accounts for intrinsic scatter and/or underestimated photometric uncertainties. The resulting constraints on the peak-to-peak variability amplitude are shown in the upper left of each panel of Figure 11.1. Differences between the filters are subtle, but there is a slight trend of larger variability amplitude at shorter wavelengths. This could be due either to wavelength-dependence in the intrinsic variability amplitude of the star (i.e. due to wavelength-dependent gravity darkening), or to a wavelength-dependent flux ratio between the star and a non-variable component.

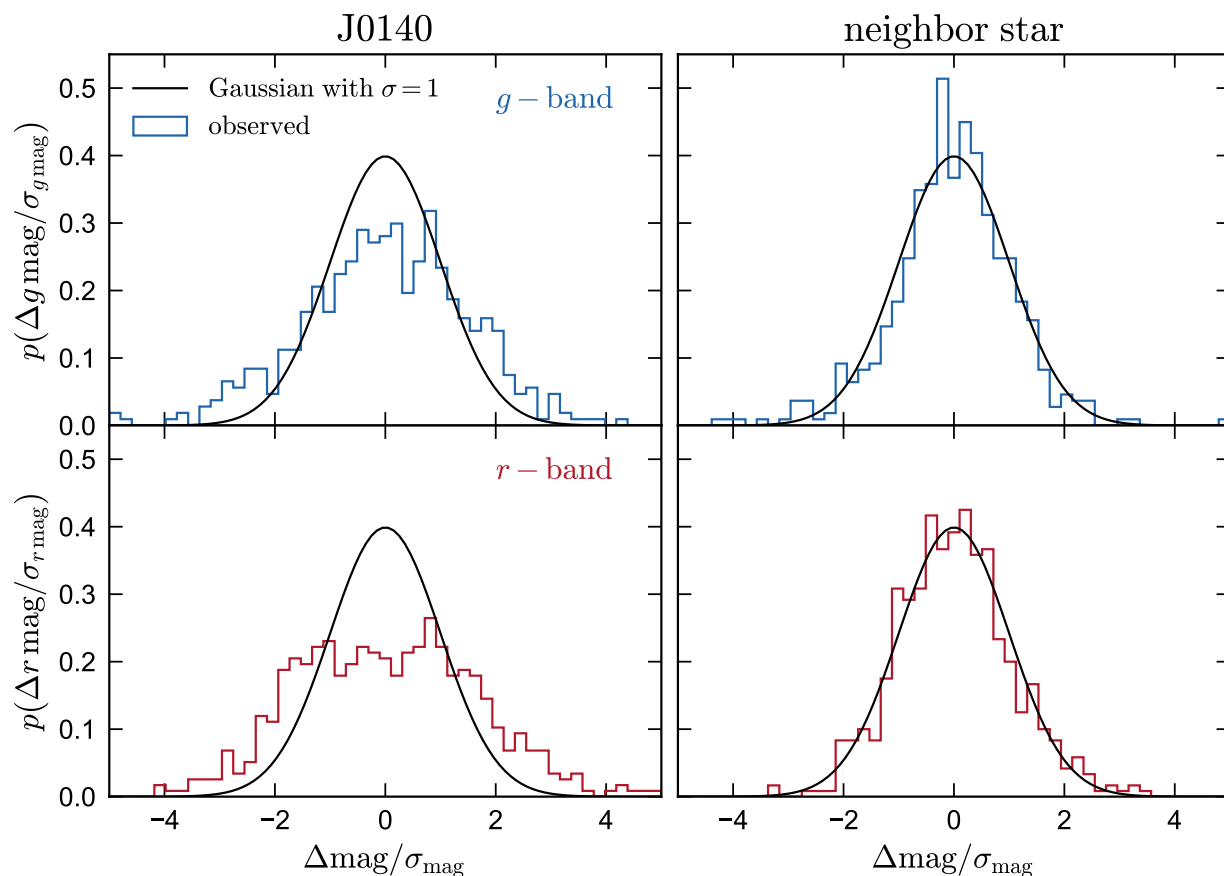


Figure 11.3: Difference between observed ZTF magnitudes and the mean magnitude at the corresponding orbital phase, normalized by the reported magnitude uncertainties. We compare distributions for J0140 (left) to those for a nearby comparison star with similar magnitude that is not significantly variable (right). For J0140, the distribution is broader than a Gaussian with $\sigma = 1$. This suggests that either (a) there is additional intrinsic scatter in the light curve, or (b) the magnitude uncertainties are underestimated. The well-behaved Gaussian distribution with $\sigma \approx 1$ for the neighbor star suggests that additional intrinsic scatter is the more likely explanation. This scatter is most likely due to pulsations of the proto-WD.

Photometric scatter

The ZTF g - and r -band light curves in Figure 11.1 appear to show somewhat more scatter than expected from the reported observational uncertainties. This is illustrated in Figure 11.3, which shows the uncertainty-normalized residuals between the observed magnitudes and the mean magnitude at their phase, which is calculated by binning the observed light curves in phase. If the magnitude were only a function of orbital phase, with no intrinsic scatter or dependence on other parameters, and the reported magnitude uncertainties were reliable, these residuals would be expected to follow a Gaussian distribution with $\sigma = 1$. The fact that the observed distribution is wider than this implies that either the magnitude uncertainties are underestimated, or there is additional variability.

The distribution of photometric residuals for a non-variable neighbor star with similar magnitude to J0140 is indeed well described by a Gaussian with $\sigma = 1$, suggesting that the reported uncertainties are reliable. We do not find strong evidence of any additional periodicity in the residuals; it appears to be stochastic. The origin of this additional variability is unknown. One possibility is that the proto-WD pulsates, as several ELM WDs are observed to (e.g. Hermes et al. 2012; Maxted et al. 2014). Indeed, the temperature and surface gravity of proto-WD (Section 11.3.5) place it within the pre-ELM WD instability strip (e.g. Córscico et al. 2016). It is also possible that the variability is due to “flickering” of an accretion disk around the massive WD, which is observed in almost all CVs (e.g. Bruch 1992) and is thought to be a result of turbulence in the disk.

To quantify the amplitude of the additional scatter, we fit the g - and r -band light curves with a Fourier series (see Section 11.3.4). Along with the Fourier amplitudes, we fit a scatter term, σ_{scatter} , which is added in quadrature to the observational uncertainties in the likelihood function. We find $\sigma_{\text{scatter}} = 2.4 \pm 0.1\%$ in the r band and $2.4 \pm 0.1\%$ in the g band. Higher-cadence photometry of the binary is required to better characterize this variability and verify whether it is due to pulsations.

Orbital ephemeris

We determined the orbital ephemeris by simultaneously fitting the normalized CSS, PTF, and ZTF light curves with a harmonic series. The long temporal baseline covers more than 35,000 orbits and thus allows for a high-precision determination of the orbital period (Table 11.1). The quoted conjunction time “ T_0 ” is the time corresponding to phase 0, when the companion (almost certainly a WD) passes in front of the proto-WD. When we analyze the radial velocities in Section 11.3.5, we find good agreement with the photometric ephemeris. This unambiguously establishes the photometric variability as tracing the orbital period of the binary.

We do not find evidence of orbital decay: fitting the full light curve and allowing for a time-dependent phase shift, we find $\dot{P}_{\text{orb}} = -4_{-3}^{+5} \times 10^{-11} \text{ s s}^{-1}$, consistent with 0. This non-detection is consistent with the theoretically expected value from gravitational waves for the

²<https://asas-sn.osu.edu/>

parameters derived in Section 11.4, which is $\dot{P}_{\text{orb,GW}} = -7.7_{-6.4}^{+3.4} \times 10^{-14} \text{ s s}^{-1}$. If magnetic braking remains efficient (see Section 11.6.1), the orbital decay due to it is expected to be significantly faster than this, $\dot{P}_{\text{orb,MB}} \approx -2.5 \times 10^{-12} \text{ s s}^{-1}$, still consistent with the non-detection.

11.3.5 KAST Spectra

We observed J0140 using the KAST double spectrograph on the 3m Shane telescope at Lick observatory in July and August 2020. Most of our observations used the 600/7500 grating on the red side and the 600/4310 grism on the blue side, with the D55 dichroic and a 2 arcsec slit. This setup results in a wavelength range of 3300–8400 Å with typical resolution (FWHM) of 4 Å on the blue side and 5 Å on the red side. Some additional spectra were obtained with a separate setup (“B”). For these observations, we used the 1200/5000 grating on the red side, the 830/3460 grism on the blue side, the D46 dichroic, and a 1 arcsec slit. This setup results in a wavelength range of 3300–7200 Å, with a gap in coverage at 4600–5400 Å. The typical resolution for this setup is 3 Å on the blue side and 2 Å on the red side.

To minimize effects of flexure on the wavelength solution, a new set of arcs was taken on-sky while tracking the target every ~ 30 minutes. The spectra were reduced using `pypeit` (Prochaska et al. 2020), which performs bias and flat field correction, cosmic ray removal, wavelength calibration, flexure correction using sky lines, sky subtraction, extraction of 1d spectra, and heliocentric RV corrections.

To balance the need to minimize orbital smearing with the need to reach high enough SNR to measure RVs at each epoch, we used 600 second exposures. This resulted in a typical per-pixel SNR of 10-15 on the red side and 5-10 on the blue side. The spectra are summarized in Table 11.5. The quoted SNR is the median per pixel on the red side; the typical SNR on the blue side is a factor of ~ 2 lower. Variations in SNR between visits are due primarily to variations in sky brightness and airmass.

Radial velocities

Comparison of the single-epoch spectra with a grid of ATLAS/SYNTHÉ spectral models (Kurucz 1970, 1979, 1993) suggested an effective temperature of $T_{\text{eff}} \approx 6500 \text{ K}$ and a surface gravity $\log(g/\text{cm s}^{-1}) \approx 4.5$. We used a model spectrum with these parameters and $[\text{Fe}/\text{H}] = 0$ as a template to estimate the RVs of each single-epoch spectrum via cross-correlation.³ RVs were measured from the red-side spectra, which have a more stable wavelength solution due to a larger number of sky emission lines. The most prominent spectral feature on the red side is the $\text{H}\alpha$ absorption line.

We fit an orbital solution to the RVs (which presumably trace the proto-WD) while holding the period fixed to the value measured from the light curve. Along with the 6

³Once we derived more accurate atmospheric parameters for the donor from the co-added spectrum (Section 11.3.5), we updated the template and re-measured RVs.




Figure 11.4: Radial velocity of the proto-WD, phased to a period of 3.81 hours. RVs are measured from Kast spectra (Table 11.5), primarily from the $H\alpha$ absorption line. Horizontal error bars show the extent of each 10 minute exposure; vertical error bars are 1σ measurement uncertainties. Red line shows the best-fit Keplerian orbit (Table 11.1).

Table 11.1: Basic observables of J0140. Uncertainties are 1σ (middle 68%). Additional constraints are in Table 11.3.

System parameters		
Right ascension (J2000)	α [deg]	14.065468
Declination (J2000)	δ [deg]	39.447740
Apparent magnitude	G	17.35
Color	$G_{BP} - G_{RP}$	0.63
<i>Gaia</i> eDR3 parallax	ϖ [mas]	0.65 ± 0.10
Zeropoint-corrected parallax	ϖ [mas]	0.68 ± 0.10
Reddening (SFD; to infinity)	$E(B - V)$	0.056
Reddening (Bayestar2019)	$E(g - r)$	0.04 ± 0.02
Parameters of the proto-WD		
Effective temperature (spectrum)	T_{eff} [K]	6800 ± 100
Surface gravity	$\log(g/(\text{cm s}^{-2}))$	4.7 ± 0.3
Metallicity	[Fe/H]	-0.1 ± 0.1
Projected rotation velocity	$v \sin i$ [km s $^{-1}$]	105 ± 20
Angular diameter	Θ [μas]	2.05 ± 0.05
Effective temperature (SED)	T_{eff} [K]	6820 ± 110
Parameters of the WD		
Angular Diameter	Θ [μas]	0.055 ± 0.013
Effective temperature	T_{eff} [K]	19700 ± 1500
Orbital parameters		
Period	P_{orb} [days]	0.1586339(2)
Conjunction time	T_0 [HMJD UTC]	59060.490(1)
RV semi-amplitude	K [km s $^{-1}$]	330 ± 5
Center-of-mass RV	γ [km s $^{-1}$]	$-53 \pm 4 \pm 30$ (sys)
eccentricity	e	< 0.017
RV scatter	s [km s $^{-1}$]	18 ± 5
Mass function	f_m [M_{\odot}]	0.591 ± 0.027

standard Keplerian orbital parameters, we fit for an intrinsic RV scatter or “jitter” term (see [El-Badry et al. 2018g](#), their Equation 9), which can represent either intrinsic RV variability not captured in the model or underestimated uncertainties. We first search for the best-fit orbital solution using simulated annealing, and then sample from the posterior in the vicinity of this solution using a Markov chain Monte Carlo method, as described in [El-Badry et al. \(2018g\)](#). Our best-fit orbital parameters are listed in Table 11.1. The eccentricity is consistent with 0. The mass function, $f_m = PK_{\text{protoWD}}^3(1 - e^2)^{3/2}/(2\pi G)$, which represents the minimum possible mass of the companion, i.e., the mass it would have if the proto-WD were a test particle and the system were viewed edge on, is consistent with a white dwarf companion. This is illustrated in Figure 11.5, which shows the dynamically implied masses of the companion for a range of proto-WD masses and inclinations.

The best-fit jitter term is $\approx 18 \text{ km s}^{-1}$. Nonzero jitter is evident in Figure 11.4, where the scatter in RV measurements is somewhat larger than the 1σ RV uncertainties suggest it should be. This jitter is most likely due to instability in the wavelength solution; it could also originate in part from the ambiguity in the definition of RVs for an ellipsoidally distorted star (e.g. [Masuda & Hirano 2021](#)), from irradiation effects (e.g. [Warner 2003](#)), or from pulsations.

Proto-WD spectrum

We reconstructed a higher-SNR spectrum of the proto-WD by coadding spectra from all epochs, shifted to the star’s rest frame. This yielded a combined spectrum with maximum SNR ≈ 60 per pixel and mean spectral resolution $R \approx 1,200$ (FWHM). We fit for the atmospheric parameters, projected rotation velocity, and metallicity of the donor using standard full spectral fitting based on a forward-model of ab-initio spectra. Rotational broadening was accounted for using the kernel from [Gray \(1992b\)](#).

We implemented *The Payne* ([Rix et al. 2016](#); [El-Badry et al. 2018g](#); [Ting et al. 2019](#)), a framework for interpolating model spectra, to predict the pseudo-continuum normalized spectra of the proto-WD. We use a 2nd order polynomial spectral model with three labels,⁴ $\ell = (T_{\text{eff}}, \log g, [\text{Fe}/\text{H}])$, which we trained on the BOSZ grid of Kurucz model spectra ([Bohlin et al. 2017](#)). In the relevant part of parameter space, the grid spacing is $\Delta T_{\text{eff}} = 250 \text{ K}$, $\Delta \log g = 0.5 \text{ dex}$, $\Delta [\text{Fe}/\text{H}] = 0.25 \text{ dex}$. For both the models and the observed spectrum, we define the pseudo-continuum as a running median with width 150 \AA . Low-resolution spectra are subject to severe line blending, so this pseudo-continuum has little to do with the true continuum. The purpose of pseudo-continuum normalization is simply to bring the observed and model spectra to the same scale, minimizing uncertainty due to unknown detector efficiency and distance. This fitting yields an effective temperature for the proto-WD of $T_{\text{eff}} \approx 6800 \text{ K}$, surface gravity $\log(g/[\text{cm s}^{-2}]) \approx 4.7$, and metallicity $[\text{Fe}/\text{H}] \approx -0.1$ (see Table 11.1).

The formal fitting uncertainties on spectroscopic labels are very small: 25 K in T_{eff} ,

⁴We also explored the possibility of determining the proto-WD’s detailed abundance pattern, since constraining the abundance of C, N, and O would provide a useful test of the evolutionary scenario proposed in Section 11.6. We found lines from these elements to be too blended for reliable abundance measurements.




Figure 11.5: Dynamically implied mass of the companion as a function of the proto WD’s mass and orbital inclination, given the observed mass function, $f_m \approx 0.59M_\odot$. Red contours show 1 and 2 sigma constraints from joint fitting of the light curve, spectra, orbit, SED, and parallax (Section 11.4). Given the observed UV excess (Figure 11.7), the $\approx 0.9M_\odot$ companion is very likely a white dwarf.

0.06 dex in $\log g$, and 0.03 dex in $[\text{Fe}/\text{H}]$. Except at very low-SNR, formal uncertainties derived from full-spectrum fitting underestimate the true parameter uncertainties (e.g. [Ness et al. 2015b](#); [Xiang et al. 2019](#)). This owes to errors in pseudo-continuum normalization, simplifications in the spectral model (e.g., fitting only 3 labels) and systematics in the model spectra on which the models are trained. We therefore inflate the uncertainties (Table 11.1) to realistic values for the resolutions and SNR of the spectrum (e.g. [Xiang et al. 2019](#)).

The strongest features in the proto-WD spectrum are the Balmer and calcium H and K lines. Other metal lines are also present, but they are subject to significant blending at low resolution. The agreement between the observed and model spectra is satisfactory for most lines. Some low-frequency undulations are evident in the residual spectrum; these likely are due to imperfect normalization. The most significant feature in the residuals is in the $\text{H}\alpha$ absorption line, which is weaker in the observed spectrum than in the model. This may be the result of contamination from weak $\text{H}\alpha$ emission originating in accretion onto the more massive WD: as we show in Section 11.4, the proto-WD is extremely close to Roche lobe-filling, so ongoing mass transfer is not unlikely. The hints of emission are, however, much weaker than found in typical CVs (e.g. [Warner 2003](#)), and there is no evidence that a disk contributes significantly to the optical SED (Section 11.3.7). If there is ongoing mass transfer, the accretion rate is likely much lower than in typical CVs at similar periods.

The observed spectrum also has significantly stronger sodium “D” absorption lines (at 5890 and 5896 Å) than the model. We observe these lines to shift coherently with the orbit of the proto-WD and can thus rule out interstellar absorption as their source. Sodium enhancement is frequently observed in evolved CV donors (e.g. [Thorstensen et al. 2002b](#)) and is thought to originate from CNO-processing in the core of the secondary near the end of its main-sequence evolution. In this scenario, sodium-enhanced material only reaches the surface of the secondary after most of the envelope has been stripped off. A similar evolutionary history can likely explain the formation of J0140 (Section 11.6).

Variability with orbital phase

We also investigated whether the spectra vary with orbital phase (Appendix 11.10). This reveals weak evidence of increased $\text{H}\alpha$ absorption at $\phi = 0.5$, as might be expected if a $\text{H}\alpha$ -emitting component were partially or fully eclipsed, or if the WD-facing side of the proto-WD were irradiated by the more massive WD. However, the apparent increase in the line’s depth is only $\approx 10\%$. Further spectroscopic follow-up with higher SNR is required to confirm whether it is real.

11.3.6 Distance

The *Gaia* eDR3 parallax ([Gaia Collaboration et al. 2020a](#)) of J0140 (Gaia DR3 source id 373857386785825408) is relatively significant, $\varpi = 0.65 \pm 0.10$. Interpreted as an inverse distance, this places J0140 at a distance of $d \approx 1.54$ kpc (1-sigma range of 1.33 to 1.84 kpc). The parallax zeropoint predicted by the fitting function from [Lindgren et al. \(2020a\)](#) for




Figure 11.6: Rest-frame normalized spectrum of the proto-WD (black) and best-fit Kurucz model spectrum (red). The strongest absorption features are the Balmer lines and Ca H & K lines. Bottom panel shows residuals. A satisfactory fit is achieved with $T_{\text{eff}} = 6800 \text{ K}$, $\log = 4.7$, and $[\text{Fe}/\text{H}] = -0.1$. Most of the structure in the residuals is broad and due to imperfect normalization. The most significant narrow feature in the residuals is in the H α line. This could arise from emission from a faint accretion disk around the WD, but may also be the result of systematics in the spectral model.

J0140 is -0.03 mas . At $G > 17$, *Gaia* eDR3 parallax uncertainties for blue sources are not overestimated on average (El-Badry et al. 2021). We thus adopt $\varpi = (0.68 \pm 0.1)$ in modeling the SED. The geometric distance from Bailer-Jones et al. (2020), which is based on the zeropoint-corrected parallax and a Milky Way model prior, is $d = 1.51^{+0.25}_{-0.18} \text{ kpc}$.

The *Gaia* eDR3 astrometric solution does not account for photocenter wobble induced by orbital motion, but this is expected to be negligible: the ratio of the projected semimajor axis to the parallax, $\theta/\varpi = a/(1\text{au})$, is only about 0.005. The renormalized unit weight error (RUWE) is 0.97, indicating an unproblematic astrometric solution.

11.3.7 Spectral energy distribution

The spectral energy distribution (SED) of J0140 provides a constraint on the temperature and angular diameter of both components. Table 11.2 summarizes the available broadband fluxes, which are taken from the GALEX (Morrissey et al. 2007), SDSS (York et al. 2000), 2MASS (Skrutskie et al. 2006), and WISE (Wright et al. 2010) surveys. Because J0140 exhibits ellipsoidal variability with amplitude of $\approx 0.3 \text{ mag}$ in the optical, we “correct” the reported magnitudes to better approximate their phase-averaged mean values. We use a Fourier model of the ZTF r -band light curve (Figure 11.1) to calculate a “optical ellipsoidal

Table 11.2: Spectral energy distribution. Magnitude lower limits are 2σ . mjd values are mid-exposure. The WISE magnitudes represent a mean from 33 exposures spread roughly uniformly in phase. The “corrected magnitude” column corrects the observed magnitudes, which occur at different orbital phases, to represent the phase-averaged mean magnitude. This column also corrects for discrepancies between the SDSS and AB magnitude systems (Section 11.3.7).

filter	$\lambda_{\text{eff}} [\mu\text{m}]$	mag	mjd	phase	optical ellipsoidal correction [mag]	corrected mag
FUV	0.154	21.90 ± 0.13	53682.4345	0.70 ± 0.04	0.11 ± 0.01	21.90 ± 0.13
NUV	0.227	20.55 ± 0.04	53682.4345	0.70 ± 0.04	0.11 ± 0.01	20.63 ± 0.08
<i>u</i>	0.359	18.72 ± 0.02	54010.4614	0.52 ± 0.04	-0.09 ± 0.03	18.59 ± 0.05
<i>g</i>	0.464	17.64 ± 0.02	54010.4630	0.53 ± 0.04	-0.08 ± 0.03	17.56 ± 0.04
<i>r</i>	0.612	17.40 ± 0.01	54010.4597	0.51 ± 0.04	-0.09 ± 0.02	17.31 ± 0.03
<i>i</i>	0.744	17.34 ± 0.01	54010.4605	0.52 ± 0.04	-0.09 ± 0.02	17.27 ± 0.03
<i>z</i>	0.890	17.34 ± 0.02	54010.4622	0.53 ± 0.04	-0.08 ± 0.03	17.29 ± 0.04
<i>J</i>	1.228	16.47 ± 0.13	51455.2390	0.86 ± 0.06	0.02 ± 0.09	16.49 ± 0.16
<i>H</i>	1.639	16.10 ± 0.22	51455.2390	0.86 ± 0.06	0.02 ± 0.09	16.12 ± 0.24
<i>K_s</i>	2.152	> 15.63	51455.2390	0.86 ± 0.06	0.02 ± 0.09	> 15.65
<i>W₁</i>	3.292	16.25 ± 0.06	~ 55400	–	–	16.25 ± 0.06
<i>W₂</i>	4.542	16.71 ± 0.25	~ 55400	–	–	16.71 ± 0.25
<i>W₃</i>	10.302	> 12.50	~ 55400	–	–	> 12.50
<i>W₄</i>	21.810	> 9.15	~ 55400	–	–	> 9.15

correction” (Table 11.2) that converts the magnitude observed at a given phase to the mean magnitude in that band.

This correction is complicated somewhat by the fact that the variability amplitude is likely wavelength dependent. Based on our fit to the spectrum, the phased light curves in Figure 11.1, and an initial fit to the SED without phase corrections, we estimate that the *r*-band variability amplitude is an adequate approximation of the variability in the optical and near-infrared bands. However, the more massive WD likely dominates the light in the UV, so we estimate that the variability amplitude is reduced by 30% in the NUV and is negligible in the FUV. The “corrected mag” column of Table 11.2 lists the magnitudes we assume in modeling the SED. In addition to corrections due to phase-variability, we also apply the empirical corrections from Eisenstein et al. (2006) to bring SDSS magnitudes to the AB magnitude system: $u_{\text{AB}} = u_{\text{SDSS}} - 0.04$, $i_{\text{AB}} = i_{\text{SDSS}} + 0.015$, and $z_{\text{AB}} = z_{\text{SDSS}} + 0.03$.

X-ray upper limit

No X-ray source has been detected at the location of J0140, but the ROSAT survey (Voges et al. 1999) obtained a 493 second exposure covering the object’s coordinates. Using the ESA upper limit server,⁵ we obtain a 3-sigma upper limit of $2.5 \times 10^{-13} \text{erg s}^{-1} \text{cm}^{-2}$ in the 0.2-2 keV range. At a distance $d = 1.5 \text{kpc}$ (Section 11.4), this corresponds to a luminosity $L_X \leq 0.018 L_\odot$ in the 0.2-2 keV range.

⁵See <http://xmmuls.esac.esa.int/upperlimitsserver/>. We model the spectrum as a power law with slope 2 and assume an absorption column density $N_{\text{H}} = 3 \times 10^{20} \text{cm}^{-2}$.




Figure 11.7: Spectral energy distribution of J0140. Red points with error bars show observed broadband fluxes from GALEX, SDSS, 2MASS, and WISE. Cyan shows a model spectrum for the proto-WD with $T_{\text{eff}} = 6825 \text{ K}$, $\log(g/\text{cm s}^{-2}) = 4.7$, and angular diameter $\Theta = 2.05 \mu\text{as}$. Blue line shows a model for the more massive WD. Gray and black lines show the total model spectrum with and without extinction. Most of the SED is dominated by the proto-WD. However, it contributes negligibly in the FUV, where the SED is likely dominated by the more massive WD. Bottom panel shows the difference between the observed bandpass-integrated magnitudes and those predicted by the total SED model.

Angular diameter and radius of the donor

We fit the SED for the effective temperatures and angular diameters, $\Theta = 2R/d$, of both components, as well as the reddening, $E(B - V)$. We predict the bandpass-averaged magnitudes of the proto-WD using synthetic photometry from MIST models (Choi et al. 2016). We fix $\log g = 4.7$ (similar to the value implied by our fit in Section 11.4) since varying $\log g$ at fixed angular diameter has minimal effects on the SED. We calculate synthetic photometry for the proto-WD using atmospheric models from Koester (2010) with $\log g = 8.5$. We assume a Cardelli et al. (1989) extinction law with total-to-selective extinction ratio $R_V = 3.1$.

The temperature and angular diameter of the proto-WD are independently well-constrained by the SED: the wavelength at which the SED peaks constrains T_{eff} , and the normalization in the Rayleigh–Jeans tail then constrains Θ . For the more massive WD, T_{eff} and Θ are highly covariant, because the peak of its SED is likely blueward of the bluest bandpass observed. To break this degeneracy and constrain the more massive WD’s temperature, we adopt a prior on its angular diameter, $\Theta_{\text{WD}}/\mu\text{as} \sim \mathcal{N}(0.055, 0.013)$. This is obtained from the *Gaia* distance and our dynamical constraint on the WD’s mass (Section 11.4), assuming the mass-radius relation from Bédard et al. (2020). We adopt a prior of $E(B - V) \sim \mathcal{N}(0.04, 0.02)$ based on the Green et al. (2019) 3D dust map, and flat priors on $T_{\text{eff, proto WD}}$, $\Theta_{\text{proto WD}}$, and $T_{\text{eff, WD}}$.

We sample from the posterior using *emcee* (Foreman-Mackey et al. 2013c), monitoring chains for convergence. The resulting constraints on the effective temperatures and angular diameters of both components are listed in Table 11.1, and the best-fit SED model is compared to the data in Figure 11.7. When combined with the *Gaia* parallax, the constraint on the angular diameter of the proto-WD translates to a constraint on its radius, $R_{\text{proto WD}} = (0.33 \pm 0.09) R_{\odot}$.

If the proto-WD is assumed to fill its Roche lobe, as the light curve suggests it very nearly does (Section 11.4), then its mean density $\bar{\rho} = 3M/(4\pi R^3)$ is given by

$$\bar{\rho} \approx 107 \text{ g cm}^{-3} (P_{\text{orb}}/\text{hr})^{-2}, \quad (11.1)$$

where M and R represent the mass and equivalent radius of the Roche-filling star (Eggleton 1983). The radius constraint from the SED then implies a proto-WD mass of $M_{\text{proto WD}} = 0.18^{+0.21}_{-0.11} M_{\odot}$.

The observed UV-to-IR SED is reproduced reasonably well. The SED of the proto-WD dominates in the optical and IR but drops off precipitously in the FUV. Under the assumption that the FUV emission is dominated by the more massive WD (without, for example, a significant contribution from accretion), we find a WD effective temperature $T_{\text{eff, WD}} = 19700 \pm 1500 \text{ K}$.

Under the assumption that the WD is accreting and has reached an equilibrium temperature for its accretion rate, its effective temperature can be related to the mass transfer rate (e.g. Townsley & Gänsicke 2009, their Equation 2). Taking $M_{\text{WD}} = 0.95 M_{\odot}$ (Section 11.4),

the implied mass transfer rate is

$$\dot{M} \lesssim 1.5 \times 10^{-10} M_{\odot} \text{ yr}^{-1}. \quad (11.2)$$

This is significantly lower than is found and predicted for most CVs at $P \approx 4$ hours, where most non-magnetic systems have $8 \lesssim -\log(\dot{M}/(M_{\odot} \text{ yr}^{-1})) \lesssim 9$. (e.g. Townsley & Gänsicke 2009; Kalomeni et al. 2016). Equation 11.2 is an upper limit, since the WD could remain hot from earlier accretion after a drop in the mass transfer rate.

11.4 Parameter constraints

To constrain the physical parameters of the binary, we combine constraints from the ZTF light curves (Figure 11.1), Kast spectrum (Figure 11.6), multi-epoch RVs (Figure 11.4), photometric SED (Figure 11.7), and the *Gaia* parallax.

11.4.1 Light curve model

We model the ZTF g - and r -band light curves simultaneously, using the binary light curve modeling software *Phoebe* (Prša & Zwitter 2005; Prša et al. 2016; Horvat et al. 2018), which builds on the Wilson & Devinney (1971) code. *Phoebe* models the surfaces of both components using a deformable mesh, self-consistently accounting for ellipsoidal variation, reflection, and eclipses. The atmosphere of the donor is modeled using *Phoenix* model atmospheres (Husser et al. 2013) with $T_{\text{eff}} = 6800$ K, and its limb darkening coefficients are calculated self-consistently from the atmosphere models. The WD, which contributes about 1% of the light in the optical, is modeled with a blackbody atmosphere. We assume it follows the mass-radius relation from Bédard et al. (2020) and has an effective temperature of 20,000 K. We fix the eccentricity to 0.

We leave the gravity darkening exponent β_1 , where $T_{\text{eff}}^4 \propto g^{\beta_1}$ (von Zeipel 1924), as a free parameter. It is common to assume $\beta_1 = 0.32$ for stars with significant convective envelopes ($T_{\text{eff}} \lesssim 6300$ K), and $\beta_1 = 1$ for stars with radiative envelopes ($T_{\text{eff}} \gtrsim 8000$ K; e.g., Lucy 1967). J0140 is likely between these two regimes, with a thin convective envelope. Observational determinations of β_1 in this temperature range are typically $0.2 \lesssim \beta_1 \lesssim 0.6$ (e.g. Claret 2015). We adopt a prior $\beta_1 \sim \mathcal{N}(0.32, 0.1)$.

The light curve provides a joint constraint on the inclination and Roche lobe filling factor, R_{protoWD}/R_L . We simultaneously fit the g - and r -band light curves for these parameters,⁶ as well as the intrinsic scatter in the light curves. Like the “jitter” term used in fitting the RVs, the intrinsic scatter term is added in quadrature to the reported photometric uncertainties;

⁶In practice, this is accomplished by fixing the masses of the two components and leaving the radius of the proto-WD free. This is a sensible choice because, in the absence of detectable eclipses, the light curve shape depends just on the Roche lobe filling factor (which depends on the star’s mean density) and on the limb darkening and gravity darkening coefficients, not on the absolute masses and radii.

it essentially forces the reduced χ^2 of the light curve fit to be 1. We sample the posterior with `emcee`; our priors and constraints are listed in the first section of Table 11.3.

Figure 11.8 compares our best-fit `Phoebe` models to the observed light curves. For illustrative purposes, we show the best-fit models when the inclination is fixed at three different values. For all inclinations $i \gtrsim 70$ degrees, equally good fits to the light curve can be obtained by decreasing the Roche lobe filling factor while increasing the inclination. For models with $i \geq 75$ degrees, the eclipse of the massive WD is visible at $\phi = 0.5$. However, the depth of the predicted eclipse is less than 1% in both bands, and the observed light curves cannot confidently distinguish between models that do and do not include eclipses. For inclinations $i < 70$ degrees, the predicted amplitude of ellipsoidal variability is weaker than what is observed, even when the proto-WD completely fills its Roche lobe. The blue curve in Figure 11.8 shows the best-fit model with $i = 65$ deg, which is confidently ruled out by the data.

The `Phoebe` models provide reasonable but not perfect fits to the light curve. Particularly in the r - band, there are coherent residuals of ≈ 0.02 mag at phases $\phi = 0.75 - 1.25$. This could be the result of systematics in the model light curves (e.g., an inaccurate limb darkening law) or additional components not included in the model (e.g. an accretion disk or hot spot around the more massive WD, or spots on the surface of the proto-WD). Because the model does not perfectly describe the data, the formal uncertainties on light curve fitting parameters may be underestimated. We therefore do not fit the light curve and other observables simultaneously. Instead, we adopt priors on the inclination and Roche lobe filling factor of the proto-WD that are based on the light curve in our combined model, as described below.

11.4.2 Combined model

We now combine constraints on the angular diameter, parallax, orbital velocity, inclination, and Roche lobe filling factor. The free parameters of the fit are the mass and radius of the proto-WD, orbital inclination, distance, and the mass of the more massive WD. The full likelihood function is:

$$\begin{aligned} \ln L = & -\frac{1}{2} \frac{[(R_{\text{proto WD}}/R_L)_{\text{model}} - (R_{\text{proto WD}}/R_L)_{\text{obs}}]^2}{\sigma_{R_{\text{proto WD}}/R_L}^2} \\ & -\frac{1}{2} \frac{(K_{\text{model}} - K_{\text{obs}})^2}{\sigma_{K,\text{obs}}^2} - \frac{1}{2} \frac{(\Theta_{\text{model}} - \Theta_{\text{obs}})^2}{\sigma_{\Theta,\text{obs}}^2} \\ & -\frac{1}{2} \frac{(\varpi_{\text{model}} - \varpi_{\text{obs}})^2}{\sigma_{\varpi,\text{obs}}^2} \end{aligned}$$

The four terms compare the proto-WD's predicted Roche lobe filling factor, RV semi-amplitude, and angular diameter, and the system parallax, to the observed constraints.




Figure 11.8: Best-fit *Phoebe* *g* (top) and *r* (bottom) light curves when the inclination is fixed to 60, 75, and 85 degrees. Models are compared to the phase-averaged ZTF data, with error bars including only random photometric uncertainties. Models are labeled by the ratio of the proto-WD equivalent radius to the equivalent Roche lobe radius. The best-fit models with $i = 75$ and $i = 85$ deg are nearly indistinguishable. Models with $i \gtrsim 75$ deg include a shallow eclipse at phase 0.5, when the massive WD is eclipsed, but the observed light curve is as yet too noisy to confirm or rule out this feature. The model with $i = 65$ deg produces too-low ellipsoidal variability amplitude, even when the proto-WD fully fills its Roche lobe. Our fit to the light curve yields $i > 70$ deg and $R_{\text{proto WD}}/R_L > 0.98$.

The proto-WD's predicted Roche lobe equivalent radius is

$$R_L \approx \frac{0.49q^{2/3}}{0.6q^{2/3} + \ln(1 + q^{1/3})} \left(\frac{P_{\text{orb}}^2 G (M_{\text{proto WD}} + M_{\text{WD}})}{4\pi^2} \right)^{1/3}, \quad (11.3)$$

where $q = M_{\text{proto WD}}/M_{\text{WD}}$, and we use the fitting function from [Eggleton \(1983\)](#). From the light curve fit, we take $(R_{\text{proto WD}}/R_L)_{\text{obs}} = 0.99$ and $\sigma_{R_{\text{proto WD}}/R_L}^2 = 0.01$.

The predicted RV semi-amplitude of the proto-WD is

$$K_{\text{model}} = \frac{[2\pi G (M_{\text{proto WD}} + M_{\text{WD}}) / P_{\text{orb}}]^{1/3}}{(1 + M_{\text{proto WD}}/M_{\text{WD}})} \sin i, \quad (11.4)$$

the model parallax is $\varpi_{\text{model}} = 1/d$, and the angular diameter is $\Theta_{\text{model}} = 2R/d$. The observational constraints on these parameters are taken from [Table 11.1](#).

Our adopted priors are listed in [Table 11.3](#). We use flat priors on all parameters except distance. We take the distance prior from [Bailer-Jones et al. \(2020\)](#):

$$P(d) = \begin{cases} \frac{1}{\Gamma(\frac{\beta+1}{\alpha})} \frac{\alpha}{L^{\beta+1}} d^\beta e^{-(d/L)^\alpha}, & \text{if } d \geq 0 \\ 0, & \text{otherwise} \end{cases} \quad (11.5)$$

Here L , α , and β are free parameters that depend on sky position and are provided by [Bailer-Jones et al. \(2020\)](#). They were calculated by fitting the distance distribution of sources in a mock *Gaia* eDR3 catalog ([Rybizki et al. 2020](#)). For the location of J0140, $L = 4.2 \times 10^{-4}$ pc, $\alpha = 0.24$, and $\beta = 7.2$. The effects of the prior on our final constraints are weak; adopting a flat prior instead changes the best-fit distance and donor radius by only a few percent.

We sample the posterior distribution using [emcee \(Foreman-Mackey et al. 2013c\)](#). We draw 4×10^5 samples using 64 walkers, after a burn-in period of 2000 steps per walker. We inspect the chains to verify convergence. The resulting constraints are listed in the lower section of [Table 11.3](#) and shown in [Figure 11.9](#). The proto-WD luminosity implied by our constraints on its radius and effective temperature is $L_{\text{proto WD}} = 0.23_{-0.06}^{+0.09} L_\odot$. The surface gravity implied by the mass and radius constraints is $\log [g / (\text{cm s}^{-2})] = 4.74 \pm 0.07$, consistent with the spectroscopic fit. We compare these parameter constraints to models in [Section 11.6](#).

11.5 Outburst frequency limits

If J0140 is still undergoing mass transfer, one might expect it to periodically exhibit dwarf nova-like outbursts, which usually occur in CVs due to disk instability (e.g. [Hameury 2020](#)). The light curve of J0140 does not show any convincing evidence of such outbursts ([Figure 11.2](#)). Here we use this fact to constrain an upper limit on the outburst frequency.

To determine whether an outburst would have been detected, we require an estimate of the outburst duration and amplitude. For normal CVs, the duration of typical outbursts

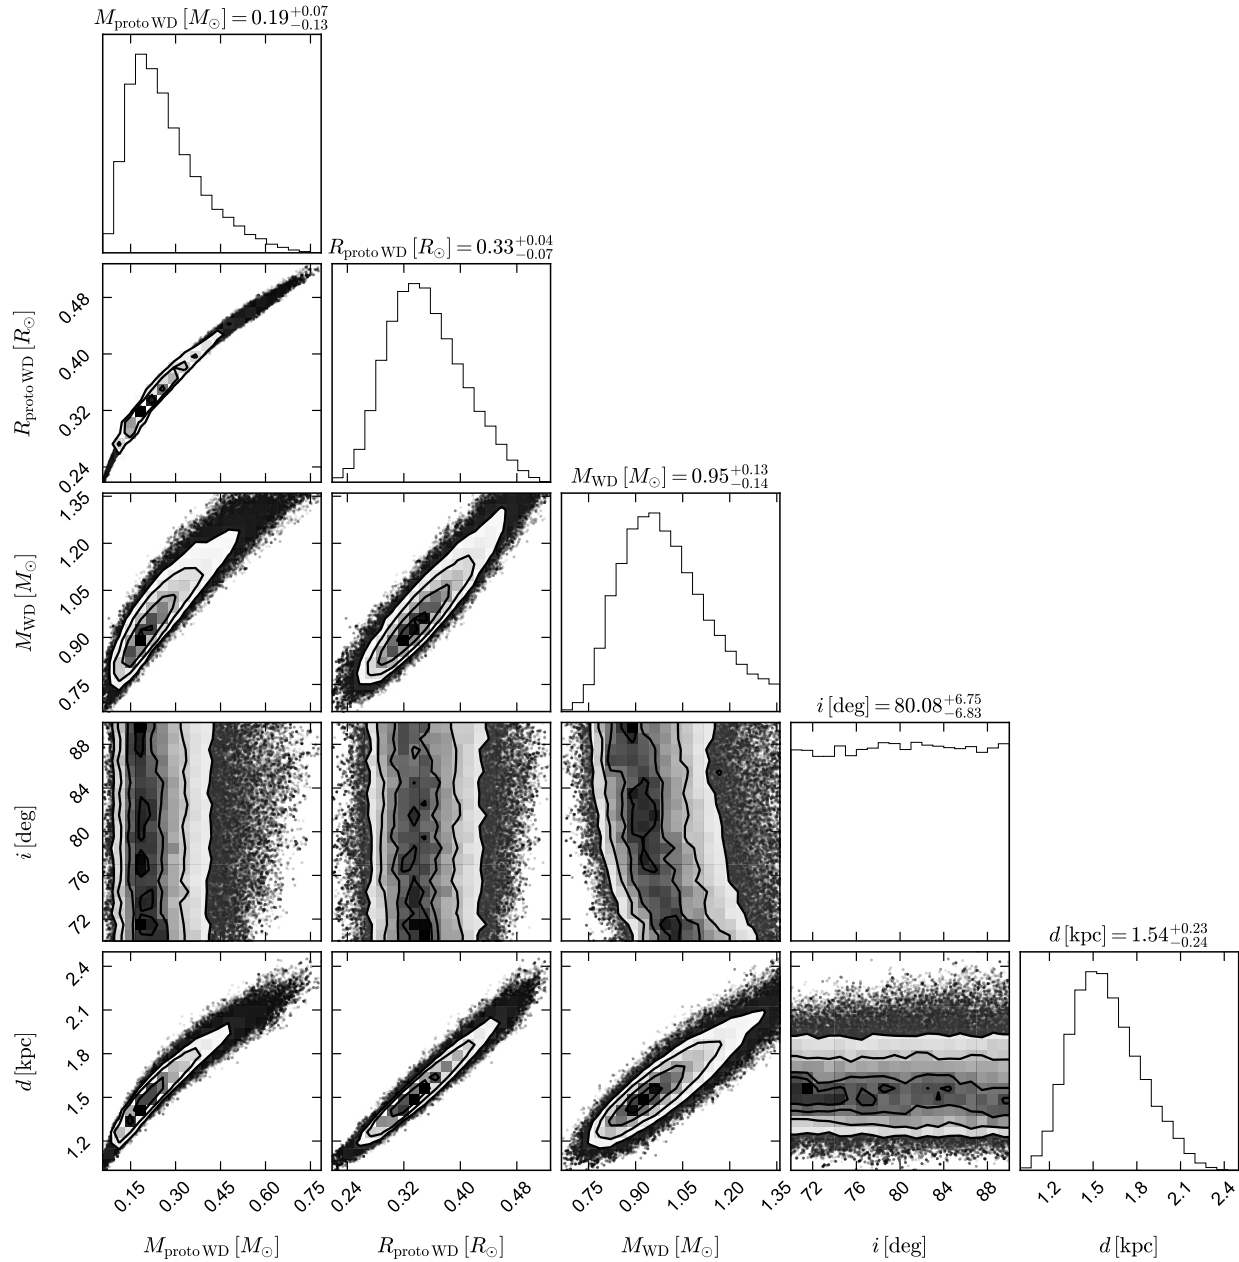


Figure 11.9: Parameter constraints from joint fitting of the light curve, SED, parallax, orbital solution, and donor spectrum. Marginalized constraints on each parameter are shown on the diagonal.

Table 11.3: Constraints from the combined fit. Uncertainties are 1σ . We first fit the light curve, adopting minimally informative priors on the inclination and Roche lobe filling factor. We then use the constraints from the light curve (namely, $i > 70$ deg and $R_{\text{proto WD}}/R_L = 0.99 \pm 0.01$) as priors during the combined fit. Finally (bottom two rows), we report constraints obtained when we require the proto-WD’s period and effective temperature to match the binary evolution models introduced in Section 11.6.

Parameter	description	units	prior	constraints
Light curve fit				
$R_{\text{proto WD}}/R_L$	Roche lobe filling factor	–	$\mathcal{U}(0, 1)$	0.99 ± 0.01
i	inclination	[deg]	$\sin(i)$	83_{-8}^{+5}
$\beta_{1, \text{proto WD}}$	gravity darkening exponent	–	$\mathcal{N}(0.32, 0.1)$	$0.29_{-0.05}^{+0.03}$
σ_{intr}	Light curve intrinsic scatter	[mag]	$\mathcal{U}(0, 1)$	0.025 ± 0.002
Combined fit				
$R_{\text{proto WD}}$	Proto-WD radius	$[R_{\odot}]$	$\mathcal{U}(0.01, 1)$	$0.33_{-0.04}^{+0.07}$
$M_{\text{proto WD}}$	Proto-WD mass	$[M_{\odot}]$	$\mathcal{U}(0.01, 1)$	$0.19_{-0.07}^{+0.13}$
M_{WD}	White dwarf mass	$[M_{\odot}]$	$\mathcal{U}(0.5, 1.36)$	$0.95_{-0.13}^{+0.14}$
i	inclination	[deg]	$\mathcal{U}(70, 90)$	80_{-7}^{+7}
d	distance	[kpc]	Equation 11.5	$1.54_{-0.23}^{+0.24}$
Combined fit including evolutionary models				
$R_{\text{proto WD}}$	Proto-WD radius	$[R_{\odot}]$	$\mathcal{U}(0.01, 1)$	0.29 ± 0.01
$M_{\text{proto WD}}$	Proto-WD mass	$[M_{\odot}]$	$\mathcal{U}(0.01, 1)$	0.15 ± 0.01

is correlated with orbital period (e.g. [Szkody & Mattei 1984](#); [Warner 2003](#)); for a period of 3.81 hours, the typical width at half-maximum, $\Delta T_{0.5}$, is ≈ 2.5 days. Similarly, the absolute magnitude of dwarf nova disks during outburst is correlated with orbital period; for a period of 3.81 hours, $M_{V,\max} \approx 4.75$ ([Warner 2003](#)). For J0141, this would correspond to a ≈ 2.5 mag brightening during outburst. It is of course not obvious that the outburst duration and amplitude of typical CVs will be a good proxy for an unusual system like J0140.

To estimate the probability that outbursts with interval T_n would have been detected by now, we perform Monte Carlo simulations in which we inject outbursts into the observed light curve (Figure 11.2). For simplicity, we model each outburst as a tophat increase in brightness of 2.5 mag that lasts 2.5 days. If the magnitude during outburst exceeds an observed magnitude or upper limit, we conclude that it would have been detected. We model the bursts as strictly periodic, with uniformly distributed orbital phase. For each T_n , we perform 1000 trials and record the fraction of trials that result in at least one burst being detected. This fraction is plotted in Figure 11.10.

For all $T_n \lesssim 1$ year, the detection efficiency is 100%. There are dips in the detection efficiency for T_n values that are integer multiples of 1 year. This is because J0140 is not observable for a few months each year, so if bursts were phased such that they occur only during these gaps, they would not be detected. Real dwarf nova outbursts are generally not completely periodic, so the detection efficiency is likely smoother than modeled in this experiment. Figure 11.10 shows that the probability of detecting at least one burst is greater than 90% for $T_n \lesssim 10^3$ days. The detection efficiency falls below 50% for $T_n \gtrsim 2500$ days.

11.6 Evolutionary history and comparison to known CVs and ELM WDs

If there is still ongoing mass transfer in J0140, then the proto-WD is an unusually evolved CV donor. If mass transfer was recently terminated, it is a just-detached ELM WD. We thus compare the system to known evolved CV secondaries and ELM WDs.

In Figure 11.11, we compare the proto-WD to CV donors from the literature that have been proposed to be evolved. These are QZ Ser ([Thorstensen et al. 2002a](#)), EI Psc (also called 1RXS J232953.9+062814; [Thorstensen et al. 2002b](#)), SDSS J170213.26 + 322954.1 ([Littlefair et al. 2006](#)), BF Eridani ([Neustroev & Zharikov 2008](#)), CSS J134052.0 + 151341 ([Thorstensen 2013](#)), HS 0218+3229 ([Rodríguez-Gil et al. 2009](#)), SDSS J001153.08-064739.2 ([Rebassa-Mansergas et al. 2014](#)), ASAS-SN 13cl ([Thorstensen 2015](#)), ASAS-SN 15cm ([Thorstensen et al. 2016](#)), ASAS-SN 14ho ([Gasque et al. 2019](#)), V1460 Her ([Ashley et al. 2020](#)), and ASAS-SN 18aan ([Wakamatsu et al. 2021](#)).

Most of these systems contain K star donors ($T_{\text{eff}} = 4500 - 5000$ K), which are warmer than is predicted for unevolved CV donors (dashed line; [Knigge 2006](#)). We note that evolved donors are common at long periods ($P_{\text{orb}} \gtrsim 6$ hours), because only evolved or unusually massive donors are large enough to already be Roche-filling at these periods. There are plenty of evolved donors with longer periods that are not shown in Figure 11.11; some of




Figure 11.10: Probability that at least one outburst would have been detected by now if the true outburst interval were a given value, based on the light curve in Figure 11.2. Sharp dips occur at multiples of 1 year, since J0140 is not observable for a few months each year; they would not be present if outbursts were not modeled as strictly periodic. The data rule out an outburst interval $T_n < 10^3$ days with 90% probability.




Figure 11.11: Effective temperatures and orbital periods of J0140 and known CV donors that have been proposed to have evolved donors (see Section 11.6 for references). All these systems have hotter donors than predicted by the semi-empirical evolutionary sequence for unevolved donors (dashed line), which we take from Knigge (2006). A majority of known CVs have temperatures close to this sequence. J0140 is considerably hotter than any previously characterized system at comparable period.

these are shown in Figure 11.12. However, evolved systems are rare at $P_{\text{orb}} \lesssim 6$ hours, and the CVs shown in Figure 11.11 represent a near-complete inventory of currently known systems with anomalously hot donors. J0140 is significantly hotter than other known CVs at comparable periods. If it is still mass-transferring, it is the most extreme evolved CV discovered to date.

11.6.1 MESA models

To understand the evolutionary history of J0140 and similar systems, we calculated a grid of binary evolution models using Modules for Experimentation in Stellar Astrophysics (MESA, version 12778; Paxton et al. 2011, 2013, 2015, 2018, 2019). The binary capabilities of MESA are described in detail by Paxton et al. (2015). We follow the mass-losing star with detailed calculations and model the more massive WD as a point mass. Roche lobe radii are computed using the fit of Eggleton (1983). Mass transfer rates during Roche lobe overflow are determined following the prescription of Kolb & Ritter (1990). The orbital separation evolves such that the total angular momentum is conserved when mass is transferred to the companion or lost, as described in Paxton et al. (2015). We model mass transfer as being fully non-conservative, with the mass that is transferred from the donor to the WD eventually being lost from the vicinity of the WD, as described by Tauris & van den Heuvel (2006, their “ β ” parameter). Although mass transfer in real CVs is not *instantaneously* non-conservative, the mass accreted by the WD is expected to be ejected from the system by classical nova explosions, which recur on a timescale that is short compared to the timescale on which the orbit evolves.

We initialize our calculations with a zero-age main sequence (ZAMS) star in a detached orbit with a point mass representing the WD. The orbit shrinks through gravitational wave radiation and magnetic braking. Magnetic braking is implemented following the prescription of Rappaport et al. (1983), which applies a torque,

$$\tau_{\text{mb}} = -6.8 \times 10^{-34} \left(\frac{M_{\text{donor}}}{M_{\odot}} \right) \left(\frac{R_{\text{donor}}}{R_{\odot}} \right)^{\gamma_{\text{mb}}} \left(\frac{P}{1 \text{ d}} \right)^{-3} \text{ dyn cm}, \quad (11.6)$$

removing angular momentum from the binary. This expression is derived under the assumption that the binary is tidally locked and the angular momentum loss from the donor through winds is the same as that from an isolated single star with the same rotation velocity (e.g. Skumanich 1972; Verbunt & Zwaan 1981). Comparison to the observed age-rotation period relation for solar-type stars suggests $\gamma_{\text{mb}} = 4$; Rappaport et al. (1983) found $2 \lesssim \gamma_{\text{mb}} \lesssim 4$ when matching their evolutionary models to the period distribution of observed CVs. We adopt $\gamma_{\text{mb}} = 3$, but our results are only weakly sensitive to this choice.

Magnetic braking is turned off when the donor becomes fully convective, because at that time the magnetic field is expected to weaken, or at least significantly change its structure (Spruit & Ritter 1983b; Donati & Landstreet 2009). We also expect magnetic braking to become inefficient when the donor loses its convective envelope (e.g. Kraft 1967). This is not implemented in the default distribution of MESA. Following Podsiadlowski et al. (2002),

we assume that the magnetic braking torque follows Equation 11.6 as long as q_{conv} , the mass fraction of the donor’s convective envelope, exceeds 0.02. For donors with $q_{\text{conv}} < 0.02$, we suppress the magnetic braking torque by multiplying Equation 11.6 by a factor $\exp(1 - 0.02/q_{\text{conv}})$. This effectively makes magnetic braking inefficient at $T_{\text{eff}} \gtrsim 6000$ K.

The “initial” separations in our calculations represent the separation after the more massive WD has formed, presumably through an episode of common envelope evolution (CEE). The separation distribution at this stage is not well-constrained observationally or theoretically, and we do not attempt to predict exactly what initial binary configurations would produce a given separation after CEE. We instead simply model a grid of separations to see which post-CEE separations could produce the observed properties of J0140.

Initializing the secondary on the ZAMS neglects its aging during the main-sequence lifetime of the massive WD’s progenitor. This is a reasonable approximation for our purposes, because, in order for the system to have gone through CEE and emerged as a WD+MS binary with a short period, the initial mass ratio of the binary must have been relatively unequal (e.g. [Soberman et al. 1997](#)). If the massive WD’s mass has not changed significantly from accretion or novae, then its current mass of $\approx 0.95M_{\odot}$ implies an initial progenitor mass between 3.5 and $5.5M_{\odot}$, depending on the adopted initial-final mass relation (e.g. [Cummings et al. 2018](#); [El-Badry et al. 2018e](#)). This corresponds to a pre-WD lifetime of at most 300 Myr.

We ran MESA tracks with initial periods ranging from 0.5 to 3 days, WD masses ranging from 0.8 to $1.2M_{\odot}$, and initial donor masses ranging from 1 to $2.5M_{\odot}$. The qualitative behavior of the models depends little on the assumed WD mass, so here we only analyze those with $M_{\text{WD}} = 0.9M_{\odot}$. Figure 11.12 shows the evolution of several tracks that all have initial donor masses of $1.1M_{\odot}$ and a range of initial orbital periods; we discuss other donor masses in Section 11.6.3. As a function of orbital period, we plot the donor mass, radius, and effective temperature, the mass of its helium core⁷, time since the beginning of the calculation, the mass transfer rate, the surface gravity of the donor, and the binary’s angular momentum loss due to magnetic braking and gravitational wave radiation. We also plot the mass, radius, and temperature of donors in known CVs with well-characterized donor parameters.⁸

The set of initial periods shown in Figure 11.12 are chosen to highlight the behavior of “normal” CVs as well as evolved systems like J0140. The track with $P_{\text{init}} = 1$ day displays behavior characteristic of “normal” CVs: the donor overflows its Roche lobe while still on the main sequence, before a helium core has formed. It grows smaller and cooler as the period shrinks, passing near most of the observed systems.⁹ At a period of about 3 hours,

⁷Defined as the region in which the helium mass fraction is $>99\%$.

⁸Effective temperatures for 91 CVs are taken from [Knigge \(2006, their Table 2\)](#). Only CVs with reliable spectroscopically determined spectral types are included. [Knigge \(2006\)](#) report spectral types, not effective temperatures; we convert them to T_{eff} using the conversion from [de Jager & Nieuwenhuijzen \(1987\)](#). We also plot masses and radii of donors in 19 eclipsing CVs from [Patterson et al. \(2005, their Table 8\)](#). We plot these, rather than larger compilations from elsewhere in the literature, because they contain masses and radii measured independently, without the assumption that the CVs follow a particular donor sequence.

⁹Even this model is somewhat warmer than the bulk of the observed population at short periods. This likely reflects the fact that the atmospheres of cool dwarfs are influenced by molecules, which are not included

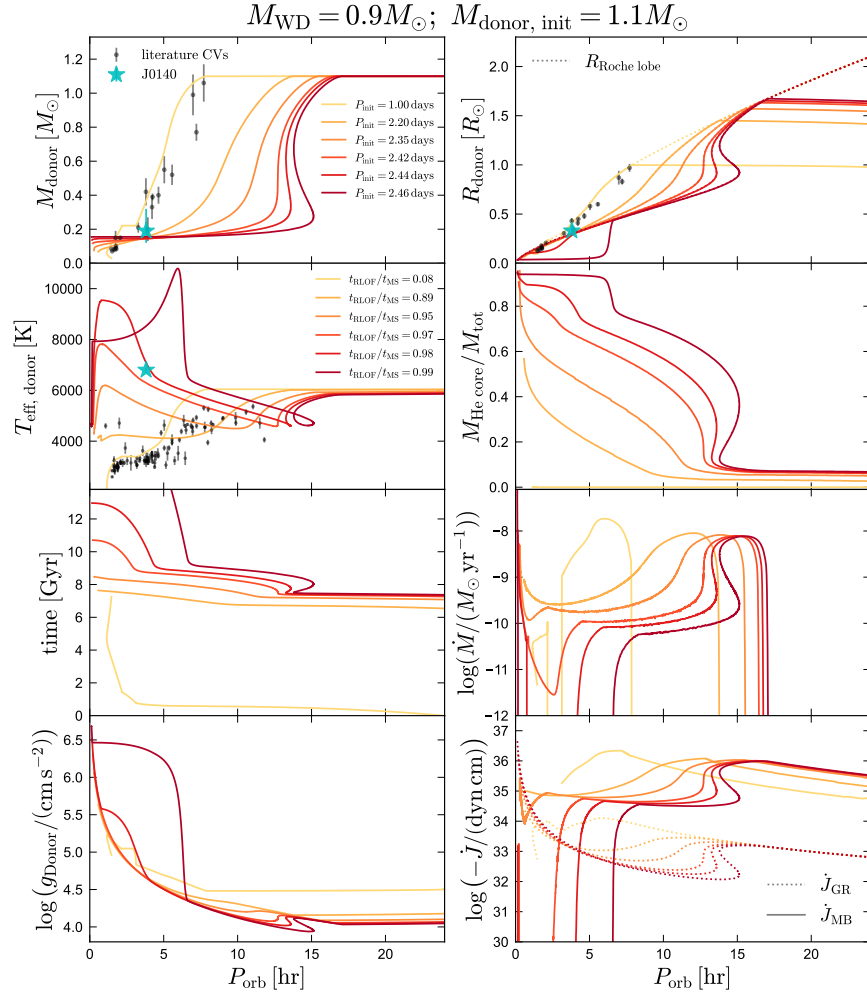


Figure 11.12: MESA binary models for CVs with evolved donors. We show 6 models, labeled by P_{init} (the initial period at which we initialize a detached WD+MS binary) and by the fraction of the donor’s normal MS lifetime that has passed at the onset of RLOF. Gray points with error bars show other CVs from the literature with well-measured masses, radii, and effective temperatures. The cyan point shows J0140. The yellow track with $P_{\text{init}} = 1$ day is representative of CVs in which the donor overflows its Roche lobe before undergoing significant nuclear evolution; it passes through the majority of the observed CV population. Darker tracks shows systems (with finely spaced periods) in which the donor overflows its Roche lobe increasingly near the end of the main sequence. By the time these systems reach $P_{\text{orb}} \sim 4$ hours, they have significant helium cores with thin H-burning envelopes. At fixed period, they have higher T_{eff} , lower masses and radii, and much lower mass transfer rates than ordinary CVs. J0140 is well described by such tracks. The most extreme evolved tracks become detached once they lose their convective envelopes and magnetic braking becomes inefficient (lower right); these become ELM WDs.

the donors become fully convective and magnetic braking ceases, slowing the orbital decay. At this point, the donor temporarily detaches from its Roche lobe and returns to thermal equilibrium as it passes through the period gap, before reestablishing Roche contact at $P_{\text{orb}} \sim 2$ hours. The donor continues to shrink until the period minimum, where it becomes degenerate. The same qualitative evolutionary scenario applies to all models that overflow their Roche lobes while on the main sequence, including models with higher- and lower-mass secondaries. For the choice of M_{WD} and $M_{\text{donor,init}}$ shown in Figure 11.12, all systems with initial periods between about 0.4 and 2 days follow similar evolution to this track.

Along with their initial periods, we label the tracks in Figure 11.12 by the fraction of the donor’s main-sequence lifetime (if it were isolated; this is $t_{\text{MS}} = 7.54$ Gyr for the $1.1 M_{\odot}$ donor shown in the figure) that has passed at the time of first Roche lobe overflow. We define the main-sequence lifetime as the time until the helium core mass reaches the Schönberg–Chandrasekhar limit, which we take to be 0.08 times the total mass of the star. Models with $t_{\text{RLOF}}/t_{\text{MS}} \lesssim 0.8$ all behave similarly to the track shown with $t_{\text{RLOF}}/t_{\text{MS}} = 0.08$.

Models with longer initial periods and $t_{\text{RLOF}}/t_{\text{MS}} \gtrsim 0.8$ behave qualitatively differently. These donors overflow their Roche lobes as a helium core is already forming. Initial RLOF occurs at a significantly longer period than in CVs with unevolved donors, because the donor’s radius has increased by $\sim 60\%$ during its main-sequence evolution. Once most of the envelope has been lost – over a period of several hundred Myr – the effective temperature of the donor begins to increase. The most evolved models, which have the thinnest hydrogen envelopes, reach the highest temperatures. By the time they reach $P_{\text{orb}} \sim 4$ hours, the evolved donors have low mass loss rates, $\dot{M} \lesssim 10^{-10} M_{\odot} \text{ yr}^{-1}$.

The mass transfer rate for the evolved models in Figure 11.12 is never particularly high, reaching at most $10^{-8} M_{\odot} \text{ yr}^{-1}$. It can, however, be significantly higher for higher initial donor masses (Section 11.6.3). In the most evolved models shown, the orbital period does not decrease monotonically with time, but increases slightly while the donor’s mass is between ≈ 0.6 and $\approx 0.3 M_{\odot}$. The orbital separation *does* decrease monotonically; the period increase reflects a rapid decrease in total mass.

The lower mass transfer rate near $P_{\text{orb}} = 4$ hours in models with evolved donors is a consequence of (a) these donors’ smaller masses and radii, which lead to weaker magnetic braking torques (Equation 11.6), and (b) the fact that the mass in evolved donors is more centrally concentrated, so a fixed decrease in radius corresponds to a smaller decrease in mass.

Once the donor reaches $T_{\text{eff}} \gtrsim 6000$ K and loses its convective envelope, magnetic braking ceases. This can be seen in the bottom panel of Figure 11.12, which shows a sharp fall-off in \dot{J}_{MB} for the evolved models that reach $T_{\text{eff}} \gtrsim 6000$ K. The binary’s orbital decay is then driven only by gravitational radiation. This leads to an immediate decrease in the mass transfer rate. In the two most evolved models, this causes the donor to detach from its

in the MESA calculations. More reliable effective temperatures for cool dwarfs can be obtained by changing the boundary condition where model atmospheres are stitched to interior models (see Choi et al. 2016, their Figure 14). We refrain from doing so here, using the `photosphere_tables` atmosphere boundary in all calculations.

Roche lobe and reestablish thermal equilibrium at a smaller radius. During this period, the donor would appear as a detached ELM WD. For the model that most closely matches the effective temperature of J0140, detachment occurs *just* at $P_{\text{orb}} = 3.8$ hours. Slightly less extreme models remain in Roche contact, either because magnetic braking never fully shuts off, or because the binary has reached short enough periods that angular momentum loss through gravitational waves is efficient by the time magnetic braking becomes inefficient.

It is important to note that there are few observational constraints on the efficiency of magnetic braking in stars like the proto-WD in J0140. The details of which models remain in Roche contact, which become detached, and when detachment occurs, depend sensitively on the adopted magnetic braking law. Characterization of similar CVs with hot and warm donors, particularly at $1 \lesssim P_{\text{orb}}/\text{hour} \lesssim 3$, will allow the efficiency of magnetic braking in hot stars to be better constrained.

The evolutionary models also allow us to refine our estimate of the mass and radius of the proto WD: plausible models with evolved donors in Figure 11.12 have $M_{\text{proto WD}} = 0.15 \pm 0.01 M_{\odot}$ and $R_{\text{proto WD}} = 0.29 \pm 0.01 R_{\odot}$. Other published evolutionary tracks for ELM WDs imply consistent values (e.g. Sun & Arras 2018; Li et al. 2019). Given the measured angular diameter, this would imply a distance $d = 1.30 \pm 0.05$ kpc. These values are consistent with, but tighter than, our constraints that are independent of evolutionary models (Section 11.4).

Figure 11.13 shows the evolution of the same MESA models on the HR diagram. The unevolved CV track with $t_{\text{RLOF}}/t_{\text{MS}} = 0.08$ closely follows the single-star main sequence (dashed black line). Most of the CVs in the sample to which we compare fall near this track, though they are on average offset slightly to lower T_{eff} . A few literature systems (gray points) are significantly hotter than the single-star main sequence at their luminosity; these are CVs thought to host evolved donors. J0140 is significantly offset from these systems. The MESA tracks that pass near it suggest it will continue evolving toward higher T_{eff} at nearly fixed luminosity before it eventually cools.

Figure 11.13 also shows a sample of spectroscopically confirmed ELM WDs assembled by Pelisoli & Vos (2019). A majority of these objects were discovered by the ELM survey (Brown et al. 2010). We estimate the luminosity of these objects from their reported T_{eff} , $\log g$, and mass.¹⁰ We note that selection effects for these objects and for CVs are very different, and we make no attempt to interpret the relative numbers of objects in the two populations. ELM WDs are found at even higher T_{eff} than J0140, and they are all detached. The two most extreme MESA tracks run through this population in temperature–luminosity space after detachment. This suggests that, whether or not it is currently mass transferring, J0140 will evolve to become a detached ELM WD. We note that a majority of the ELM WDs plotted have somewhat longer orbital periods than J0140. Such systems are thought to form if $t_{\text{RLOF}}/t_{\text{MS}} > 1$; i.e., if the donor’s core has reached the Schönberg–Chandrasekhar limit before first RLOF. In this case, the donor radius and separation expand in response

¹⁰That is, we calculate $L = L_{\odot} (R/R_{\odot})^2 (T_{\text{eff}}/T_{\text{eff},\odot})^4$, with $(R/R_{\odot})^2 = G(M/M_{\odot})(g/g_{\odot})^{-1}$. Since no mass uncertainty is reported for most of these objects, we assume an uncertainty of $0.03M_{\odot}$ when calculating the uncertainty in L .




Figure 11.13: HR diagram. The same MESA tracks from Figure 11.12 are compared to other CV donors (gray points) and to extremely low mass (ELM) WDs. Most normal CV donors fall near the single-star main sequence, which also overlaps the yellow model with early RLOF. Tracks in which the donor overflows its Roche lobe near the end of its main-sequence evolution (orange and red) diverge from the main sequence toward higher T_{eff} , approaching the ELM white dwarfs. The most extreme models become detached in this phase (Figure 11.12). J0140 is likely transitioning between the two populations, evolving to higher temperature at near-constant luminosity.




Figure 11.14: Future evolution of J0140. Colored tracks show the same MESA models from Figures 11.12 and 11.13. Normal CVs reach a period minimum at about 80 minutes, when they become degenerate. Systems with evolved donors are smaller when they become degenerate (due to higher mass and lower Y_e), allowing them to reach much shorter minimum periods. These models suggest that J0140 will reach a minimum period $P_{\text{orb}} \lesssim 15$ minutes, when (if it avoids a merger) it will be classified as an AM CVn binary.

to mass loss, and the system never reaches short periods (e.g. Kalomeni et al. 2016; Sun & Arras 2018).

11.6.2 Future evolution

Irrespective of whether angular momentum loss is driven by magnetic braking or gravitational waves, the orbit of J0140 will continue to decay. Figure 11.14 shows the evolution of the MESA models from Figures 11.12 and 11.13 at short periods. We run the calculations until the donor mass falls below $0.03 M_{\odot}$, noting that in the most evolved model, this takes longer than the age of the Universe.

All models eventually reach a period minimum and then begin to widen again. The mass

ratio is low enough that a merger is unlikely, at least under the simplified treatment of mass transfer and accretion onto the WD assumed here (e.g. Marsh et al. 2004). The minimum period is much shorter for the evolved tracks, which reach minimum periods between 5 and 30 minutes, compared to ≈ 80 minutes for the unevolved CV track. Mass transfer rates at the period minimum are high but not catastrophic, reaching $1.5 \times 10^{-7} M_{\odot} \text{ yr}$ in the most evolved model. Models that reach short periods evolve through the period minimum relatively quickly; for example, the most evolved model spends only 20 Myr with $P_{\text{orb}} < 20$ minutes. The period minimum of evolved CVs was explored extensively by Podsiadlowski et al. (2003), and the behavior of our models is qualitatively similar to their predictions.

The shorter period minimum for evolved donors can be understood as follows. The period minimum is set by the radius of the donor at the time when it becomes degenerate. Unevolved donors become degenerate at the hydrogen burning limit, $M_{\text{donor}} \approx 0.08 M_{\odot}$. Due to modest departures from thermal equilibrium, their radius at this time is $R_{\text{donor}} \approx 0.12 R_{\odot}$, about 20% larger than the radius of a main-sequence star of the same mass. Equation 11.1 then predicts a period minimum of 77 minutes, close to what is found in the MESA models. Evolved donors have smaller radii. The radius of cold, non-relativistic degenerate objects scales as $R \sim M^{-1/3} Y_e^{5/3}$, where Y_e is the mean number of electrons per nucleon. Taking $Y_e \approx 0.5$ for evolved donors that are mostly helium and $Y_e \approx 0.87$ for normal donors with Solar composition, the radius of a $0.16 M_{\odot}$ evolved donor (corresponding to the most evolved model in Figure 11.14) at the period minimum is predicted to be ≈ 3.2 times smaller than that of a $0.08 M_{\odot}$ hydrogen donor. This corresponds to a 64 times higher mean density and an 8 times shorter minimum period (Equation 11.1). The actual minimum period is ≈ 13 times shorter, likely a result of the fact that near the period minimum, the unevolved model is inflated by mass loss-driven departures from thermal equilibrium, while the evolved model (being detached just prior to the onset of mass transfer and the minimum orbital period) is not.

Near its period minimum, J0140 may appear as an ‘‘AM-CVn’’ binary, (e.g. Solheim 2010). An ‘‘evolved CV’’ channel for producing AM-CVn binaries is discussed in Podsiadlowski et al. (2003). Because a relatively fine-tuned scenario is required to produce an ultra-compact binary from a CV, most studies have concluded that the evolved CV channel likely only forms a small fraction of ultracompact binaries (van der Sluys et al. 2005; Solheim 2010; Goliach & Nelson 2015). Nevertheless, all plausible models we construct for J0140 eventually reach short periods and either remain mass-transferring or resume mass transfer after the period minimum.

Other possible formation channels for AM-CVn binaries include a WD+WD channel (Tutukov & Yungelson 1979) and a He-star-WD channel (Iben & Tutukov 1991; Brooks et al. 2015); which channel dominates is not known. If J0140 becomes detached before the period minimum, it might be counted instead as contributing to the WD+WD channel, for which a few detached progenitor systems have been identified (Kilic et al. 2014). If J0140 becomes detached, it will evolve to become very similar to these systems. Near the period minimum, the system would then resemble some of the ultracompact WD binaries recently discovered by ZTF (e.g. Burdge et al. 2019, 2020).

The fate of the system after the period minimum is unclear. If a merger is avoided, accretion of helium onto the surface of the more massive WD will lead to unstable thermonuclear flashes, analogous to classical novae on normal CVs. These flashes are predicted to become less frequent and more energetic as the accretion rate decreases, eventually culminating in a “final flash” that may fuse radioactive elements and could be observed as a rapidly-evolving transient (a “Ia” supernova; Bildsten et al. 2007). Whether a detonation occurs in this scenario depends on the time-dependence of the accretion (Piersanti et al. 2015).

An alternative possibility is that, following classical nova-like flashes on the surface of the accreting WD, dynamical friction will shrink the orbit and cause the binary to merge (Shen 2015). The outcome of such a merger would depend on the masses of the components. If the accreting WD has a carbon-oxygen core and the total mass exceeds $\approx 1.2M_{\odot}$, then a SNe Ia may be triggered through surface He detonation and subsequent C detonation in the core by a converging shock (e.g. Livne 1990; Pakmor et al. 2013; Shen & Bildsten 2014). For lower total masses (which are more likely, given our constraints), the merger might form an R Coronae Borealis star (e.g. Clayton 2012).

11.6.3 Varying model parameters

Here we briefly summarize how the evolution of the MESA models changes when the problem setup is varied.

- *Shorter or longer initial periods:* Shorter initial periods behave basically identically to the track with $P_{\text{init}} = 1$ day. In binaries with longer initial periods, the donor overflows its Roche lobe after the core has reached the Schönberg–Chandrasekhar limit ($t_{\text{RLOF}}/t_{\text{MS}} > 1$). These systems never reach short periods, because mass loss causes the donor’s radius to increase rather than decrease. The orbit widens in response to mass loss, leading to further expansion and mass loss. Such systems eventually form helium white dwarfs, with a well-characterized relation between final orbital period and mass (e.g. Rappaport et al. 1995). The division between such “divergent” systems and converging systems that become CVs is called the “bifurcation limit” and has been studied elsewhere (Pylyser & Savonije 1988; Podsiadlowski et al. 2003).
- *Higher or lower initial donor masses:* For significantly lower initial donor masses, the donor cannot terminate its main-sequence evolution within a Hubble time, excluding the formation of a CV with an evolved donor. For higher initial masses, the donor has a radiative envelope on the main sequence, making magnetic braking inefficient. This does not qualitatively change the system’s evolution once mass transfer begins, because mass loss leads to a drop in T_{eff} and efficient magnetic braking. However, due to inefficient magnetic braking and a shorter main-sequence lifetime, donors with higher initial masses are able to reach short periods and become CVs only if the initial period is short (e.g., $P_{\text{init}} \lesssim 1$ day for an initial donor mass of $2M_{\odot}$). In addition, the mass transfer rates shortly after RLOF are much higher (up to $10^{-6}M_{\odot} \text{ yr}^{-1}$) in models with more massive donors. This allows for stable hydrogen burning on the surface of

the massive WD, potentially producing a supersoft X ray source or SNe Ia (van den Heuvel et al. 1992; Han & Podsiadlowski 2004).

- *Different magnetic braking law:* We also experimented with models in which magnetic braking is not turned off when the donor’s envelope becomes radiative, as is currently the default behavior in MESA. The most significant difference between these models and those presented in Figures 11.12-11.14 is that their orbital decay does not slow down when the donor heats up. As a result, even the most evolved models that reach short periods remain mass transferring (though at low rates; $\dot{M} < 10^{-10} M_{\odot} \text{ yr}^{-1}$).

11.7 Summary and Discussion

We have presented discovery data for LAMOST J0140355+392651 (J0140), an unusual binary with period $P_{\text{orb}} = 3.81$ hours containing a white dwarf and a hot, low-mass secondary. Our main conclusions are as follows:

1. *Basic properties:* The system’s optical light curve is dominated by ellipsoidal variability (Figure 11.1). The optical spectrum and SED of J0140 are dominated by the low-mass star, which appears to be a proto-WD (Figures 11.6 and 11.7). The massive WD dominates in the UV. Fitting the SED, *Gaia* parallax, optical spectrum, and radial velocities allows us to constrain physical properties of the system, including the proto-WD’s mass, radius, and temperature, the WD mass, the orbital inclination, and the distance (Figure 11.9 and Table 11.3).
2. *Comparison to CVs and ELM WDs:* The properties of J0140 are transitional between known populations of CVs and ELM WDs. The effective temperature of the secondary, $T_{\text{eff}} = 6800 \pm 100$, is much higher than in any known CVs at similar periods: most CVs near $P_{\text{orb}} = 4$ hours have $T_{\text{eff}} \sim 3500$ K, and the hottest donors discovered previously had $T_{\text{eff}} \sim 5000$ K (Figure 11.11). The secondary’s mass and radius are also smaller than those of normal CV donors at comparable period (Figure 11.12). The system has not been observed in outburst; if it does experience outbursts, their recurrence time is likely at least a few years (Figures 11.2 and 11.10), implying that the mass transfer rate is very low. On the other hand, the secondary is cooler and more bloated than known ELM WDs (Figure 11.13).
3. *Evolutionary history:* We calculated a set of MESA binary evolution models to investigate the formation and future evolution of J0140 (Figures 11.12-11.14). These models reproduce the secondary’s high temperature, small mass and radius, and low mass transfer rate if the secondary overflowed its Roche lobe *just* at the end of its main-sequence lifetime (within a few percent). In this case, all that is left of the secondary at $P_{\text{orb}} \sim 4$ hours is a helium core with a thin hydrogen-burning envelope. Some of these models remain as mass-transferring CVs for the rest of their evolution, while others become detached as they grow hotter and temporarily appear as detached ELM

WDs (Figure 11.13). Unlike normal CVs, which reach a minimum period of ≈ 80 minutes, models with evolved donors reach much shorter minimum periods (5-30 minutes), where they will appear as AM CVn binaries (Figure 11.14).

Further observations are necessary to better understand the nature of the system. In particular, higher-resolution and higher-SNR spectra will allow a deeper search for emission features associated with accretion, which will enable a more conclusive determination of whether there is ongoing mass transfer. UV observations with the *Hubble Space Telescope* will provide direct constraints on the temperature of the WD. Such observations will also constrain the nitrogen and carbon abundance of the secondary (e.g. Gänsicke et al. 2003; Ashley et al. 2020). This will test the evolutionary scenario proposed in Section 11.6, which predicts that both the donor and WD should bear evidence of CNO processing: nitrogen and helium enhancement, and carbon depletion. Finally, higher-cadence photometry will characterize the short-timescale variability that is thus far manifest only as additional scatter in the light curve (Figure 11.3). If this variability is due to pulsations, it will be useful for probing the internal structure of the WD (e.g. Hermes et al. 2012).

J0140 is clearly a rare system, containing the hottest secondary among of order 1000 spectroscopically confirmed CVs with $P_{\text{orb}} \lesssim 6$ hours. We expect, however, that there are numerous similar systems with comparable or brighter magnitude yet to be discovered. J0140 exhibits neither outbursts nor strong emission lines, the two observables through which CVs are most commonly identified. Existing samples of CVs are thus likely biased against objects like J0140. We suspect that similar objects can be efficiently identified by searching the light curves of sources that fall blueward of the main sequence in the color-magnitude diagram (e.g. Figure 11.17) for large-amplitude ellipsoidal variation. Light curves will be dominated by ellipsoidal variation only if the donor dominates the optical light, which at short periods will only occur if the donor is unusually hot. We have initiated such a search and will present its yields in future work.

Connection to EL CVn binaries and sdA stars

We note that the proto-WD in J0140 is in a similar evolutionary state to the bloated stripped stars found in EL CVn-type binaries (e.g. Maxted et al. 2014). In EL CVn binaries, the companions are main-sequence stars rather than WDs. Like the ELM WDs discovered by the ELM survey, the stripped stars in known EL CVn binaries are somewhat hotter and more compact than the proto-WD in J0140, likely because they are further evolved.

Another population to which J0140 is closely related is the “sdA” stars, a heterogeneous collection of stars discovered by Kepler et al. (2016) with A-type spectra that appeared to have higher surface gravities than main-sequence stars of similar temperature. A majority of the sdA stars turned out to be main-sequence stars with overestimated $\log g$ (e.g. Hermes et al. 2017). However, a small fraction appear to be genuine pre-ELM WDs, and some of these have temperatures and surface gravities similar to the proto-WD in J0140 (Pelisoli et al. 2019). This raises the question whether their evolutionary status is similar to that of J0140.

To investigate this possibility, we analyzed the ZTF light curves of the 50 “high probability” pre-ELM WDs identified by Pelisoli et al. (2019) within the SDSS sdA sample. 42 of these 50 sources have high-quality ZTF light curves. Only one of them shows clear high-amplitude variability,¹¹ and while likely somewhat tidally distorted, it is not close to filling its Roche lobe. We therefore conclude that the sdA pre-ELM WD candidates identified by Pelisoli et al. (2019) have longer orbital periods than J0140, have already shrunk well within their Roche lobes, and likely formed from donors that initiated mass transfer *after* the end of their main-sequence evolution.

11.7.1 Is J0140 detached or mass-transferring?

The lack of outbursts and strong emission lines in J0140 raises the question whether the system is a detached binary or a CV. Some of the MESA models we calculate (Figure 11.12) do indeed become detached at $P_{\text{orb}} \lesssim 4$ hours, because magnetic braking becomes inefficient after the donor loses its convective envelope.

The relatively high temperature of the more massive WD inferred from the FUV emission ($T_{\text{eff}} \approx 20,000$ K) suggests that if the binary is detached, it must have ceased mass transfer recently.¹² Similarly, the large observed amplitude of ellipsoidal variation implies that the proto-WD is very close to Roche lobe overflow: our light curve fit (Section 11.4.1) yields a Roche lobe filling factor $R_{\text{proto WD}}/R_L > 0.98$. It is possible that detachment occurred recently and the donor is still almost Roche lobe filling, but this would require some fine tuning.

Long outburst recurrence timescales are expected at low accretion rate, because a longer period of sustained accretion is required to build up enough mass in the disk for it to become unstable (e.g. Warner 2003; Hameury 2020). The intrinsic distribution of CV outburst timescales is difficult to constrain observationally, because there is a strong detection bias toward systems with frequent outbursts (e.g. Breedt et al. 2014). Recently, Yu et al. (2019) discovered a long-period CV ($P_{\text{orb}} \sim 9$ hours) with only one outburst in 4 years of nearly continuous Kepler data. The expected future evolution of that system is similar to what we infer for J0140. V1460 Her, another CV proposed to contain an evolved donor, experienced only two confirmed outbursts in 14 years (Ashley et al. 2020). Since J0140 appears to be the most extreme member of the evolved CV population discovered so far, we cannot rule out that outbursts occur on a long recurrence timescale.

Given the unusual nature of the system, it is also worth considering whether accretion could be occurring without *any* outbursts. At sufficiently low accretion rates, cooling is inefficient and a thin disk is not expected to form. For J0140, the minimum accretion rate below which even the outer disk is in the radiatively inefficient regime is $\dot{M} \lesssim \text{few} \times$

¹¹SDSS J 132713.01+382514.0, with reported $T_{\text{eff}} = 7967 \pm 18$ K, shows what appears to be ellipsoidal variability with $P_{\text{orb}} = 2.00$ hours and peak-to-peak variability amplitude $\approx 8\%$ (compared to $\approx 25\%$ in J0140). The variability is also consistent with a reflection effect binary with $P_{\text{orb}} = 1.00$ hours.

¹²This T_{eff} is really an upper limit, since some of the FUV emission could come from a disk. But a disk could not be invoked if the binary were detached.

Table 11.4: LAMOST DR5 parameters for J0140, estimated from spectra obtained in two separate visits. A third visit (obsid 1012095) had too-low SNR to yield useful RVs or stellar parameters.

parameter	units	visit 1	visit 2
obsid	–	184214089	353912095
mjd	d	56621	57278
g -band SNR	–	14.9	19.7
RV	km s^{-1}	-5 ± 51	-377 ± 62
T_{eff}	K	6586 ± 246	6394 ± 829
$\log(g/(\text{cm s}^{-2}))$	–	4.37 ± 0.39	4.22 ± 1.19
[Fe/H]	dex	-0.17 ± 0.23	-0.16 ± 0.77
spectral type	–	F0	F3

$10^{-11} M_{\odot} \text{ yr}$ (e.g. Menou et al. 1999), which is plausible for the accretion rates predicted by the MESA models in Section 11.6.1.

We conclude that J0140 is either currently undergoing mass transfer at an unusually low rate, or recently became detached, but further observations are needed to confidently distinguish between these possibilities.

Acknowledgements

We thank Boris Gänsicke and Tom Marsh for helpful comments. We are grateful to the staff at Lick Observatory for their assistance in obtaining the Kast spectra, and to the developers of `pypeit` and `Phoebe` for maintaining and making public the codes. KE acknowledges support from an NSF graduate research fellowship and a Hellman fellowship from UC Berkeley. This research benefited from meetings supported by the Gordon and Betty Moore Foundation through grant GBMF5076. EQ was supported in part by a Simons Investigator award from the Simons Foundation. KJS received support for this work from NASA through the Astrophysics Theory Program (NNX17AG28G). We thank Geoff Tabin and In-Hei Hahn for their hospitality during the writing of this manuscript.

Data Availability

Data in this paper are available upon reasonable request to the corresponding author.

11.8 Summary of LAMOST data

Table 11.4 lists the parameters of J0140 as reported in LAMOST DR5 and measured from combined spectra. These spectra are expected to be subject to significant orbital smearing, so parameters derived from them should be interpreted with caution.

The first visit (obsid 184214089) was split between two 2700 second subexposures, while the second (obsid 353912095) was split between three 1800 second subexposures. RVs measured from the subexposures are shown in Figure 11.15. There is clearly significant RV variation (up to 600 km s^{-1}) between subexposures of a single visit, implying that the spectra produced by coadding these sub-exposures will be smeared out. Indeed, the sub-exposure spectra are also subject to significant smearing, as the RV varies by more than 300 km s^{-1} between adjacent sub-exposures.

Figure 11.15 also shows the RVs predicted for the orbital ephemeris derived in Section 11.3.4 from the light curve and our follow-up spectroscopy. The LAMOST RVs are not used in deriving the ephemeris, but they are in reasonably good agreement with it. All 5 subexposure RVs are offset above our orbital solution, by $\approx 40 \text{ km s}^{-1}$ on average. This may indicate a zeropoint offset in our orbital solution compared to the LAMOST RV scale, since our follow-up RVs are not absolutely calibrated. This would have no effect on the inferred physical parameters of the binary.

11.9 Summary of Kast spectra

Table 11.5 summarizes the multi-epoch spectra of J0140 we obtained with the Kast spectrograph at Lick observatory.

11.10 Phase-dependent spectra

To investigate the possibility of variability in the spectra with orbital phase, we binned the single-epoch spectra in phase before coadding. Figure 11.16 compares the coadded spectra at $\phi \approx 0$ (when the WD passes in front of the donor), $\phi \approx 0.5$ (when the donor passes in front of the WD), and $\phi \approx 0.25$ or $\phi \approx 0.75$ (when both components are viewed side-on). In all cases, we include spectra within 0.1 of the indicated phase. Most lines have consistent strengths across phase bins. The $\text{H}\alpha$ line, however, appears somewhat stronger at $\phi \approx 0.5$ than at other phases. This is the phase at which the WD and disk (if it exists) are behind the donor, so the increased $\text{H}\alpha$ absorption at $\phi \approx 0.5$ is consistent with a scenario in which the disk contributes weak $\text{H}\alpha$ emission that is fully or partially eclipsed at $\phi \approx 0.5$. An alternative explanation is that irradiation by the hot WD gives rise to emission on one side of the donor, which would be hidden from view at $\phi \approx 0.5$. Higher-SNR spectra are required to confirm the phase-dependence of the $\text{H}\alpha$ line, and, if it is confirmed, to further elucidate its origin.

11.11 CMD position

Figure 11.17 shows J0140 on the *Gaia* eDR3 color-magnitude diagram. We compare it to other stars with good astrometry within 2 degrees (which are subject to a similar extinction)




Figure 11.15: Radial velocities measured from individual LAMOST sub-exposures for two separate visits. Horizontal errorbars show exposure time. Sub-exposure spectra are normally coadded (without any RV shifts) to produce the visit spectra from which stellar parameters are derived. For J0140, there are large RV shifts between visits, so the combined spectrum is subject to orbital smearing. Red line shows the radial velocity predicted by the orbital solution obtained from the light curve.

Table 11.5: Summary of Kast spectra. HMJD is mid-exposure time. Setup is described in Section 11.3.5.

HMJD UTC	phase	exptime [s]	setup	SNR	RV [km s ⁻¹]
59060.3371	0.04	600	B	5.6	-118 ± 21
59060.3444	0.08	600	B	7.9	-176 ± 19
59060.3518	0.13	600	B	8.8	-296 ± 13
59060.3643	0.21	600	B	10.3	-378 ± 12
59060.3717	0.25	600	B	11.7	-378 ± 10
59060.3790	0.30	600	B	11.3	-386 ± 12
59060.3940	0.39	600	B	10.8	-243 ± 12
59060.4015	0.44	600	B	10.1	-137 ± 12
59061.4040	0.76	600	A	11.5	311 ± 19
59061.4113	0.81	600	A	10.4	232 ± 30
59061.4186	0.85	600	A	10.9	246 ± 20
59061.4305	0.93	600	A	9.3	105 ± 21
59061.4378	0.97	600	A	8.8	32 ± 24
59062.3340	0.62	600	A	10.4	188 ± 16
59062.3414	0.67	600	A	10.5	274 ± 33
59062.3487	0.72	600	A	9.3	260 ± 19
59062.3601	0.79	600	A	10.5	239 ± 42
59062.3674	0.83	600	A	9.6	189 ± 55
59062.3747	0.88	600	A	8.5	-15 ± 54
59064.3512	0.34	600	A	5.4	-287 ± 56
59064.3585	0.39	600	A	6.0	-301 ± 54
59064.3658	0.43	600	A	6.4	-132 ± 38
59064.4701	0.09	600	A	13.6	-215 ± 32
59064.4774	0.14	600	A	13.9	-276 ± 24
59069.3438	0.81	600	A	17.6	240 ± 12
59069.3511	0.86	600	A	16.9	194 ± 13
59069.3584	0.90	600	A	15.2	70 ± 13
59069.3708	0.98	600	A	10.5	-32 ± 21
59069.3781	0.03	600	A	13.5	-155 ± 16
59069.3854	0.07	600	A	13.6	-232 ± 12
59069.3927	0.12	600	A	14.5	-315 ± 18
59069.4047	0.20	600	A	16.0	-355 ± 15
59069.4120	0.24	600	A	16.3	-391 ± 15
59069.4193	0.29	600	A	16.4	-361 ± 13
59069.4266	0.33	600	A	15.8	-348 ± 20
59069.4385	0.41	600	A	14.9	-208 ± 19
59069.4458	0.45	600	A	14.7	-146 ± 7
59069.4531	0.50	600	A	14.6	-58 ± 10
59069.4604	0.55	600	A	14.3	57 ± 10
59069.4715	0.62	600	A	14.8	179 ± 10
59069.4788	0.66	600	A	14.3	257 ± 20




Figure 11.16: Coadded spectra of J0140, produced by stacking single-epoch spectra with phases within 0.1 of the values specified in the legend. The strongest lines are labeled. No strong phase dependence is obvious, but there are hints of increased H α absorption near $\phi = 0.5$. This may be due to the full or partial eclipse of a H α emitting component at $\phi = 0.5$.




Figure 11.17: J0140 on the *Gaia* color-magnitude diagram (CMD). In the left panel, black points show sources within 2 degrees of J0140 with clean and precise photometry and astrometry; in the right panel, they show sources over the whole sky but with $d < 100$ pc. J0140 is significantly bluer than normal main-sequence stars of the same absolute magnitude. Right panel shows the predicted individual CMD positions of the proto-WD and WD from the SED models in Figure 11.7, including effects of extinction. Even without the light contributions of the WD, the donor is off the main sequence.

and to the *Gaia* 100 pc sample.¹³ J0140 falls below the main sequence; that is, it is bluer than a main-sequence star of the same absolute magnitude. In the right panel, we show the predicted individual CMD positions of the donor, disk, and WD. As in Figure 11.7, we caution that the SED of the disk and WD are poorly constrained.

Curiously, the CMD position of J0140 is quite similar to that of “normal” CVs with the same period (e.g. [Abril et al. 2020](#); [Abrahams et al. 2020](#)). Most CVs with $P_{\text{orb}} \sim 4$ hours, however, are completely dominated by the accretion disk in the optical, even in quiescence. At $P_{\text{orb}} \sim 4$ hours, the total luminosity of donors on the CV donor sequence is only $0.015L_{\odot}$ ([Knigge 2006](#)), about 10% of the luminosity of the donor in J0140. In J0140, the donor dominates. The system has a lower accretion rate – and correspondingly, a fainter disk – than ordinary CVs of its period. This compensates for its unusually high donor luminosity, leading to an unremarkable total absolute magnitude.

¹³For the 100 pc sample, we adopt the photometric and astrometric cuts suggested by [Lindgren et al. \(2018\)](#). For the stars within 2 degrees, we use the same quality cuts but only require `parallax_over_error` > 5 and `parallax` > 0.5 .

Bibliography

- 1997, ESA Special Publication, Vol. 1200, The HIPPARCOS and TYCHO catalogues. Astrometric and photometric star catalogues derived from the ESA HIPPARCOS Space Astrometry Mission
- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016, *Phys. Rev. D*, **93**, 122003
- . 2017, *Phys. Rev. Lett.*, **119**, 161101
- . 2019, *ApJ*, **882**, L24
- Abdul-Masih, M., Banyard, G., Bodensteiner, J., et al. 2019, arXiv e-prints, arXiv:1912.04092
- . 2020, *Nature*, **580**, E11
- Abrahams, E. S., Bloom, J. S., Mowlavi, N., et al. 2020, arXiv e-prints, arXiv:2011.12253
- Abrams, N. S., & Takada, M. 2020, arXiv e-prints, arXiv:2006.05578
- Abril, J., Schmidtobreick, L., Ederoclite, A., & López-Sanjuan, C. 2020, *MNRAS*, **492**, L40
- Abt, H. A., & Levy, S. G. 1976, *ApJS*, **30**, 273
- Adams, S. M., Kochanek, C. S., Gerke, J. R., Stanek, K. Z., & Dai, X. 2017, *MNRAS*, **468**, 4968
- Aerts, C., & Eyser, L. 2000, in *Astronomical Society of the Pacific Conference Series*, Vol. 210, *Delta Scuti and Related Stars*, ed. M. Breger & M. Montgomery, 113
- Aerts, C., Simón-Díaz, S., Groot, P. J., & Degroote, P. 2014, *A&A*, **569**, A118
- Aitken, R. G., & Doolittle, E. 1932, *New General Catalogue of Double Stars within 120° of the North Pole*
- Allard, F. 2014, in *IAU Symposium, Vol. 299, Exploring the Formation and Evolution of Planetary Systems*, ed. M. Booth, B. C. Matthews, & J. R. Graham, 271
- Allard, F., Homeier, D., & Freytag, B. 2012, *Philosophical Transactions of the Royal Society of London Series A*, **370**, 2765
- Allende Prieto, C., Majewski, S. R., Schiavon, R., et al. 2008, *Astronomische Nachrichten*, **329**, 1018
- Andersen, J. 1991, *A&A Rev.*, **3**, 91
- Andrews, J. J., Agüeros, M. A., Belczynski, K., et al. 2012, *ApJ*, **757**, 170
- Andrews, J. J., Agüeros, M. A., Gianninas, A., et al. 2015, *ApJ*, **815**, 63
- Andrews, J. J., Chanamé, J., & Agüeros, M. A. 2017, *MNRAS*, **472**, 675
- . 2018a, *Research Notes of the American Astronomical Society*, **2**, 29
- . 2018b, *MNRAS*, **473**, 5393

- Ansdell, M., Williams, J. P., Trapman, L., et al. 2018, *ApJ*, 859, 21
- Arenou, F., Luri, X., Babusiaux, C., et al. 2018a, ArXiv e-prints, [arXiv:1804.09375](#)
- . 2018b, *A&A*, 616, A17
- Arias, M. L., Zorec, J., Cidale, L., et al. 2006, *A&A*, 460, 821
- Artymowicz, P., Clarke, C. J., Lubow, S. H., & Pringle, J. E. 1991, *ApJ*, 370, L35
- Ashley, R. P., Marsh, T. R., Breedt, E., et al. 2020, *MNRAS*, 499, 149
- Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013a, *A&A*, 558, A33
- . 2013b, *A&A*, 558, A33
- Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018a, *AJ*, 156, 123
- Astropy Collaboration, Price-Whelan, A. M., SipHocz, B. M., et al. 2018b, *aj*, 156, 123
- Augusteijn, T., van der Hooft, F., de Jong, J. A., & van Paradijs, J. 1996, *A&A*, 311, 889
- Badenes, C., Maoz, D., & Ciardullo, R. 2015, *ApJ*, 804, L25
- Badenes, C., Mazzola, C., Thompson, T. A., et al. 2017, ArXiv e-prints, [arXiv:1711.00660](#)
[astro-ph.SR]
- . 2018, *ApJ*, 854, 147
- Bagnuolo, Jr., W. G., & Gies, D. R. 1991, *ApJ*, 376, 266
- Bahcall, J. N., Hut, P., & Tremaine, S. 1985, *ApJ*, 290, 15
- Bahcall, J. N., & Soneira, R. M. 1981, *ApJ*, 246, 122
- Bailer-Jones, C. A. L., Rybizki, J., Fouesneau, M., Demleitner, M., & Andrae, R. 2020, ArXiv e-prints, [arXiv:2012.05220](#)
- Baluev, R. V. 2009, *MNRAS*, 393, 969
- Banik, I., & Zhao, H. 2018, *MNRAS*, 480, 2660
- Bate, M. R. 2000, *MNRAS*, 314, 33
- . 2009, *MNRAS*, 392, 590
- . 2012, *MNRAS*, 419, 3115
- . 2014, *MNRAS*, 442, 285
- . 2018, *MNRAS*, 475, 5618
- . 2019, *MNRAS*, 484, 2341
- Bate, M. R., & Bonnell, I. A. 1997, *MNRAS*, 285, 33
- Bédard, A., Bergeron, P., Brassard, P., & Fontaine, G. 2020, *ApJ*, 901, 93
- Belczynski, K., Hirschi, R., Kaiser, E. A., et al. 2019, ArXiv e-prints, [arXiv:1911.12357](#)
- Bellm, E. C., Kulkarni, S. R., Graham, M. J., et al. 2019, *PASP*, 131, 018002
- Belokurov, V., Penoyre, Z., Oh, S., et al. 2020, *MNRAS*, 496, 1922
- Bergemann, M., Lind, K., Collet, R., Magic, Z., & Asplund, M. 2012, *MNRAS*, 427, 27
- Bergeron, P., Dufour, P., Fontaine, G., et al. 2019, *ApJ*, 876, 67
- Bergeron, P., Wesemael, F., Dufour, P., et al. 2011, *ApJ*, 737, 28
- Bethe, H. A. 1939, *Physical Review*, 55, 434
- Beuermann, K., Baraffe, I., Kolb, U., & Weichhold, M. 1998, *A&A*, 339, 518
- Bhattacharya, D., & van den Heuvel, E. P. J. 1991, *Phys. Rep.*, 203, 1
- Bildsten, L., Shen, K. J., Weinberg, N. N., & Nelemans, G. 2007, *ApJ*, 662, L95
- Binney, J., & Tremaine, S. 2008a, *Galactic Dynamics: Second Edition* (Princeton University Press)

- . 2008b, *Galactic Dynamics: Second Edition*
- Blasius, T. D., Monnier, J. D., Tuthill, P. G., Danchi, W. C., & Anderson, M. 2012, *MNRAS*, **426**, 2652
- Bodensteiner, J., Shenar, T., Mahy, L., et al. 2020, arXiv e-prints, arXiv:2006.10770
- Bohlin, R. C., Mészáros, S., Fleming, S. W., et al. 2017, *AJ*, **153**, 234
- Bonatto, C., Lima, E. F., & Bica, E. 2012, *A&A*, **540**, A137
- Bovy, J. 2016a, *ApJ*, **817**, 49
- . 2016b, *ApJ*, **817**, 49
- Branch, D. 1976, *ApJ*, **210**, 392
- Branch, D., Livio, M., Yungelson, L. R., Boffi, F. R., & Baron, E. 1995, *PASP*, **107**, 1019
- Branch, M. A., Coleman, T. F., & Li, Y. 1999, *SIAM Journal on Scientific Computing*, **21**, 1
- Brandeker, A., & Cataldi, G. 2019a, *A&A*, **621**, A86
- . 2019b, *A&A*, **621**, A86
- Brandt, T. D. 2018, *ApJS*, **239**, 31
- Breedt, E., Gänsicke, B. T., Drake, A. J., et al. 2014, *MNRAS*, **443**, 3174
- Bressan, A., Marigo, P., Girardi, L., et al. 2012, *MNRAS*, **427**, 127
- Brinch, C., Jørgensen, J. K., Hogerheijde, M. R., Nelson, R. P., & Gressel, O. 2016, *ApJ*, **830**, L16
- Brooks, J., Bildsten, L., Marchant, P., & Paxton, B. 2015, *ApJ*, **807**, 74
- Brown, W. R., Kilic, M., Allende Prieto, C., & Kenyon, S. J. 2010, *ApJ*, **723**, 1072
- Brown, W. R., Kilic, M., Kenyon, S. J., & Gianninas, A. 2016, *ApJ*, **824**, 46
- Brown, W. R., Kilic, M., Kosakowski, A., et al. 2020, *ApJ*, **889**, 49
- Bruch, A. 1992, *A&A*, **266**, 237
- Buder, S., Asplund, M., Duong, L., et al. 2018, *MNRAS*, **478**, 4513
- Buder, S., Sharma, S., Kos, J., et al. 2020, arXiv e-prints, arXiv:2011.02505
- Burdge, K. B., Fuller, J., Phinney, E. S., et al. 2019, *ApJ*, **886**, L12
- Burdge, K. B., Prince, T. A., Fuller, J., et al. 2020, *ApJ*, **905**, 32
- Burgasser, A. J. 2007, *AJ*, **134**, 1330
- Burgasser, A. J., Liu, M. C., Ireland, M. J., Cruz, K. L., & Dupuy, T. J. 2008, *ApJ*, **681**, 579
- Burkert, A., Bate, M. R., & Bodenheimer, P. 1997, *MNRAS*, **289**, 497
- Buta, R. J., & Smith, M. A. 1979, *ApJ*, **232**, 213
- Caballero, J. A. 2009, *A&A*, **507**, 251
- Calamida, A., Corsi, C. E., Bono, G., et al. 2008, *Mem. Soc. Astron. Italiana*, **79**, 347
- Cantrell, A. G., & Dougan, T. J. 2014, *MNRAS*, **445**, 2028
- Cardelli, J. A., Clayton, G. C., & Mathis, J. S. 1989, *ApJ*, **345**, 245
- Carney, B. W. 1983, *AJ*, **88**, 623
- Carr, B. J., & Sakellariadou, M. 1999, *ApJ*, **516**, 195
- Casares, J., Negueruela, I., Ribó, M., et al. 2014, *Nature*, **505**, 378
- Catalán, S., Isern, J., García-Berro, E., et al. 2008, *A&A*, **477**, 213
- Chambers, K. C., Magnier, E. A., Metcalfe, N., et al. 2016, arXiv e-prints, arXiv:1612.05560

- Chanamé, J., & Gould, A. 2004, *ApJ*, **601**, 289
- Chanamé, J., & Ramírez, I. 2012, *ApJ*, **746**, 102
- Chanamé, J., & Ramírez, I. 2012, *ApJ*, **746**, 102
- Chapellier, E., Sadsaoud, H., Valtier, J. C., et al. 1998, *A&A*, **331**, 1046
- Chen, X., & Han, Z. 2008, *MNRAS*, **387**, 1416
- Chen, Y., Girardi, L., Bressan, A., et al. 2014, *MNRAS*, **444**, 2525
- Cheng, S., Cummings, J. D., Ménard, B., & Toonen, S. 2020, *ApJ*, **891**, 160
- Choi, J., Conroy, C., Ting, Y.-S., et al. 2018, *ApJ*, **863**, 65
- Choi, J., Dotter, A., Conroy, C., et al. 2016, *ApJ*, **823**, 102
- Chojnowski, S. D., Nguyen, D. C., Nidever, D. L., et al. 2015, in American Astronomical Society Meeting Abstracts, Vol. 225, American Astronomical Society Meeting Abstracts, 340.05
- Chojnowski, S. D., Labadie-Bartz, J., Rivinius, T., et al. 2018, *ApJ*, **865**, 76
- Claret, A. 2015, *A&A*, **577**, A87
- Clark, G. W. 1975, *ApJ*, **199**, L143
- Clarke, C. J. 2020, *MNRAS*, **491**, L72
- Clayton, G. C. 2012, *Journal of the American Association of Variable Star Observers (JAAVSO)*, **40**, 539
- Close, L. M., Richer, H. B., & Crabtree, D. R. 1990, *AJ*, **100**, 1968
- Comerón, F., Reipurth, B., Yen, H.-W., & Connelley, M. S. 2018, *A&A*, **612**, A73
- Connelley, M. S., Reipurth, B., & Tokunaga, A. T. 2008, *AJ*, **135**, 2526
- Coronado, J., Rix, H.-W., Trick, W. H., et al. 2020, *MNRAS*, **495**, 4098
- Coronado, J., Sepúlveda, M. P., Gould, A., & Chanamé, J. 2018, ArXiv e-prints, [arXiv:1804.06851 \[astro-ph.SR\]](https://arxiv.org/abs/1804.06851)
- Correa-Otto, J. A., Calandra, M. F., & Gil-Hutton, R. A. 2017, *A&A*, **600**, A59
- Córsico, A. H., Althaus, L. G., Serenelli, A. M., et al. 2016, *A&A*, **588**, A74
- Cox, A. N. 2000, *Allen's astrophysical quantities*
- Crawford, J. A. 1955, *ApJ*, **121**, 71
- Cui, X.-Q., Zhao, Y.-H., Chu, Y.-Q., et al. 2012a, *Research in Astronomy and Astrophysics*, **12**, 1197
- . 2012b, *Research in Astronomy and Astrophysics*, **12**, 1197
- Cummings, J. D., Kalirai, J. S., Tremblay, P. E., Ramirez-Ruiz, E., & Choi, J. 2018, *ApJ*, **866**, 21
- Czekala, I., Chiang, E., Andrews, S. M., et al. 2019, arXiv e-prints, arXiv:1906.03269
- Czekala, I., Mandel, K. S., Andrews, S. M., et al. 2017, *ApJ*, **840**, 49
- Dachs, J., Eichendorf, W., Schleicher, H., et al. 1981, *A&AS*, **43**, 427
- Dachs, J., Hummel, W., & Hanuschik, R. W. 1992, *A&AS*, **95**, 437
- Dachs, J., Rohe, D., & Loose, A. S. 1990, *A&A*, **238**, 227
- Davis, D. S., Richer, H. B., King, I. R., et al. 2008, *MNRAS*, **383**, L20
- Davis, S., Richer, H. B., Coffey, J., et al. 2006, in *Bulletin of the American Astronomical Society*, Vol. 38, American Astronomical Society Meeting Abstracts, 1215
- de Jager, C., & Nieuwenhuijzen, H. 1987, *A&A*, **177**, 217

- de Mink, S. E., Sana, H., Langer, N., Izzard, R. G., & Schneider, F. R. N. 2014, *ApJ*, 782, 7
- De Rosa, R. J., Patience, J., Wilson, P. A., et al. 2014, *MNRAS*, 437, 1216
- De Silva, G. M., Freeman, K. C., Bland-Hawthorn, J., et al. 2015, *MNRAS*, 449, 2604
- de Val-Borro, M., Gahm, G. F., Stempels, H. C., & Pepliński, A. 2011, *MNRAS*, 413, 2679
- Desidera, S., Gratton, R. G., Scuderi, S., et al. 2004, *A&A*, 420, 683
- Dhital, S., West, A. A., Stassun, K. G., & Bochanski, J. J. 2010, *AJ*, 139, 2566
- Dhital, S., West, A. A., Stassun, K. G., Schluns, K. J., & Massey, A. P. 2015, *AJ*, 150, 57
- Díaz, C. G., González, J. F., Levato, H., & Grosso, M. 2011, *A&A*, 531, A143
- Dolphin, A. E. 2002, *MNRAS*, 332, 91
- Donati, J. F., & Landstreet, J. D. 2009, *ARA&A*, 47, 333
- Doroshenko, V., Ducci, L., Santangelo, A., & Sasaki, M. 2014, *A&A*, 567, A7
- Dorval, J., Boily, C. M., Moraux, E., & Roos, O. 2017, *MNRAS*, 465, 2198
- Dotter, A. 2016, *ApJS*, 222, 8
- Dotter, A., Conroy, C., Cargile, P., & Asplund, M. 2017, *ApJ*, 840, 99
- Drake, A. J., Djorgovski, S. G., Mahabal, A., et al. 2009, *ApJ*, 696, 870
- Driebe, T., Schoenberner, D., Bloeker, T., & Herwig, F. 1998, *A&A*, 339, 123
- Duchêne, G. 1999, *A&A*, 341, 547
- Duchêne, G., & Kraus, A. 2013, *ARA&A*, 51, 269
- Duchêne, G., Lacour, S., Moraux, E., Goodwin, S., & Bouvier, J. 2018, *MNRAS*, 478, 1825
- Dulaney, N. A., Richardson, N. D., Gerhartz, C. J., et al. 2017, *ApJ*, 836, 112
- Dupuy, T. J., & Liu, M. C. 2011, *ApJ*, 733, 122
- Duquennoy, A., & Mayor, M. 1991, *A&A*, 248, 485
- Dziembowski, W. A., Moskalik, P., & Pamyatnykh, A. A. 1993, *MNRAS*, 265, 588
- Eddington, A. S. 1924, *MNRAS*, 84, 308
- . 1926, *The Internal Constitution of the Stars*
- Eggenberger, A., Halbwegs, J.-L., Udry, S., & Mayor, M. 2004, in *Revista Mexicana de Astronomia y Astrofisica*, vol. 27, Vol. 21, *Revista Mexicana de Astronomia y Astrofisica Conference Series*, ed. C. Allen & C. Scarfe, 28
- Eggleton, P. P. 1983, *ApJ*, 268, 368
- Eggleton, P. P., Fitchett, M. J., & Tout, C. A. 1989, *ApJ*, 347, 998
- Eisenstein, D. J., Liebert, J., Harris, H. C., et al. 2006, *ApJS*, 167, 40
- Eisner, J. A., Arce, H. G., Ballering, N. P., et al. 2018, *ApJ*, 860, 77
- El-Badry, K. 2019, *MNRAS*, 482, 5018
- El-Badry, K., & Quataert, E. 2020, *MNRAS*, 493, L22
- El-Badry, K., & Rix, H.-W. 2018a, *MNRAS*, arXiv:1807.06011 [astro-ph.SR]
- . 2018b, *MNRAS*, 480, 4884
- . 2018c, *MNRAS*, 480, 4884
- . 2019a, *MNRAS*, 482, L139
- . 2019b, *MNRAS*, 482, L139
- El-Badry, K., Rix, H.-W., & Heintz, T. M. 2021, arXiv e-prints, arXiv:2101.05282
- El-Badry, K., Rix, H.-W., Tian, H., Duchêne, G., & Moe, M. 2019, *MNRAS*, 489, 5822
- El-Badry, K., Rix, H.-W., Ting, Y.-S., et al. 2018a, *MNRAS*, 473, 5043

- . 2018b, *MNRAS*, **473**, 5043
- . 2018c, *MNRAS*, **473**, 5043
- El-Badry, K., Rix, H.-W., & Weisz, D. R. 2018d, *ApJ*, **860**, L17
- . 2018e, *ApJ*, **860**, L17
- . 2018f, *ApJ*, **860**, L17
- El-Badry, K., Ting, Y.-S., Rix, H.-W., et al. 2018g, *MNRAS*, **476**, 528
- Eldridge, J. J., Stanway, E. R., Breivik, K., et al. 2019, arXiv e-prints, arXiv:1912.03599
- . 2020, *MNRAS*
- Eldridge, J. J., Stanway, E. R., Xiao, L., et al. 2017, *PASA*, **34**, e058
- Evans, D. W., Riello, M., De Angeli, F., et al. 2018, ArXiv e-prints, [arXiv:1804.09368](https://arxiv.org/abs/1804.09368) [[astro-ph.IM](#)]
- Fabrizius, C., Luri, X., Arenou, F., et al. 2020, arXiv e-prints, arXiv:2012.06242
- Faherty, J. K., Bochanski, J. J., Gagné, J., et al. 2018, *ApJ*, **863**, 91
- Famaey, B., & McGaugh, S. S. 2012, *Living Reviews in Relativity*, **15**, 10
- Farihi, J. 2016, *New A Rev.*, **71**, 9
- Farris, B. D., Duffell, P., MacFadyen, A. I., & Haiman, Z. 2014, *ApJ*, **783**, 134
- Fellhauer, M., Lin, D. N. C., Bolte, M., Aarseth, S. J., & Williams, K. A. 2003, *ApJ*, **595**, L53
- Fernandez, M. A., Covey, K. R., De Lee, N., et al. 2017, *PASP*, **129**, 084201
- Fischer, D. A., & Marcy, G. W. 1992, *ApJ*, **396**, 178
- Flewelling, H. A., Magnier, E. A., Chambers, K. C., et al. 2016, arXiv e-prints, arXiv:1612.05243
- Ford, E. B., Kozinsky, B., & Rasio, F. A. 2000, *ApJ*, **535**, 385
- Foreman-Mackey, D. 2016, *The Journal of Open Source Software*, **24**
- Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013a, *PASP*, **125**, 306
- . 2013b, *PASP*, **125**, 306
- . 2013c, *PASP*, **125**, 306
- Fouesneau, M., Rix, H.-W., von Hippel, T., Hogg, D. W., & Tian, H. 2018, ArXiv e-prints, [arXiv:1802.06663](https://arxiv.org/abs/1802.06663) [[astro-ph.SR](#)]
- . 2019, *ApJ*, **870**, 9
- Freeman, K., & Bland-Hawthorn, J. 2002, *ARA&A*, **40**, 487
- Fregeau, J. M., Richer, H. B., Rasio, F. A., & Hurley, J. R. 2009, *ApJ*, **695**, L20
- Gaia Collaboration, Brown, A. G. A., Vallenari, A., et al. 2018a, ArXiv e-prints, [arXiv:1804.09365](https://arxiv.org/abs/1804.09365)
- . 2020a, arXiv e-prints, arXiv:2012.01533
- Gaia Collaboration, Prusti, T., de Bruijne, J. H. J., et al. 2016, *A&A*, **595**, A1
- Gaia Collaboration, Babusiaux, C., van Leeuwen, F., et al. 2018b, ArXiv e-prints, [arXiv:1804.09378](https://arxiv.org/abs/1804.09378) [[astro-ph.SR](#)]
- Gaia Collaboration, Smart, R. L., Sarro, L. M., et al. 2020b, arXiv e-prints, arXiv:2012.02061
- Gänsicke, B. T., Szkody, P., de Martino, D., et al. 2003, *ApJ*, **594**, 443
- Gao, S., Zhao, H., Yang, H., & Gao, R. 2017, *MNRAS*, **469**, L68
- García Pérez, A. E., Allende Prieto, C., Holtzman, J. A., et al. 2016, *AJ*, **151**, 144

- Garnavich, P. M. 1988, [ApJ](#), **335**, L47
- Gasque, L. C., Hening, C. A., Hviding, R. E., et al. 2019, [AJ](#), **158**, 156
- Ge, H., Hjellming, M. S., Webbink, R. F., Chen, X., & Han, Z. 2010, [ApJ](#), **717**, 724
- Geha, M., Brown, T. M., Tumlinson, J., et al. 2013, [ApJ](#), **771**, 29
- Ghez, A. M., Neugebauer, G., & Matthews, K. 1993, [AJ](#), **106**, 2005
- Giannuzzi, M. A. 1987, [Ap&SS](#), **135**, 245
- Gies, D. R., Bagnuolo, William G., J., Ferrara, E. C., et al. 1998, [ApJ](#), **493**, 440
- Giesers, B., Dreizler, S., Husser, T.-O., et al. 2018, [MNRAS](#), **475**, L15
- Giesers, B., Kamann, S., Dreizler, S., et al. 2019, [A&A](#), **632**, A3
- Glebocki, R., Gnacinski, P., & Stawikowski, A. 2000, *Acta Astron.*, **50**, 509
- Godoy-Rivera, D., & Chanamé, J. 2018, ArXiv e-prints, [arXiv:1807.00009 \[astro-ph.SR\]](#)
- Goliasch, J., & Nelson, L. 2015, [ApJ](#), **809**, 80
- González, J. F., & Levato, H. 2006, [A&A](#), **448**, 283
- Goodricke, J. 1783, *Philosophical Transactions of the Royal Society of London Series I*, **73**, 474
- Goodwin, S. P. 2010, *Philosophical Transactions of the Royal Society of London Series A*, **368**, 851
- Gould, A., Bahcall, J. N., Maoz, D., & Yanny, B. 1995, [ApJ](#), **441**, 200
- Gould, A., & Chanamé, J. 2004, [ApJS](#), **150**, 455
- Gould, A., & Eastman, J. 2006, ArXiv Astrophysics e-prints, [astro-ph/0610799](#)
- Graham, M. J., Kulkarni, S. R., Bellm, E. C., et al. 2019, [PASP](#), **131**, 078001
- Gray, D. F. 1992a, *The observation and analysis of stellar photospheres.*, Vol. 20
- . 1992b, *The observation and analysis of stellar photospheres.*, Vol. 20
- Green, G. M., Schlafly, E., Zucker, C., Speagle, J. S., & Finkbeiner, D. 2019, [ApJ](#), **887**, 93
- Greenstein, J. L. 1986, [AJ](#), **92**, 867
- Grevesse, N., & Sauval, A. J. 1998, [Space Sci. Rev.](#), **85**, 161
- Griest, K., Whitmore, J. B., Wolfe, A. M., et al. 2010, [ApJ](#), **708**, 158
- Griffin, R. E. M., & Griffin, R. F. 2010, [MNRAS](#), **402**, 1675
- Groh, J. H., Farrell, E., Meynet, G., et al. 2019, ArXiv e-prints, [arXiv:1912.00994](#)
- Grupp, F. 2004, [A&A](#), **420**, 289
- Gullikson, K., Kraus, A., & Dodson-Robinson, S. 2016, [AJ](#), **152**, 40
- Hachisu, I., Kato, M., & Nomoto, K. 1999, [ApJ](#), **522**, 487
- Hadrava, P. 1995, *A&AS*, **114**, 393
- Hainich, R., Ramachandran, V., Shenar, T., et al. 2019, [A&A](#), **621**, A85
- Halbwachs, J. L. 1988, [Ap&SS](#), **142**, 139
- Halbwachs, J.-L., Kiefer, F., Arenou, F., et al. 2017a, ArXiv e-prints, [arXiv:1710.02301 \[astro-ph.IM\]](#)
- Halbwachs, J. L., Mayor, M., Udry, S., & Arenou, F. 2003, [A&A](#), **397**, 159
- Halbwachs, J.-L., Arenou, F., Boffin, H. M. J., et al. 2017b, ArXiv e-prints, [arXiv:1710.02017 \[astro-ph.SR\]](#)
- Hameury, J. M. 2020, *Advances in Space Research*, **66**, 1004
- Han, Z., & Podsiadlowski, P. 2004, [MNRAS](#), **350**, 1301

- Han, Z., Podsiadlowski, P., Maxted, P. F. L., Marsh, T. R., & Ivanova, N. 2002, *MNRAS*, **336**, 449
- Hanawa, T., Ochi, Y., & Ando, K. 2010, *ApJ*, **708**, 485
- Hanuschik, R. W. 1988, *A&A*, 190, 187
- Hanuschik, R. W., Hummel, W., Sutorius, E., Dietle, O., & Thimm, G. 1996, *A&AS*, 116, 309
- Hartman, Z. D., & Lépine, S. 2020, *ApJS*, **247**, 66
- Hawkins, K., Lucey, M., Ting, Y.-S., et al. 2020, *MNRAS*, **492**, 1164
- Heber, U. 2016, *PASP*, **128**, 082001
- Heggie, D. C. 1975, *MNRAS*, **173**, 729
- Hensberge, H., Ilijčić, S., & Torres, K. B. V. 2008, *A&A*, **482**, 1031
- Hermes, J. J., Gänsicke, B. T., & Breedt, E. 2017, in *Astronomical Society of the Pacific Conference Series*, Vol. 509, 20th European White Dwarf Workshop, ed. P. E. Tremblay, B. Gaensicke, & T. Marsh, 453
- Hermes, J. J., Montgomery, M. H., Winget, D. E., et al. 2012, *ApJ*, **750**, L28
- Hernandez, X., Cortés, R. A. M., Allen, C., & Scarpa, R. 2018, *ArXiv e-prints*, [arXiv:1810.08696](https://arxiv.org/abs/1810.08696)
- Hernandez, X., Jiménez, M. A., & Allen, C. 2012, *European Physical Journal C*, **72**, 1884
- Herschel, W. 1782, *Philosophical Transactions of the Royal Society of London Series I*, 72, 112
- . 1803, *Philosophical Transactions of the Royal Society of London Series I*, 93, 339
- Hertzsprung, E. 1923, *Bull. Astron. Inst. Netherlands*, 2, 15
- Herwig, F. 2000, *A&A*, 360, 952
- Hettinger, T., Badenes, C., Strader, J., Bickerton, S. J., & Beers, T. C. 2015, *ApJ*, **806**, L2
- Heyl, J. 2007, *MNRAS*, **381**, L70
- . 2008a, *MNRAS*, **390**, 622
- Heyl, J., & Penrice, M. 2009, *MNRAS*, **397**, L79
- Heyl, J. S. 2008b, *MNRAS*, **385**, 231
- Hills, J. G. 1975, *AJ*, **80**, 809
- Hinkel, N. R., Timmes, F. X., Young, P. A., Pagano, M. D., & Turnbull, M. C. 2014, *AJ*, **148**, 54
- Hioki, T., Itoh, Y., Oasa, Y., et al. 2007, *AJ*, **134**, 880
- Hjellming, M. S., & Webbink, R. F. 1987, *ApJ*, **318**, 794
- Ho, A. Y. Q., Ness, M. K., Hogg, D. W., et al. 2017, *ApJ*, **836**, 5
- Höfner, S., & Olofsson, H. 2018, *A&A Rev.*, **26**, 1
- Hogeveen, S. J. 1992, *Ap&SS*, **196**, 299
- Hogg, D. W., Casey, A. R., Ness, M., et al. 2016, *ApJ*, **833**, 262
- Holberg, J. B., & Bergeron, P. 2006, *AJ*, **132**, 1221
- Hollands, M. A., Tremblay, P. E., Gaensicke, B. T., Gentile-Fusillo, N. P., & Toonen, S. 2018, *ArXiv e-prints*, [arXiv:1805.12590](https://arxiv.org/abs/1805.12590) [*astro-ph.SR*]
- Holtzman, J. A., Shetrone, M., Johnson, J. A., et al. 2015, *AJ*, **150**, 148
- Horvat, M., Conroy, K. E., Pablo, H., et al. 2018, *ApJS*, **237**, 26

- Howard, A. W., Johnson, J. A., Marcy, G. W., et al. 2010, *ApJ*, **721**, 1467
- Huang, S.-S. 1963, *ApJ*, **138**, 471
- Hubeny, I., & Lanz, T. 1995, *ApJ*, **439**, 875
- . 2011, Synspec: General Spectrum Synthesis Program
- . 2017, arXiv e-prints, arXiv:1706.01859
- Hulse, R. A., & Taylor, J. H. 1975, *ApJ*, **195**, L51
- Hummel, W., & Dachs, J. 1992, *A&A*, **262**, L17
- Hummel, W., & Vrancken, M. 2000, *A&A*, **359**, 1075
- Hunter, J. D. 2007, *Computing In Science & Engineering*, **9**, 90
- Hurley, J., & Tout, C. A. 1998, *MNRAS*, **300**, 977
- Hurley, J. R., Tout, C. A., & Pols, O. R. 2002a, *MNRAS*, **329**, 897
- . 2002b, *MNRAS*, **329**, 897
- Husser, T. O., Wende-von Berg, S., Dreizler, S., et al. 2013, *A&A*, **553**, A6
- Hut, P. 1981, *A&A*, **99**, 126
- Hwang, H.-C., Ting, Y.-S., Schlaufman, K. C., Zakamska, N. L., & Wyse, R. F. G. 2020, *MNRAS*, arXiv:2010.02920 [astro-ph.SR]
- Iben, I. J., & Tutukov, A. V. 1984, *ApJS*, **54**, 335
- Iben, Icko, J., & Tutukov, A. V. 1986, *ApJ*, **311**, 742
- . 1991, *ApJ*, **370**, 615
- Iglesias-Marzoa, R., López-Morales, M., & Jesús Arévalo Morales, M. 2015, *PASP*, **127**, 567
- Ilijic, S. 2004, in *Astronomical Society of the Pacific Conference Series*, Vol. 318, *Spectroscopically and Spatially Resolving the Components of the Close Binary Stars*, ed. R. W. Hilditch, H. Hensberge, & K. Pavlovski, 107
- Ilijic, S. 2017, *FDBinary: A tool for spectral disentangling of double-lined spectroscopic binary stars*
- Ilijic, S., Hensberge, H., Pavlovski, K., & Freyhammer, L. M. 2004, in *Astronomical Society of the Pacific Conference Series*, Vol. 318, *Spectroscopically and Spatially Resolving the Components of the Close Binary Stars*, ed. R. W. Hilditch, H. Hensberge, & K. Pavlovski, 111
- Irrgang, A., Geier, S., Kreuzer, S., Pelisoli, I., & Heber, U. 2019, arXiv e-prints, arXiv:1912.08338
- . 2020, *A&A*, **633**, L5
- Ivanova, N., Justham, S., Chen, X., et al. 2013, *A&A Rev.*, **21**, 59
- Izzard, R. G., Dermine, T., & Church, R. P. 2010, *A&A*, **523**, A10
- Jao, W.-C., Mason, B. D., Hartkopf, W. I., Henry, T. J., & Ramos, S. N. 2009, *AJ*, **137**, 3800
- Jaschek, C., & Jaschek, M. 1959, *ZAp*, **48**, 263
- . 1983, *A&A*, **117**, 357
- Jaschek, M., & Egret, D. 1982, in *IAU Symposium*, Vol. 98, *Be Stars*, ed. M. Jaschek & H. G. Groth, 261
- Jeans, J. H. 1919, *MNRAS*, **79**, 408
- . 1924, *MNRAS*, **85**, 2

- . 1925, *MNRAS*, **85**, 914
- Jetsu, L., Porceddu, S., Lyytinen, J., et al. 2013, *ApJ*, **773**, 1
- Jiang, Y.-F., & Tremaine, S. 2010, *MNRAS*, **401**, 977
- Jurić, M., Ivezić, Ž., Brooks, A., et al. 2008a, *ApJ*, **673**, 864
- Jurić, M., Ivezić, Ž., Brooks, A., et al. 2008b, *ApJ*, **673**, 864
- Justesen, A. B., & Albrecht, S. 2020, *A&A*, **642**, A212
- Kalirai, J. S., Ventura, P., Richer, H. B., et al. 2001, *AJ*, **122**, 3239
- Kalogera, V., & Webbink, R. F. 1996, *ApJ*, **458**, 301
- Kalomeni, B., Nelson, L., Rappaport, S., et al. 2016, *ApJ*, **833**, 83
- Kashi, A., & Soker, N. 2011, *MNRAS*, **417**, 1466
- Kato, T., & Osaki, Y. 2014, *PASJ*, **66**, L5
- Katoh, N., Itoh, Y., Toyota, E., & Sato, B. 2013, *AJ*, **145**, 41
- Katz, B., & Dong, S. 2012, ArXiv e-prints, [arXiv:1211.4584](https://arxiv.org/abs/1211.4584) [astro-ph.SR]
- Katz, D., Sartoretti, P., Cropper, M., et al. 2018, ArXiv e-prints, [arXiv:1804.09372](https://arxiv.org/abs/1804.09372) [astro-ph.IM]
- . 2019, *A&A*, **622**, A205
- Kaufer, A., Stahl, O., Tubbesing, S., et al. 1999, *The Messenger*, **95**, 8
- Kepler, S. O., Pelisoli, I., Koester, D., et al. 2016, *MNRAS*, **455**, 3413
- Kervella, P., Arenou, F., Mignard, F., & Thévenin, F. 2019, *A&A*, **623**, A72
- Khokhlov, S. A., Miroshnichenko, A. S., Zharikov, S. V., et al. 2018, *ApJ*, **856**, 158
- Kilic, M., Hermes, J. J., Gianninas, A., et al. 2014, *MNRAS*, **438**, L26
- Knigge, C. 2006, *MNRAS*, **373**, 484
- Knigge, C., Baraffe, I., & Patterson, J. 2011, *ApJS*, **194**, 28
- Koch, R. H., & Hrivnak, B. J. 1981, *AJ*, **86**, 438
- Kochanek, C. S., Shappee, B. J., Stanek, K. Z., et al. 2017, *PASP*, **129**, 104502
- Koester, D. 1987, *ApJ*, **322**, 852
- . 2010, *Mem. Soc. Astron. Italiana*, **81**, 921
- Koester, D., Gänsicke, B. T., & Farihi, J. 2014, *A&A*, **566**, A34
- Kolb, U., & Ritter, H. 1990, *A&A*, **236**, 385
- Kollmeier, J. A., Zasowski, G., Rix, H.-W., et al. 2017, arXiv e-prints, [arXiv:1711.03234](https://arxiv.org/abs/1711.03234)
- Konacki, M. 2005, *ApJ*, **626**, 431
- Kopal, Z. 1956, *Annales d’Astrophysique*, **19**, 298
- Kordopatis, G., Gilmore, G., Steinmetz, M., et al. 2013, *AJ*, **146**, 134
- Koubský, P., Kotková, L., Votruba, V., Šlechta, M., & Dvořáková, Š. 2012, *A&A*, **545**, A121
- Kounkel, M., Covey, K., Moe, M., et al. 2019, *AJ*, **157**, 196
- Kouwenhoven, M. B. N., Goodwin, S. P., Parker, R. J., et al. 2010, *MNRAS*, **404**, 1835
- Kowalski, P. M., & Saumon, D. 2006, *ApJ*, **651**, L137
- Kozai, Y. 1962, *AJ*, **67**, 591
- Kraft, R. P. 1962, *ApJ*, **135**, 408
- . 1967, *ApJ*, **150**, 551
- Kratter, K. M., Matzner, C. D., Krumholz, M. R., & Klein, R. I. 2010, *ApJ*, **708**, 1585
- Kraus, A. L., Ireland, M. J., Martinache, F., & Hillenbrand, L. A. 2011, *ApJ*, **731**, 8

- Kroupa, P. 1995, *MNRAS*, 277, 1507
- Kroupa, P. 1998, *MNRAS*, 298, 231
- Kroupa, P. 2001a, *MNRAS*, 322, 231
- . 2001b, *MNRAS*, 322, 231
- Kroupa, P., & Burkert, A. 2001, *ApJ*, 555, 945
- Krumholz, M. R., & Thompson, T. A. 2007, *ApJ*, 661, 1034
- Kuiper, G. P. 1935, *PASP*, 47, 15
- . 1938, *ApJ*, 88, 472
- . 1941, *ApJ*, 93, 133
- Kunder, A., Kordopatis, G., Steinmetz, M., et al. 2017, *AJ*, 153, 75
- Kurucz, R. L. 1970, SAO Special Report, 309
- . 1979, *ApJS*, 40, 1
- Kurucz, R. L. 1992, in IAU Symposium, Vol. 149, The Stellar Populations of Galaxies, ed. B. Barbuy & A. Renzini, 225
- . 1993, SYNTHE spectrum synthesis programs and line data
- Lam, C. Y., Lu, J. R., Hosek, Matthew W., J., Dawson, W. A., & Golovich, N. R. 2020, *ApJ*, 889, 31
- Lanz, T., & Hubeny, I. 2007, *ApJS*, 169, 83
- Latham, D. W., Stefanik, R. P., Torres, G., et al. 2002, *AJ*, 124, 1144
- Latham, D. W., Tonry, J., Bahcall, J. N., Soneira, R. M., & Schechter, P. 1984, *ApJ*, 281, L41
- Law, N. M., Kulkarni, S. R., Dekany, R. G., et al. 2009, *PASP*, 121, 1395
- Leinert, C., Zinnecker, H., Weitzel, N., et al. 1993, *A&A*, 278, 129
- Lépine, S., & Bongiorno, B. 2007a, *AJ*, 133, 889
- . 2007b, *AJ*, 133, 889
- Li, C., de Grijs, R., & Deng, L. 2013, *MNRAS*, 436, 1497
- Li, Z., Chen, X., Chen, H.-L., & Han, Z. 2019, *ApJ*, 871, 148
- Lindgren, L., Hernandez, J., Bombrun, A., et al. 2018, ArXiv e-prints, [arXiv:1804.09366](https://arxiv.org/abs/1804.09366) [[astro-ph.IM](https://arxiv.org/abs/1804.09366)]
- Lindgren, L., Bastian, U., Biermann, M., et al. 2020a, arXiv e-prints, [arXiv:2012.01742](https://arxiv.org/abs/2012.01742)
- Lindgren, L., Klioner, S. A., Hernández, J., et al. 2020b, arXiv e-prints, [arXiv:2012.03380](https://arxiv.org/abs/2012.03380)
- Littlefair, S. P., Dhillon, V. S., Marsh, T. R., & Gänsicke, B. T. 2006, *MNRAS*, 371, 1435
- Liu, J., Zhang, H., Howard, A. W., et al. 2019, arXiv e-prints, [arXiv:1911.11989](https://arxiv.org/abs/1911.11989)
- Liu, J., Zheng, Z., Soria, R., et al. 2020, arXiv e-prints, [arXiv:2005.12595](https://arxiv.org/abs/2005.12595)
- Livne, E. 1990, *ApJ*, 354, L53
- Lodieu, N., Zapatero Osorio, M. R., & Martín, E. L. 2009, *A&A*, 499, 729
- Lomax, O., Whitworth, A. P., Hubber, D. A., Stamatellos, D., & Walch, S. 2015, *MNRAS*, 447, 1550
- Longhitano, M., & Binggeli, B. 2010, *A&A*, 509, A46
- Lucy, L. B. 1967, *ZAp*, 65, 89
- . 2006, *A&A*, 457, 629
- Lucy, L. B., & Ricco, E. 1979, *AJ*, 84, 401

- Lutz, T. E., & Kelker, D. H. 1973, *PASP*, **85**, 573
- Luyten, W. J. 1971, *Ap&SS*, **11**, 49
- 1979a, New Luyten catalogue of stars with proper motions larger than two tenths of an arcsecond; and first supplement; NLTT. (Minneapolis (1979)); Label 12 = short description; Label 13 = documentation by Warren; Label 14 = catalogue
- 1979b, Proper Motion Survey, University of Minnesota, 52
- 1988, *Ap&SS*, **142**, 17
- Machida, M. N. 2008, *ApJ*, **682**, L1
- Machida, M. N., Omukai, K., Matsumoto, T., & Inutsuka, S.-I. 2009, *MNRAS*, **399**, 1255
- Maíz Apellániz, J., Pantaleoni González, M., & Barbá, R. H. 2021, arXiv e-prints, arXiv:2101.10206
- Majewski, S. R., Schiavon, R. P., Frinchaboy, P. M., et al. 2017, *AJ*, **154**, 94
- Makarov, V. V., Zacharias, N., & Hennessy, G. S. 2008, *ApJ*, **687**, 566
- Mallada, E., & Fernandez, J. A. 2001, in *Revista Mexicana de Astronomia y Astrofisica*, vol. 27, Vol. 11, *Revista Mexicana de Astronomia y Astrofisica Conference Series*, 27
- Margon, B., Ford, H. C., Katz, J. I., et al. 1979, *ApJ*, **230**, L41
- Marks, M., & Kroupa, P. 2012, *A&A*, **543**, A8
- Marks, M., Kroupa, P., & Oh, S. 2011, *MNRAS*, **417**, 1684
- Marsh, T. R., Nelemans, G., & Steeghs, D. 2004, *MNRAS*, **350**, 113
- Masci, F. J., Laher, R. R., Rusholme, B., et al. 2019, *PASP*, **131**, 018003
- Masuda, K., & Hirano, T. 2021, arXiv e-prints, arXiv:2103.05216
- Mathieu, R. D. 1994, *ARA&A*, **32**, 465
- Matijevič, G., Zwitter, T., Munari, U., et al. 2010, *AJ*, **140**, 184
- Matijevič, G., Zwitter, T., Bienaymé, O., et al. 2011, *AJ*, **141**, 200
- Matsumoto, T., Saigo, K., & Takakuwa, S. 2019, *ApJ*, **871**, 36
- Maxted, P. F. L., Serenelli, A. M., Marsh, T. R., et al. 2014, *MNRAS*, **444**, 208
- Mazeh, T., Goldberg, D., Duquennoy, A., & Mayor, M. 1992, *ApJ*, **401**, 265
- Mazeh, T., Simon, M., Prato, L., Markus, B., & Zucker, S. 2003, *ApJ*, **599**, 1344
- McClintock, J. E., & Remillard, R. A. 2006, *Black hole binaries*, Vol. 39, 157
- McLaughlin, D. B. 1924, *ApJ*, **60**, 22
- Menou, K., Narayan, R., & Lasota, J.-P. 1999, *ApJ*, **513**, 811
- Merle, T., Van Eck, S., Jorissen, A., et al. 2017, ArXiv e-prints, arXiv:1707.01720 [[astro-ph.SR](#)]
- Michalik, D., Lindegren, L., & Hobbs, D. 2015, *A&A*, **574**, A115
- Milne, E. A. 1930, *MNRAS*, **91**, 4
- Minor, Q. E. 2013, *ApJ*, **779**, 116
- Miranda, R., Muñoz, D. J., & Lai, D. 2017, *MNRAS*, **466**, 1170
- Moe, M., & Di Stefano, R. 2013, *ApJ*, **778**, 95
- 2017a, *ApJS*, **230**, 15
- 2017b, *ApJS*, **230**, 15
- Moe, M., & Kratter, K. M. 2018, *ApJ*, **854**, 44
- Moe, M., Kratter, K. M., & Badenes, C. 2018, ArXiv e-prints, arXiv:1808.02116 [[astro-ph.SR](#)]

- Moeckel, N., & Bate, M. R. 2010, *MNRAS*, 404, 721
- Moeckel, N., & Clarke, C. J. 2011, *MNRAS*, 415, 1179
- Monroy-Rodríguez, M. A., & Allen, C. 2014, *ApJ*, 790, 159
- Moody, M. S. L., Shi, J.-M., & Stone, J. M. 2019, *ApJ*, 875, 66
- Morrissey, P., Conrow, T., Barlow, T. A., et al. 2007, *ApJS*, 173, 682
- Muñoz, D. J., Miranda, R., & Lai, D. 2019, *ApJ*, 871, 84
- Muno, M. P., & Mauerhan, J. 2006, *ApJ*, 648, L135
- Murphy, S. J., Moe, M., Kurtz, D. W., et al. 2018, *MNRAS*, 474, 4322
- Murray, C. D., & Correia, A. C. M. 2010, *Keplerian Orbits and Dynamics of Exoplanets*, ed. S. Seager, 15
- Naoz, S., Farr, W. M., Lithwick, Y., Rasio, F. A., & Teyssandier, J. 2013, *MNRAS*, 431, 2155
- Narayan, R., Paczynski, B., & Piran, T. 1992, *ApJ*, 395, L83
- Nelson, A. F., & Marzari, F. 2016, *ApJ*, 827, 93
- Ness, M., Hogg, D. W., Rix, H.-W., Ho, A. Y. Q., & Zasowski, G. 2015a, *ApJ*, 808, 16
- Ness, M., Hogg, D. W., Rix, H. W., Ho, A. Y. Q., & Zasowski, G. 2015b, *ApJ*, 808, 16
- Neustroev, V. V., & Zharikov, S. 2008, *MNRAS*, 386, 1366
- Nidever, D. L., Holtzman, J. A., Allende Prieto, C., et al. 2015, *AJ*, 150, 173
- Nieva, M. F., & Przybilla, N. 2007, *A&A*, 467, 295
- . 2012, *A&A*, 539, A143
- Nordström, B., Mayor, M., Andersen, J., et al. 2004, *A&A*, 418, 989
- Ochi, Y., Sugimoto, K., & Hanawa, T. 2005, *ApJ*, 623, 922
- Oelkers, R. J., Stassun, K. G., & Dhital, S. 2017, *AJ*, 153, 259
- Offner, S. S. R., Kratter, K. M., Matzner, C. D., Krumholz, M. R., & Klein, R. I. 2010, *ApJ*, 725, 1485
- Oh, S., Price-Whelan, A. M., Hogg, D. W., Morton, T. D., & Spergel, D. N. 2017a, *AJ*, 153, 257
- . 2017b, *AJ*, 153, 257
- Okazaki, A. T. 1991, *PASJ*, 43, 75
- Öpik, E. 1924, *Publications of the Tartu Astrofizica Observatory*, 25, 1
- Osaki, Y. 1996, *PASP*, 108, 39
- Paczynski, B. 1976, in *IAU Symposium, Vol. 73, Structure and Evolution of Close Binary Systems*, ed. P. Eggleton, S. Mitton, & J. Whelan, 75
- Paczynski, B. 1977, *ApJ*, 216, 822
- Pakmor, R., Kromer, M., Taubenberger, S., & Springel, V. 2013, *ApJ*, 770, L8
- Parker, R. J., Goodwin, S. P., Kroupa, P., & Kouwenhoven, M. B. N. 2009, *MNRAS*, 397, 1577
- Parker, R. J., & Reggiani, M. M. 2013, *MNRAS*, 432, 2378
- Patterson, J. 1984, *ApJS*, 54, 443
- Patterson, J., Kemp, J., Harvey, D. A., et al. 2005, *PASP*, 117, 1204
- Pavlovski, K., & Hensberge, H. 2010, in *Astronomical Society of the Pacific Conference Series, Vol. 435, Binaries - Key to Comprehension of the Universe*, ed. A. Prša & M. Zejda,

207

- Paxton, B., Bildsten, L., Dotter, A., et al. 2011, [ApJS](#), **192**, 3
- Paxton, B., Cantiello, M., Arras, P., et al. 2013, [ApJS](#), **208**, 4
- Paxton, B., Marchant, P., Schwab, J., et al. 2015, [ApJS](#), **220**, 15
- Paxton, B., Schwab, J., Bauer, E. B., et al. 2018, [ApJS](#), **234**, 34
- Paxton, B., Smolec, R., Schwab, J., et al. 2019, [ApJS](#), **243**, 10
- Pelisoli, I., Bell, K. J., Kepler, S. O., & Koester, D. 2019, [MNRAS](#), **482**, 3831
- Pelisoli, I., Kepler, S. O., & Koester, D. 2018a, [MNRAS](#), **475**, 2480
- Pelisoli, I., Kepler, S. O., Koester, D., et al. 2018b, [MNRAS](#), **478**, 867
- Pelisoli, I., & Vos, J. 2019, [MNRAS](#), **488**, 2892
- Perpinyà-Vallès, M., Rebassa-Mansergas, A., Gänsicke, B. T., et al. 2019, [MNRAS](#), **483**, 901
- Peters, G. J., Pewett, T. D., Gies, D. R., Touhami, Y. N., & Grundstrom, E. D. 2013, [ApJ](#), **765**, 2
- Peters, G. J., Wang, L., Gies, D. R., & Grundstrom, E. D. 2016, [ApJ](#), **828**, 47
- Pickering, E. C. 1890, *The Observatory*, **13**, 80
- Piersanti, L., Yungelson, L. R., & Tornambé, A. 2015, [MNRAS](#), **452**, 2897
- Pinsonneault, M. H., & Stanek, K. Z. 2006, [ApJ](#), **639**, L67
- Pittordis, C., & Sutherland, W. 2018, [MNRAS](#), **480**, 1778
- . 2019, [MNRAS](#), **488**, 4740
- Podsiadlowski, P., Han, Z., & Rappaport, S. 2003, [MNRAS](#), **340**, 1214
- Podsiadlowski, P., Rappaport, S., & Pfahl, E. D. 2002, [ApJ](#), **565**, 1107
- Poeckert, R., & Marlborough, J. M. 1979, [ApJ](#), **233**, 259
- Pourbaix, D., Tokovinin, A. A., Batten, A. H., et al. 2004, [A&A](#), **424**, 727
- Poveda, A., & Allen, C. 2004, in *Revista Mexicana de Astronomia y Astrofisica*, vol. 27, Vol. 21, *Revista Mexicana de Astronomia y Astrofisica Conference Series*, ed. C. Allen & C. Scarfe, 49
- Poveda, A., Herrera, M. A., Allen, C., Cordero, G., & Lavalley, C. 1994, *Rev. Mexicana Astron. Astrofis.*, **28**, 43
- Price-Whelan, A. M., & Goodman, J. 2018, [ApJ](#), **867**, 5
- Price-Whelan, A. M., Hogg, D. W., Foreman-Mackey, D., & Rix, H.-W. 2017, [ApJ](#), **837**, 20
- Prochaska, J. X., Hennawi, J. F., Westfall, K. B., et al. 2020, arXiv e-prints, arXiv:2005.06505
- Prša, A., Pepper, J., & Stassun, K. G. 2011, [AJ](#), **142**, 52
- Prša, A., & Zwitter, T. 2005, [ApJ](#), **628**, 426
- Prša, A., Conroy, K. E., Horvat, M., et al. 2016, [ApJS](#), **227**, 29
- Pustynnik, I. 1998, [Astronomical and Astrophysical Transactions](#), **15**, 357
- Pylyser, E., & Savonije, G. J. 1988, [A&A](#), **191**, 57
- Qiu, D., Tian, H.-J., Wang, X.-D., et al. 2020, arXiv e-prints, arXiv:2012.04890
- Quinn, D. P., Wilkinson, M. I., Irwin, M. J., et al. 2009, [MNRAS](#), **396**, L11
- Raghavan, D., McAlister, H. A., Henry, T. J., et al. 2010, [ApJS](#), **190**, 1
- Raguzova, N. V. 2001, [A&A](#), **367**, 848
- Rahoma, W. A., Abd El-Salam, F. A., & Ahmed, M. K. 2009, [Journal of Astrophysics and Astronomy](#), **30**, 187

- Rappaport, S., Podsiadlowski, P., Joss, P. C., Di Stefano, R., & Han, Z. 1995, *MNRAS*, **273**, 731
- Rappaport, S., Verbunt, F., & Joss, P. C. 1983, *ApJ*, **275**, 713
- Rastegaev, D. A. 2010, *AJ*, **140**, 2013
- Rebassa-Mansergas, A., Gänsicke, B. T., Rodríguez-Gil, P., Schreiber, M. R., & Koester, D. 2007, *MNRAS*, **382**, 1377
- Rebassa-Mansergas, A., Parsons, S. G., Copperwheat, C. M., et al. 2014, *ApJ*, **790**, 28
- Reggiani, M. M., & Meyer, M. R. 2011, *ApJ*, **738**, 60
- Reid, I. N., & Gizis, J. E. 1997, *AJ*, **113**, 2246
- Reimers, D. 1975, *Memoires of the Societe Royale des Sciences de Liege*, **8**, 369
- Reipurth, B., Guimarães, M. M., Connelley, M. S., & Bally, J. 2007, *AJ*, **134**, 2272
- Reipurth, B., & Mikkola, S. 2012, *Nature*, **492**, 221
- Reis, I., Poznanski, D., Baron, D., Zasowski, G., & Shahaf, S. 2017, ArXiv e-prints, [arXiv:1711.00022 \[astro-ph.IM\]](https://arxiv.org/abs/1711.00022)
- Remillard, R. A., & McClintock, J. E. 2006, *ARA&A*, **44**, 49
- Ren, J., Luo, A., Li, Y., et al. 2013, *AJ*, **146**, 82
- Retterer, J. M., & King, I. R. 1982, *ApJ*, **254**, 214
- Riaz, B., Gizis, J. E., & Samaddar, D. 2008, *ApJ*, **672**, 1153
- Ricker, G. R., Winn, J. N., Vanderspek, R., et al. 2015, *Journal of Astronomical Telescopes, Instruments, and Systems*, **1**, 014003
- Riello, M., De Angeli, F., Evans, D. W., et al. 2020, arXiv e-prints, [arXiv:2012.01916](https://arxiv.org/abs/2012.01916)
- Rivinius, T., Baade, D., Hadrava, P., Heida, M., & Klement, R. 2020, *A&A*, **637**, L3
- Rivinius, T., Carciofi, A. C., & Martayan, C. 2013, *A&A Rev.*, **21**, 69
- Rix, H.-W., Ting, Y.-S., Conroy, C., & Hogg, D. W. 2016, *ApJ*, **826**, L25
- Robin, A. C., Reylé, C., Derrière, S., & Picaud, S. 2003, *A&A*, **409**, 523
- Rodríguez-Gil, P., Torres, M. A. P., Gänsicke, B. T., et al. 2009, *A&A*, **496**, 805
- Rosseland, S. 1926, *ApJ*, **63**, 342
- Rowell, N., Davidson, M., Lindegren, L., et al. 2020, arXiv e-prints, [arXiv:2012.02069](https://arxiv.org/abs/2012.02069)
- Russell, H. N. 1929, *ApJ*, **70**, 11
- Ryans, R. S. I., Dufton, P. L., Rolleston, W. R. J., et al. 2002, *MNRAS*, **336**, 577
- Rybizki, J., Demleitner, M., Fouesneau, M., et al. 2018, *PASP*, **130**, 074101
- Rybizki, J., Demleitner, M., Bailer-Jones, C., et al. 2020, *PASP*, **132**, 074501
- Sadavoy, S. I., & Stahler, S. W. 2017, *MNRAS*, **469**, 3881
- Salim, S., & Gould, A. 2003, *ApJ*, **582**, 1011
- Sartoretti, P., Katz, D., Cropper, M., et al. 2018, ArXiv e-prints, [arXiv:1804.09371](https://arxiv.org/abs/1804.09371)
- Savedoff, M. P. 1966, *AJ*, **71**, 369
- Scarpa, R., Ottolina, R., Falomo, R., & Treves, A. 2017, *International Journal of Modern Physics D*, **26**, 1750067
- Schatzman, E. 1962, *Annales d'Astrophysique*, **25**, 18
- Schlegel, D. J., Finkbeiner, D. P., & Davis, M. 1998, *ApJ*, **500**, 525
- Schlesinger, K. J., Johnson, J. A., Lee, Y. S., et al. 2010, *ApJ*, **719**, 996
- Schootemeijer, A., Götberg, Y., de Mink, S. E., Gies, D., & Zapartas, E. 2018, *A&A*, **615**,

A30

- Seabroke, G. M., & Gilmore, G. 2007, *MNRAS*, **380**, 1348
- Sesar, B., Ivezić, Ž., & Jurić, M. 2008, *ApJ*, **689**, 1244
- Shahaf, S., & Mazeh, T. 2019, arXiv e-prints, arXiv:1905.13239
- Sharma, S., Bland-Hawthorn, J., Johnston, K. V., & Binney, J. 2011, *ApJ*, **730**, 3
- Sharma, S., Bland-Hawthorn, J., Binney, J., et al. 2014, *ApJ*, **793**, 51
- Sharma, S., Stello, D., Buder, S., et al. 2017, ArXiv e-prints, arXiv:1707.05753 [astro-ph.SR]
- Shatsky, N., & Tokovinin, A. 2002, *A&A*, **382**, 92
- Shaya, E. J., & Olling, R. P. 2011, *ApJS*, **192**, 2
- Shen, K. J. 2015, *ApJ*, **805**, L6
- Shen, K. J., & Bildsten, L. 2014, *ApJ*, **785**, 61
- Shen, R. F., Matzner, C. D., Howard, A. W., & Zhang, W. 2019, arXiv e-prints, arXiv:1911.12581
- Shenar, T., Bodensteiner, J., Abdul-Masih, M., et al. 2020, arXiv e-prints, arXiv:2004.12882
- Shi, J.-M., Krolik, J. H., Lubow, S. H., & Hawley, J. F. 2012, *ApJ*, **749**, 118
- Silvestri, N. M., Oswalt, T. D., & Hawley, S. L. 2002, *AJ*, **124**, 1118
- Simon, K. P., & Sturm, E. 1994, *A&A*, **281**, 286
- Simon, M., & Obbie, R. C. 2009, *AJ*, **137**, 3442
- Simón-Díaz, S., & Herrero, A. 2014, *A&A*, **562**, A135
- Simón-Díaz, S., Maíz Apellániz, J., Lennon, D. J., et al. 2019, arXiv e-prints, arXiv:1912.07255
- Simons, D. A., Henry, T. J., & Kirkpatrick, J. D. 1996, *AJ*, **112**, 2238
- Simpson, J. D., Martell, S. L., Da Costa, G., et al. 2018, ArXiv e-prints, arXiv:1804.05894
- Sion, E. M., Oswalt, T. D., Liebert, J., & Hintzen, P. 1991, *AJ*, **101**, 1476
- Skrutskie, M. F., Cutri, R. M., Stiening, R., et al. 2006, *AJ*, **131**, 1163
- Skumanich, A. 1972, *ApJ*, **171**, 565
- Smak, J. 1969, *Acta Astron.*, **19**, 155
- Smith, D. A., & Dhillon, V. S. 1998, *MNRAS*, **301**, 767
- Smith, J. A. 1997, PhD thesis, FLORIDA INSTITUTE OF TECHNOLOGY
- Smith, M. A., Fitch, W. S., Africano, J. L., et al. 1984, *ApJ*, **282**, 226
- Soberman, G. E., Phinney, E. S., & van den Heuvel, E. P. J. 1997, *A&A*, **327**, 620
- Soderblom, D. R. 2010, *ARA&A*, **48**, 581
- Söderhjelm, S. 2000, *Astronomische Nachrichten*, **321**, 165
- . 2007, *A&A*, **463**, 683
- Solheim, J. E. 2010, *PASP*, **122**, 1133
- Sollima, A. 2019, arXiv e-prints, arXiv:1904.05646
- Spruit, H. C. 1998, *A&A*, **333**, 603
- Spruit, H. C., & Ritter, H. 1983a, *A&A*, **124**, 267
- . 1983b, *A&A*, **124**, 267
- Stamatellos, D., & Whitworth, A. P. 2009, *MNRAS*, **392**, 413
- Stebbins, J. 1910, *ApJ*, **32**, 185
- Steinmetz, M., Zwitter, T., Siebert, A., et al. 2006, *AJ*, **132**, 1645

- Sterzik, M. F., & Durisen, R. H. 1998, *A&A*, 339, 95
- Stone, N., Metzger, B. D., & Loeb, A. 2015, *MNRAS*, 448, 188
- Struve, F. G. W. 1837, *Stellarum duplicium et multiplicium mensurae micrometricae per magnum Fraunhoferi tubum annis a 1824 ad 1837 in Specula Dorpatensi institutae...*
- Sun, M., & Arras, P. 2018, *ApJ*, 858, 14
- Szkody, P., & Mattei, J. A. 1984, *PASP*, 96, 988
- Taam, R. E., & Spruit, H. C. 2001, *ApJ*, 561, 329
- Takakuwa, S., Saigo, K., Matsumoto, T., et al. 2017, *ApJ*, 837, 86
- Tang, Y., MacFadyen, A., & Haiman, Z. 2017, *MNRAS*, 469, 4258
- Tauris, T. M., & van den Heuvel, E. P. J. 2006, *Formation and evolution of compact stellar X-ray sources*, Vol. 39, 623
- Terrien, R. C., Mahadevan, S., Deshpande, R., et al. 2014, *ApJ*, 782, 61
- Thompson, T. A. 2011, *ApJ*, 741, 82
- Thompson, T. A., Kochanek, C. S., Stanek, K. Z., et al. 2019, *Science*, 366, 637
- . 2020, arXiv e-prints, arXiv:2005.07653
- Thorstensen, J. R. 2013, *PASP*, 125, 506
- . 2015, *PASP*, 127, 351
- Thorstensen, J. R., Alper, E. H., & Weil, K. E. 2016, *AJ*, 152, 226
- Thorstensen, J. R., Fenton, W. H., Patterson, J., et al. 2002a, *PASP*, 114, 1117
- Thorstensen, J. R., Fenton, W. H., Patterson, J. O., et al. 2002b, *ApJ*, 567, L49
- Tian, H.-J., El-Badry, K., Rix, H.-W., & Gould, A. 2020, *ApJS*, 246, 4
- Ting, Y.-S., Conroy, C., & Goodman, A. 2015, *ApJ*, 807, 104
- Ting, Y.-S., Conroy, C., & Rix, H.-W. 2016a, *ApJ*, 826, 83
- . 2016b, *ApJ*, 816, 10
- Ting, Y.-S., Conroy, C., Rix, H.-W., & Cargile, P. 2017a, *ApJ*, 843, 32
- . 2018, ArXiv e-prints, arXiv:1804.01530 [astro-ph.SR]
- . 2019, *ApJ*, 879, 69
- Ting, Y.-S., & Rix, H.-W. 2018, arXiv e-prints, arXiv:1808.03278
- Ting, Y.-S., Rix, H.-W., Conroy, C., Ho, A. Y. Q., & Lin, J. 2017b, ArXiv e-prints, arXiv:1708.01758 [astro-ph.SR]
- Tobin, J. J., Chandler, C. J., Wilner, D. J., et al. 2013, *ApJ*, 779, 93
- Tobin, J. J., Looney, L. W., Li, Z.-Y., et al. 2016, *ApJ*, 818, 73
- Tokovinin, A. 2011, *AJ*, 141, 52
- . 2014a, *AJ*, 147, 87
- . 2014b, *AJ*, 147, 87
- . 2017a, *MNRAS*, 468, 3461
- . 2017b, *MNRAS*, 468, 3461
- Tokovinin, A., Hartung, M., & Hayward, T. L. 2010, *AJ*, 140, 510
- Tokovinin, A., & Kiyaveva, O. 2016, *MNRAS*, 456, 2070
- Tokovinin, A., Thomas, S., Sterzik, M., & Udry, S. 2006, *A&A*, 450, 681
- Tokovinin, A. A. 2000, *A&A*, 360, 997
- Tokovinin, A. A., & Smekhov, M. G. 2002a, *A&A*, 382, 118

- . 2002b, [A&A](#), **382**, 118
- Toonen, S., Hamers, A., & Portegies Zwart, S. 2016, [Computational Astrophysics and Cosmology](#), **3**, 6
- Toonen, S., Hollands, M., Gänsicke, B. T., & Boekholt, T. 2017, [A&A](#), **602**, A16
- Toonen, S., Perets, H. B., & Hamers, A. S. 2018, [A&A](#), **610**, A22
- Torres, G., Andersen, J., & Giménez, A. 2010, [A&A Rev.](#), **18**, 67
- Tout, C. A., & Hall, D. S. 1991, [MNRAS](#), **253**, 9
- Townsend, R. H. D. 1997, [MNRAS](#), **284**, 839
- Townsley, D. M., & Gänsicke, B. T. 2009, [ApJ](#), **693**, 1007
- Traven, G., Matijević, G., Zwitter, T., et al. 2017, [ApJS](#), **228**, 24
- Tremblay, P.-E., Bergeron, P., & Gianninas, A. 2011, [ApJ](#), **730**, 128
- Trimble, V. 1974, [AJ](#), **79**, 967
- . 1987, [Astronomische Nachrichten](#), **308**, 343
- . 1990, [MNRAS](#), **242**, 79
- Troup, N. W., Nidever, D. L., De Lee, N., et al. 2016, [AJ](#), **151**, 85
- Tutukov, A. V., Fedorova, A. V., Ergma, E. V., & Yungelson, L. R. 1985, *Soviet Astronomy Letters*, **11**, 52
- Tutukov, A. V., & Yungelson, L. R. 1979, *Acta Astron.*, **29**, 665
- van Albada, T. S. 1968, *Bull. Astron. Inst. Netherlands*, **20**, 57
- van Biesbroeck, G. 1916, [AJ](#), **29**, 173
- van den Heuvel, E. P. J., Bhattacharya, D., Nomoto, K., & Rappaport, S. A. 1992, [A&A](#), **262**, 97
- van den Heuvel, E. P. J., & Tauris, T. M. 2020, arXiv e-prints, arXiv:2005.04896
- van der Sluys, M. V., Verbunt, F., & Pols, O. R. 2005, [A&A](#), **431**, 647
- Van Der Walt, S., Colbert, S. C., & Varoquaux, G. 2011, ArXiv e-prints, [arXiv:1102.1523 \[cs.MS\]](#)
- Veras, D., Wyatt, M. C., Mustill, A. J., Bonsor, A., & Eldridge, J. J. 2011, [MNRAS](#), **417**, 2104
- Verbunt, F., & Phinney, E. S. 1995, [A&A](#), **296**, 709
- Verbunt, F., & Zwaan, C. 1981, [A&A](#), **100**, L7
- Voges, W., Aschenbach, B., Boller, T., et al. 1999, [A&A](#), **349**, 389
- von Zeipel, H. 1924, [MNRAS](#), **84**, 665
- Vrancken, M., Butler, K., & Becker, S. R. 1996, [A&A](#), **311**, 661
- Waelkens, C. 1991, [A&A](#), **246**, 453
- Waelkens, C., Aerts, C., Kestens, E., Grenon, M., & Eyser, L. 1998, [A&A](#), **330**, 215
- Wakamatsu, Y., Thorstensen, J. R., Kojiguchi, N., et al. 2021, [PASJ](#), [arXiv:2102.04104 \[astro-ph.SR\]](#)
- Wang, L., Gies, D. R., & Peters, G. J. 2017, [ApJ](#), **843**, 60
- . 2018, [ApJ](#), **853**, 156
- Warner, B. 1995, *Cataclysmic variable stars*, Vol. 28
- . 2003, *Cataclysmic Variable Stars*
- Wasserman, I., & Weinberg, M. D. 1991, [ApJ](#), **382**, 149

- Webster, B. L., & Murdin, P. 1972, *Nature*, 235, 37
- Wegner, W. 2002, *Baltic Astronomy*, 11, 1
- Weidemann, V. 1977, *A&A*, 59, 411
- Weinberg, M. D., Shapiro, S. L., & Wasserman, I. 1987, *ApJ*, 312, 367
- Whelan, J., & Iben, Icko, J. 1973, *ApJ*, 186, 1007
- White, R. J., & Ghez, A. M. 2001, *ApJ*, 556, 265
- Widmark, A., Leistedt, B., & Hogg, D. W. 2018, *ApJ*, 857, 114
- Willems, B., & Kolb, U. 2004, *A&A*, 419, 1057
- Wilson, R. E., & Devinney, E. J. 1971, *ApJ*, 166, 605
- Winters, J. G., Charbonneau, D., Henry, T. J., et al. 2020, arXiv e-prints, arXiv:2011.09409
- Woosley, S. E. 2017, *ApJ*, 836, 244
- Wright, E. L., Eisenhardt, P. R. M., Mainzer, A. K., et al. 2010, *AJ*, 140, 1868
- Xiang, M., Ting, Y.-S., Rix, H.-W., et al. 2019, *ApJS*, 245, 34
- Yoo, J., Chanamé, J., & Gould, A. 2004, *ApJ*, 601, 311
- Yoon, S. C., Woosley, S. E., & Langer, N. 2010, *ApJ*, 725, 940
- York, D. G., Adelman, J., Anderson, John E., J., et al. 2000, *AJ*, 120, 1579
- Young, M. D., & Clarke, C. J. 2015, *MNRAS*, 452, 3085
- Yu, Z., Thorstensen, J. R., Rappaport, S., et al. 2019, *MNRAS*, 489, 1023
- Zapatero Osorio, M. R., & Martín, E. L. 2004, *A&A*, 419, 167
- Zasowski, G., Johnson, J. A., Frinchaboy, P. M., et al. 2013, *AJ*, 146, 81
- Zhao, G., Zhao, Y.-H., Chu, Y.-Q., Jing, Y.-P., & Deng, L.-C. 2012, *Research in Astronomy and Astrophysics*, 12, 723
- Zhao, H., Li, B., & Bienaymé, O. 2010, *Phys. Rev. D*, 82, 103001
- Ziegler, C., Law, N. M., Baranec, C., et al. 2018, ArXiv e-prints, arXiv:1806.10142 [astro-ph.EP]
- Ziegler, C., Law, N. M., Baranec, C., Riddle, R. L., & Fuchs, J. T. 2015, *ApJ*, 804, 30
- Zijlstra, A. A. 2006, in *IAU Symposium, Vol. 234, Planetary Nebulae in our Galaxy and Beyond*, ed. M. J. Barlow & R. H. Méndez, 55
- Zinn, J. C. 2021, arXiv e-prints, arXiv:2101.07252
- Zorec, J., & Briot, D. 1997, *A&A*, 318, 443
- Zucker, S., & Mazeh, T. 1994, *ApJ*, 420, 806