Title

Mapping the Kinematic Structure of the Milky Way Halo with the HALO7D Survey

Author

Cunningham, Emily Clifford

Publication Date

2019

UNIVERSITY OF CALIFORNIA
SANTA CRUZ

MAPPING THE KINEMATIC STRUCTURE OF THE MILKY WAY HALO WITH THE HALO7D SURVEY

A dissertation submitted in partial satisfaction of the
requirements for the degree of

Doctor of Philosophy

in

ASTRONOMY & ASTROPHYSICS

by

Emily Clifford Cunningham

June 2019

The Dissertation of Emily Clifford Cunningham
is approved:

Professor Constance Rockosi, Chair

Professor Puragra Guhathakurta

Alis Deason, Ph.D.

Lori Kletzer
Vice Provost and Dean of Graduate Studies

Table of Contents

List of Figures vi

List of Tables xvi

Abstract xix

Acknowledgments xxi

1 Introduction **1**
- 1.1 Outline . . . 3

2 Isotropic at the Break? 3D Kinematics of Milky Way Halo Stars in the Foreground of M31 **6**
- 2.1 Introduction . . . 6
- 2.2 Dataset . . . 11
  - 2.2.1 HST Imaging: Proper Motions . . . 11
  - 2.2.2 Keck/DEIMOS Spectra . . . 11
- 2.3 Velocity Ellipsoid Parameter Estimation . . . 21
- 2.4 Results . . . 23
- 2.5 Discussion . . . 27
  - 2.5.1 A Galactic Shell . . . 27
  - 2.5.2 Dual Stellar Halo: In-Situ Star Formation . . . 28
  - 2.5.3 Future Work . . . 29
- 2.6 Conclusions . . . 31

3 HALO7D I: The Line of Sight Velocities of Distant Main Sequence Stars in the Milky Way Halo **33**
- 3.1 Introduction . . . 33
  - 3.1.1 HALO7D: A Deep, Pencil Beam Complement to Gaia . . . 35
  - 3.1.2 Need for VELOCIRAPTOR: Challenges of Deep Slit Spectroscopy 37
- 3.2 Data: The HALO7D Keck Program . . . 40
  - 3.2.1 Survey Fields . . . 40
  - 3.2.2 Halo Star Candidate Selection . . . 41

3.2.3 Extragalactic Targets 45
3.2.4 Observations 45
3.2.5 Spectroscopically Confirmed Contaminants 49
3.3 Hierarchical Bayesian LOS Velocities: Velociraptor 53
3.3.1 Definitions 56
3.3.2 Single-Mode: Modeling a Single Spectrum 57
3.3.3 Hierarchical Modeling 60
3.4 Results 64
3.4.1 LOS Velocity Distributions 64
3.4.2 Comparison with Other Tracers 68
3.5 Comparison with Simulations 76
3.5.1 Dispersion Estimates 81
3.5.2 Accretion Histories 83
3.6 Conclusions 86

4 HALO7D II: The Halo Velocity Ellipsoid and Velocity Anisotropy with Distant Main Sequence Stars 90
4.1 Introduction 90
4.2 Dataset 98
4.2.1 Keck/DEIMOS Spectroscopy 98
4.2.2 HST Proper Motions 99
4.3 Modeling The Halo Velocity Ellipsoid 111
4.3.1 Disk Model 112
4.3.2 Halo Model 113
4.3.3 Absolute Magnitudes 114
4.3.4 Full Posterior 114
4.4 Results 117
4.4.1 Spherically Averaged Estimates 117
4.4.2 Individual Fields 120
4.4.3 Disk Contamination 121
4.4.4 The β Radial Profile 123
4.5 Comparison with Other Studies 126
4.6 Comparison with Simulations 130
4.7 Conclusions 134

5 Conclusions 136
5.1 Summary 136
5.2 Future Directions 138
5.2.1 Mapping Velocity Anisotropy with Gaia 138
5.2.2 The Milky Way Halo Main Sequence with *HST* Proper Motions 140
5.2.3 Anisotropy Variation in Cosmological Simulations 141

A Testing Velociraptor 144

B Bayesian Mixture Modeling for PM Measurements 151
B.1 Proper Motion Model 151
  B.1.1 Gaussian Mixture Models written with Indicator Variables . . . 155
  B.1.2 Gibbs Sampling Algorithm 156

C Fake Data Testing for Ellipsoid Model 159
C.1 Sensitivity to Sample Size 163

Bibliography 165

List of Figures

2.1 Visual representation of the evolution of our sample, from the 23 CMD-selected halo star candidates from D13 to our sample of 13 stars. The different symbols represent the classification of the stars based on their proper motions: red triangles are M31 star candidates, while blue squares are Milky Way stars (see Fig. 3 of D13). The pink circle denotes the object classified as a potential Milky Way disk star in D13. 8

2.2 Two sample stellar spectra from our sample: one of our higher signal-to-noise spectra (S/N per Å=26) is shown in purple and a lower signal-to-noise spectrum (S/N per Å=8) is shown in blue. The noise spectrum for the lower signal-to-noise spectrum is shown on the bottom in pink. The spectra have been normalized, smoothed with a Gaussian kernel with $\sigma = 3$, and plotted with a vertical offset. We mask the chip gap in each spectrum, as well as the telluric A band at 7600 Åand the telluric B band at 6875 Å. The dip in the spectra between 6700 Åand 7200 Åis an instrumental feature. Spikes in the noise spectrum correspond to night sky emission lines. The signal to noise ratios are computed at Hα. . . . 9

2.3 Cumulative histogram of LOS velocities (in the Galactocentric frame) of the 13 halo stars in our sample (black). The overplotted blue line shows the CDF for the most likely value for σ_{LOS} for the full sample (see Sec. 2.3), with the shaded blue region indicating the 68% confidence region. An approximate CDF for the Triangulum-Andromeda Stream (TriAnd) is shown in red ($v_0 \sim 50$ km s^{-1}, $\sigma \sim 15$ km s^{-1}). The pink line shows the CDF when the LOS velocity distribution is modelled as a double Gaussian, with TriAnd ($\sim$ 20%) and the field halo treated as separate components. 13

2.4 Projections of our 3D kinematic sample, color coded by LOS velocity in the GSR frame. Likely TriAnd members are indicated by triangles. The ellipses show the 2D projection of the 3D velocity ellipsoid; the mean sample distance (20 kpc; see Table 2.2, Section 2.3) was used to convert from km s^{-1} to mas yr^{-1} for the proper motion axes. Ellipses are drawn to enclose 2σ. 14

2.5 Projections of the posterior probability distribution for our four free parameters, when the full sample of 13 objects was used. Contours are shown at 0.5, 1, 1.5 and 2 σ, respectively. The top panel in each column shows the 1D marginalized PDF for each parameter, with peaks and 68 % confidence intervals indicated by dashed vertical lines. We acknowledge the use of triangle.py (Foreman-Mackey et al. 2014) to produce this figure. 20

2.6 Radial velocity anisotropy profile. The "break radius" of the Milky Way stellar halo is shown by the blue shaded region ($16 \lesssim r/\mathrm{kpc} \lesssim 26$; Deason et al. 2011). Our measurement of β, from the 3D kinematics of $N = 13$ stars in the radial range $18 \lesssim r/\mathrm{kpc} \lesssim 32$, is shown in pink. Solar neighborhood measurements, using full 3D velocity information, find a radially biased β, shown in red (Bond et al. 2010; Smith et al. 2009). The remaining error bars show estimates of β for distant ($D \gtrsim 10$ kpc) halo stars using LOS velocity distributions (Sirko et al. 2004; Kafle et al. 2012; Deason et al. 2012; King et al. 2015). 26

3.1 Footprints of the four HALO7D fields. Black points are the positions of halo star candidates selected from the CANDELS catalogs. Filled points were observed with DEIMOS, while empty circles denote halo star candidates that were not observed. Green points in EGS and COSMOS are halo star candidates selected from the IRAC and UltraVista catalogs, respectively. DEIMOS mask outlines are drawn in purple. Grey squares indicate one epoch of HST imaging that is used for measuring PMs of these same stars (see Paper II). 43

3.2 CMDs for stars in the Besançon Galaxy Model, in a 1 deg^2 field of view centered on the coordinates of EGS. The green density maps show the CMD locations of the disk stars, with the number of disk stars in each CMD bin indicated by the colorbar. Halo stars are shown in magenta; only one out of five halo stars are shown for clarity. The HALO7D selection boxes are shown in blue. Stars were assigned priority based on their positions in these two CMDs: stars were assigned top priority if they fell within both solid selection boxes, and lowest priority if they fell into only one of the dashed selection boxes. 46

3.3 CMDs for stars in the four HALO7D fields. Selection boxes are shown in blue. Stars were assigned priority based on their positions in these two CMDs: stars were assigned top priority if they fell within both solid selection boxes, and lowest priority if they fell into only one of the dashed selection boxes. Magnitudes in F606W and F814W are computed in the STMAG system. We note that the bright stars in GOODS-S have been masked out in the catalog used for selection. 47

3.4 Signal to noise per angstrom for all HALO7D masks, as a function of V-band magnitude. Grey dashed lines indicate the predicted signal to noise with 8 hours of exposure time. 50

3.5 Illustration of the velocity offset caused by slit miscentering. Blue and pink lines show spectra of the same star taken during different observing runs; lefthand panels show the Hα region of the spectrum, and righthand panels show the telluric A-band region. Top panels show the raw spectra, in the observed frame, uncorrected for heliocentric motion as well as slit miscentering; the Hα lines and telluric absorption lines are clearly misaligned. The lower panels show the spectra with heliocentric and A-band corrections applied. 55

3.6 Histograms of posterior samples for the corrected velocity of a HALO7D target ($m_{F606W} = 22.0$) from seven observations. The top panel shows the posterior samples for the velocities when the spectra are modeled independently: note that these are the PDFs for $v = v_{raw} - v_{Aband} - v_{helio}$. The bottom panel shows the PDFs for the individual observations once they have been combined into the hierarchical model. The final PDF for the corrected velocity (thick black line), incorporating all observations, thus folds in all information and uncertainty from all observations of a star. 63

3.7 Velocity uncertainties for the HALO7D sample as a function of m_{F606W} apparent magnitude. Velocity uncertainties are the 16th and 84th percentiles of the posterior distributions. 64

3.8 Cumulative LOS velocity histograms in the four HALO7D fields. Velocities are given with respect to the Galactic Standard of Rest (GSR). Black lines indicate the CDFs for the data: for each of the 100 black lines, velocity values were drawn from the posterior distributions for the measurements. The width of each step thus demonstrates the velocity uncertainty for that data point. The CDF for the disk model is shown in green, and the CDF for the halo model, using the posterior median values for $\langle v_{LOS} \rangle, \sigma_{LOS}$ is shown in pink. Blue lines show the CDF for the mixture of the disk and halo populations, for 100 draws from the joint posterior distribution for $f_{\mathrm{Disk}}, \langle v_{LOS} \rangle, \sigma_{LOS}$. 70

3.9 LOS velocity histograms in the four HALO7D fields. Shown in pink are the resulting velocity distributions from 50 draws from the posteriors for $\langle v_{LOS} \rangle$ and σ_{LOS}. The green line indicates the disk distribution. The parameters of the disk velocity distribution are fixed; only the fraction to the total contribution is allowed to vary. Bold lines show the corresponding distributions for the median posterior values of $f_{\rm disk}$, $\langle v_{LOS} \rangle$ and σ_{LOS}. 71

3.10 Posterior samples for the mean LOS velocity (right), LOS velocity dispersion (middle) and the fraction of disk contamination (right) for all four fields. All fields have mean LOS velocity consistent with 0 km/s, and dispersions consistent with one another. 72

3.11 Normalized PDFs for absolute magnitude for six different choices of $m_{F606W} - m_{F814W}$ color. These PDFs are derived from the KDE constructed from the VandenBerg et al. (2006) isochrones, weighted by a Salpeter IMF and the approximate age and metallicity distributions of the halo. 74

3.12 Cumulative distance distributions for the four HALO7D fields. Distance distributions are computed as given by Equation 3.13, using colors and assuming a MW stellar density profile (Deason et al. 2011). All fields have $\langle D \rangle \sim 20$ kpc. 74

3.13 LOS velocity dispersions of the four HALO7D fields, plotted as a function of mean Galactocentric radius. Vertical errorbars show the 16-84% quantiles of the marginalized posterior. We compare our LOS dispersions with results from other studies: the black dashed line indicates best-fit LOS dispersion profile from Xue et al. (2008), measured from BHBs in SDSS. The grey dashed line indicates the best-fit dispersion profile from Brown et al. (2010) study, using BHB and BS stars as tracers. The black connected points show the resulting dispersion profile from the SEGUE K-giant survey Xue et al. 2016. The HALO7D dispersions are consistent with predictions from other tracers. 75

3.14 LOS velocity distributions from one square degree "observations" of the BJ05 accreted stellar halos generated with Galaxia. Different colored histograms denote the observations in the different HALO7D fields. Seven out of the eleven BJ05 halos show consistent velocity distributions across the four fields. Three halos show consistency across three fields with one field dominated by substructure. Halo17 shows four distinct LOS distributions across the four fields. 77

3.15 Examples of posterior samples from our dispersion analysis on our mock datasets. Each panel shows histograms of posterior samples for the 1D marginalized posterior distributions for the velocity dispersion (same as second panel of Figure 3.10) for a single mock HALO7D dataset. Each colored histogram represents a different HALO7D field. The first two panels show examples of consistent estimates of the LOS velocity dispersion for Halo07 and Halo09. The third panel shows an example of posterior estimates for Halo17: here, the cold peaks in GOODS-N and GOODS-S have caused the estimates of σ_{LOS} to be lower in these fields than in the other two fields. The fourth panel shows results from a mock dataset analysis for Halo20, in which the kinematically cold substructure at 400 km s^{-1}) and at 0 km s^{-1} result in a higher dispersion estimate in the EGS field. 78

3.16 Age vs orbit circularity for the accretion events making up Halo07, Halo09, Halo17 and Halo20. Points are scaled by accretion event mass. Grey crosses indicate still-bound satellites. Colored points indicate the "dominant" satellites in the mock HALO7D samples. For Halo07 and Halo09, the same satellite is dominant across all four fields. In Halo17, one satellite dominates two fields while the other two fields are dominated by distinct satellites. For Halo20, three out of the four HALO7D fields are dominated by the most massive satellite, while the EGS field is dominated by a low mass, recent accretion event. The recent accretion events in Halo17 and Halo20 are responsible for the variable velocity distributions, while early, massive accretion results in consistent velocity distributions for Halo07 and Halo09. 78

3.17 Fraction of stars contributed to a given line of sight as a function of accretion time, for Halo07 (gold), Halo09 (magenta), Halo17 (purple) and Halo20 (light blue). Point sizes are scaled by the mass of the accretion event. Only events that contribute $> 10\%$ of the stars are shown. Different shape orientations denote different sightlines. Halo09's most massive satellite dominates the mock HALO7D sample along all sightlines, contributing 20-40% of the stars. The second most massive accreted satellite, which was accreted at the same time, also contributes 10-20% along all four sightlines. In Halo07, several relatively massive accretion events that were accreted around the same time all contribute between 10-20% of the stars in the four sightlines. These two accretion histories give rise to consistent velocity distributions along the different sightlines. In contrast, Halo17 has experienced recent accretion of low mass satellites. These recent accretion events dominate the four HALO7D sightlines, and create cold peaks in the LOS velocity distributions. In the case of Halo20, the most massive satellite contributes strongly to the samples across all four sightlines; however, as with Halo17, two lower mass, recent accretion events contribute $> 15\%$ of the stars along the EGS sightline, creating two cold peaks in the velocity distribution. 79

4.1 The multi-epoch *HST*/ACS footprints of the four HALO7D fields. Different colors indicate the positions of each ACS chip in the different *HST* programs used to measure PMs in this work. HALO7D spectroscopic targets are indicated by black points; filled points indicate targets for which we successfully measured a PM, whereas empty circles indicate targets for which we could not measure a PM. 97

4.2 Top panels: Proper motion diagrams of the four HALO7D fields. Black points indicate stars that were in the *HST* images that were not HALO7D spectroscopic targets — primarily MW disk stars. HALO7D spectroscopic targets are shown in pink. Light blue points indicate spectroscopically confirmed WDs, while red MW disk star contaminants are shown in orange. In the EGS PM diagram, the inset shows PMs and errorbars for six spectroscopically confirmed quasars observed in the EGS field. Lower panels: PM diagrams from the Besançon Galaxy Model, from 1 square degree fields centered on the coordinates of the HALO7D fields. Grey crosses indicate Besançon disk stars, whereas magenta crosses are Besançon halo stars. 100

4.3 Comparison of the *HST* measured PMs for sources brighter than $v \sim 21$ with their *Gaia* counterparts, if available. Top panels show the comparison of PMW $= -\mu^*_{\rm RA}$; lower panels show the comparison of PMN $= \mu_{Dec}$. The two datasets show excellent agreement, with *HST* PMs generally having lower errorbars. For clarity, figures are zoomed in to the regions of highest target density; there are several high |PM| sources that are beyond the range of the figures (2/27 stars in GOODS-N; 6 high |PMW| and 4 high |PMN| out of 39 matched sources in GOODS-S; and 7 high |PMW| and 4 high |PMN| out of 56 sources in EGS). In addition, the number of targets that overlap across the two datasets is not necessarily a function of the area of the field, because of *Gaia*'s scanning law; for example, there are only 6 sources with reported PMs in COSMOS, while EGS, GOODS-S and GOODS-N have 56, 39, and 27 matches, respectively. 105

4.4 CMDs of the four HALO7D fields, in the STMAG bands F606W and F814W. Gray lines indicate the HALO7D selection boxes for these bandpasses; see Paper I for more details on target selection. Stars used in this analysis (main sequence stars with successful LOS velocity and PM measurements) are plotted as colored circles, with the colors corresponding to the logarithm of the magnitude of their proper motions ($|PM|^2 = (\mu_l \cos b)^2 + \mu_b^2$). For the CMDs over the full color range of these fields, please see Figure 3 of Paper I. 106

4.5 Proper motion diagrams of the four HALO7D fields, in Galactic coordinates, color coded by LOS velocity in the Galactocentric frame. Solid black lines indicate the solar reflex motion along each sightline; the squares indicate the implied mean PM along each line of sight for $D = 5, 10, 20, 50$ kpc (with mean PM at larger distances approaching ($\mu_l \cos(b)$, $\mu_b) = (0, 0)$. Dotted lines indicate $\mu_l \cos(b) = 0$ mas yr^{-1}, $\mu_b = 0$ mas yr^{-1}. 107

4.6 Summary of posterior results for spherically averaged samples. Left-hand panels: histograms of marginalized posterior samples for the four parameters of the halo velocity ellipsoid. Each of the estimates shown combines targets from all four survey fields. The estimates using the full HALO7D sample are shown in purple, while the blue histograms show the resulting estimates from three apparent magnitude bins. Upper righthand panel: CDFs of the distances of the full sample and the three apparent magnitude bins. Lower righthand panel: posterior distributions for β. 116

4.7 Square roots of the second moments of the radial (circular points) and tangential (square points) velocity distributions, as a function of mean Galactocentric radius. Different colors indicate the results from the analysis of the full sample (purple) as well as the three apparent magnitude bins (same colors as in Figure 4.6). 117

4.8 Summary of posterior results for the four fields. Left-hand panels: histograms of marginalized posterior samples for the four parameters of the halo velocity ellipsoid. Each colored histogram represents a different HALO7D field. Upper righthand panel: CDFs for the distances of the stars in the four fields. Lower righthand panel: posterior distributions for β. 120

4.9 Posterior distributions for the disk contamination in the four HALO7D fields for each of the models. Black histograms indicate the posterior distributions for the fraction of disk stars when the fields are modeled individually. Colored histograms indicate the full sample (purple histograms) and the spherically averaged estimates in different apparent magnitude bins (as in Figures 4.6 and 4.7). 122

4.10 The Milky Way's radial anisotropy profile, β, as measured from 3D kinematics. Colored points indicate results from this work, while grey points indicate findings from previous work and other studies. The grey diamond shows the anisotropy estimate from Bond et al. (2010), using main sequence stars from SDSS, and the grey square shows the estimate from Cunningham et al. (2016), using 13 MW MSTO stars along the line of sight towards M31. Gray crosses show the estimates of β from Lancaster et al. 2019, using the overall estimate of β consisting of the mixture of a metal-rich, radially biased population and an isotropic, metal-poor population, from BHBs in Gaia DR2. The two recent estimates for β from MW globular clusters, using HST PMs and *Gaia* PMs, are shown with triangles (Sohn et al. 2018, Watkins et al. 2019). Square shaped points are results from individual fields, while our spherically averaged results are shown as circles. When using small fields to estimate anisotropy, β varies from mildly tangential (e.g., GOODS-S, M31) to strongly radial (e.g., EGS). However, the spherically averaged estimates are all consistently $\beta \sim 0.6$ (and consistent with solar neighborhood and GC estimates), and the posterior means increase as a function of mean sample distance. . . 125

4.11 Maps of velocity anisotropy of the stellar halo in two *Latte* FIRE-2 simulations of Milky Way-mass galaxies: m12i (left) and m12m (right). The top panels show stars in the distance range 35 kpc $< D <$ 50 kpc, while the lower panels show 50 kpc $< D <$ 70 kpc. Black points show the positions of star particles in Galactic coordinates. In each halo, the sky has been subdivided into patches, and the color of the large circle in each patch indicates the velocity anisotropy for that subset of stars. Within a given distance range, each halo shows variation in its velocity anisotropy across the sky. Variation as a function of distance is also evident. In addition, the median and spread in β both vary from halo to halo: many more of the star particles in m12m are on tangentially biased orbits than in m12i. 129

A.1 Results from testing Velociraptor on fake data. Top panel: errorbars show the mean recovered velocity, and the standard deviation of the recovered velocities, for 90 runs of Velociraptor in single-mode, as a function of signal to noise. Middle panel: resulting distributions of recovered velocities when Velociraptor is run with six observations in hierarchical mode. Note that here the x-axis refers to the signal to noise of a single observation. Lower panel: velocity error (computed as half the difference between the 84th and 16th percentiles) in single-mode (purple) and hierarchical mode with six observations (orange). 146

A.2 Full corner plot for all 11 parameters for a fake spectrum with S/N=10. The true parameter values used to generate the fake spectrum are shown in blue. The absorption line strength parameters for a given spectral region are covariant with the continuum level; this is expected based on how we have parameterized the absorption lines. 147

A.3 Traces for all 11 single-mode parameters, for a fake spectrum generated to have S/N=10. For clarity, we show traces for only 20 randomly selected walkers. Black dashed lines indicate the true values of the model parameters used to generate this fake spectrum. 148

A.4 Corner plot for the corrected velocity, the additional uncertainty σ_v^2, and the six raw velocities for six fake spectra that each have S/N=5. Note that we are only showing projections here for 8 out of the 68 parameters in this model. This particular run of Velociraptor ran for with 800 walkers for 500 steps. For this figure, we excluded the first 3000 steps as burn-in, and thinned the chain, showing every 50th sample. 149

A.5 Traces for the corrected velocity, the additional uncertainty σ_v^2, and the six raw velocities for six fake spectra that each have S/N=5. For clarity, we show traces for only 20 randomly selected walkers. Truths are shown as black dashed lines. Because of the complexity of the model and the large number of free parameters, the chains do not mix efficiently, and the sampler needs to be run for many iterations. Note that the true value for σ_v^2 is not recovered in this case, because all of our fake spectra were generated to have exactly the same velocity. 150

B.1 An example of the linear transformation method on two images, j8pu44cvq (taken in 2004) and jboa38c2q (from 2011). Axes represent the change in pixels, in the distortion-corrected frame (u, v), for objects in the two images, after applying the six parameter linear transformation. Grey crosses indicate the change in positions for the galaxies initially classified as "bad"; black points are the positions of "good" galaxies used in the reference frame. The change in positions of the "good" galaxies are clustered at $(0, 0)$. Pink stars show change in the positions of the stars in these two images. 154

C.1 Resulting projections of posterior samples for fake GOODS-N data. This fake sample contained 100 halo stars and 50 disk stars. The true values of the distributions used to generate the data are shown in blue. 162

C.2 Top panel: Distributions of posterior medians for the halo velocity ellipsoid parameters recovered from 30 fake datasets, each with 100 stars. Parameter values used to generate the fake data are shown as blue vertical dashed lines. Bottom panel: histograms of the the error estimates for each parameter. 164

C.3 Histograms of posterior medians for the estimates of β from fake data testing. Lefthand panel: distribution of β estimates from 30 fake datasets, each containing 100 stars, with PM uncertainties of 0.2 mas yr^{-1}. Righthand panel: the estimates of β from 100 fake datasets, each containing 20 stars, with PM errors of 0.20 mas yr^{-1}. For both sets of fake datasets, radial velocity uncertainties were assigned as a function of apparent magnitude (see Figure 7 of Paper I). None out of the 100 fake datasets yielded a posterior median estimate of $\beta < 0.14$; the minimum posterior median estimate was $\beta_{\rm min} = 0.20$. 164

List of Tables

2.1 The properties of the candidate halo stars with measured 3D kinematics used in this analysis. We give the right ascension (RA) and declination (DEC), *HST*/ACS STMAG magnitudes, PMs in Galactic coordinates and LOS velocity (in the Galactocentric frame). The RA, DEC and magnitudes come from Brown et al. (2009), and the proper motions derive from the study by Sohn et al. (2012). The LOS velocity measurements are described in Section 2.2.2. Potential TriAnd members are indicated by an asterisk. The signal to noise ratios are computed at Hα. 10

2.2 Summary of our main results. We give the velocity ellipsoid in Galactic and spherical coordinate systems and the resulting velocity anisotropy, both for when we include all 13 stars and for when we exclude the 2 stars that are likely TriAnd members. We also give the approximate location of our three *HST* fields in the plane of the sky, as well as the average heliocentric and Galactocentric distances for our sample (which are unchanged to within 0.5 kpc when TriAnd members are excluded). For the latter quantities we list two uncertainties, the first being the error in the mean, and the second being the root-mean-square spread of the sample. 19

3.1 Coordinates of the four CANDELS fields studied in HALO7D. These fields were chosen for their deep, multi-epoch *HST* photometry. The listed field area corresponds to the field area covered with multi-epoch imaging. Catalogs indicated with daggers were used in the secondary target selection; see Section 3.2.2. 42

3.2 Summary of the masks observed through HALO7D. In GOODS-N, the masks GN3 and GN4 contain the same MW targets, but different extragalactic targets. In GOODS-S, three sets of extragalactic targets were observed alongside each mask of MW targets. 51

3.3 Summary of the progression of the HALO7D sample, from CMD selected targets to objects used in kinematic analysis. In COSMOS and EGS, we first indicate halo star candidate selected from the CANDELS catalogs, followed by candidates from the secondary catalogs for the first three columns of the table (see Section 3.2.2). 52

3.4 Summary of model parameters (and priors) for different levels of our hierarchical model. In the first level of the model, we model a region of the spectrum containing an absorption line feature, such as Hα, the telluric A-band region, or the CaT region. In the next level, we use multiple spectral regions to estimate the corrected velocity of a star from a single spectroscopic observation. Finally, in our hierarchical model, we incorporate multiple spectroscopic observations into our estimate of the corrected velocity of the star. 54

3.5 Summary of results from the modeling of the LOS velocity distributions. Posterior medians are quoted, with error bars giving the 16/84 percentile credible regions. We list the predicted fraction of disk contamination from Besançon; because we have included no color or apparent magnitude information into our estimate of the disk contamination, we expect our f_{disk} estimates are higher than the actual disk contamination in our sample. We list the parameters for the disk model for each field: the skew α, the location ζ, and the scale ω. These values were derived by fitting a skew normal distribution to the velocities of disk stars that fall in the HALO7D selection box in the Besançon Galaxy Model. Finally, we list the average distance to each field, as computed in Equation 3.13. 69

4.1 Galactic coordinates, projection of the Sun's velocity in Galactic coordinates, and the median PM error (in Galactic coordinates) for the four HALO7D fields. Quoted median PM errors are the errors in a single component (e.g., $\mu_l \cos(b)$ or μ_b; we find both components of PM have the same median errorbars, to within 0.005 mas yr^{-1}, within a given field). 104

4.2 Summary of the *HST* programs used for the PM measurements in this paper. 110

4.3 Summary of the estimates of the parameters of the halo velocity ellipsoid, for the full sample, the apparent magnitude bins, and the individual HALO7D fields. Posterior medians are quoted, with errorbars giving the 16/84 percentiles. 118

B.1 Parameters of the conditional posterior distributions for the image transformation parameters. Conditional posteriors for all 6 image transformation parameters are normal distributions with mean μ and variance V. Sums are over all objects in an image. σ_j=0.02 for stars,σ_j=0.1 for "good" galaxies, and σ_j=3. for "bad" galaxies. 158

Abstract

Mapping the Kinematic Structure of the Milky Way Halo with the HALO7D Survey

by

Emily Clifford Cunningham

The structure and kinematics of the Milky Way (MW) stellar halo provide a unique archaeological record of the MW's formation, past evolution, and accretion history. We present studies of the kinematic structure of the MW stellar halo with the Halo Assembly in Lambda-CDM: Observations in 7 Dimensions (HALO7D) survey. The HALO7D dataset consists of Keck II/DEIMOS spectroscopy and *Hubble Space Telescope*-measured proper motions of Milky Way (MW) halo main-sequence turnoff stars in the CANDELS fields. We first present the HALO7D pilot study, in which we use 13 distant main-sequence MW halo stars make the first estimate of the velocity anisotropy β using 3D kinematic information outside of the solar neighborhood. Next, we present the spectroscopic component of the full HALO7D survey, and discuss target selection, observing strategy, and survey properties. We present a new method of measuring line-of-sight (LOS) velocities by combining multiple spectroscopic observations of a given star, utilizing Bayesian hierarchical modeling. We estimate the LOS velocity dispersions in the four fields and find that they are consistent with one another. We perform mock HALO7D surveys using the synthetic survey software Galaxia to "observe" the Bullock & Johnston (2005) accreted stellar halos. Based on these simulated datasets, the consistent LOS velocity distributions across the four HALO7D fields indicates that the HALO7D sample is dominated by stars from the same massive (or few relatively massive) accretion event(s). Finally, we present the proper motions for the HALO7D sample, and use the 3D kinematic measurements to estimate β. We find that β varies from field to field, which suggests that the halo is not phase mixed at $\langle r \rangle = 23$ kpc. We explore the β variation across the skies of two stellar halos from the *Latte* suite of FIRE-2 simulations, finding that both simulated galaxies show β variation over a similar range to the variation observed across the four HALO7D fields. The accretion histories of the two simulated galaxies result in different β variation patterns; spatially mapping β is thus a way forward in characterizing the accretion history of the Galaxy.

Acknowledgments

I've had a fantastic time as a graduate student at UCSC. It's been a time of intense personal and intellectual growth, and has also been a lot of fun. The fact that I've had such a positive experience is entirely thanks to the wonderful people who have educated me, worked with me, supported me, and laughed with me over these last six years.

First of all, I would like to thank my amazing advisors and mentors. I truly feel that I won the grad school lottery by entering UC Santa Cruz while Alis Deason was a Hubble Fellow. I could not have dreamt of a better mentor; in addition to being undeniably brilliant, she has always pushed me when I needed her to and been patient with me when I struggled. I've learned so much from her, and I'm so grateful to have had the opportunity to be her student. I am thrilled to continue to have her as a collaborator and friend. I also owe a tremendous debt of gratitude to my faculty mentors, Raja Guhathakurta and Connie Rockosi. As an entering grad student, I had little conception of what I wanted to pursue during my PhD, but had vague notions of wanting to observe at Keck Observatory. I had no idea I would be handed such a huge Keck program, and have the opportunity to learn from two such accomplished observers. I also have thank Raja for all he's taught me about being an educator and a mentor. I had so much fun teaching with him at Castilleja, and am so grateful for the opportunities he gave me to co-mentor students with him. Furthermore, I want to thank Connie and Raja not just for their scientific support but their personal support. From the moment I entered UCSC as a grad student, when I happened to be in the midst of a personal crisis, they have shown me only kindness and have gone above and beyond to support me. While they are all talented and successful scientists, Connie, Raja and Alis are also wonderful people whom I'm so glad to know.

I am also indebted to the amazing community of grad students at UC Santa Cruz that have supported me over the course of my PhD. I'm grateful for the amazing, lifelong friendships I've formed during my years in Santa Cruz, and was consistently blown away by the brilliance of my peers. In my class of only two students, I have to thank my fellow (and only) classmate Chris Mankovich, who was always supportive as we tackled the hurdles of grad school alongside one another. I would like to thank the amazing Claire Dorman, not only for training me on DEIMOS but also for being a wonderful friend whose kindness and thoughtfulness towards others is unparallaled. And she always makes me laugh. I'd like to thank Alexa Villaume for being such a great roommate at many conferences over the years and sharing so many adventures with me. She and Chris are also outstanding movie buddies. I'd like to thank Phil Macias, whose amazing physical intuition continues to inspire me and really helped me through the core classes during my first year of grad school. I can't thank Alexander Rudy enough for being such a generous person, both as a friend and as a colleague; I'm fairly certain that every grad student who overlapped at UCSC with Alex can say that he really helped them at some point in their career. I also can't thank him enough for introducing me to Esther Wallace, whom I love dearly and am honored to have as a friend. I'd like to thank Jenn Burt and Katie Hamren, who were my hosts when I visited UC Santa Cruz as a prospective student; they lured me here with promises of

Zumba, pedicures and cocktails, and they made good on all of those promises.

I of course have to thank my family, who have helped me build my intellectual confidence ever since I began to talk. From practicing my five times tables on his shoulders to giving me math puzzles in the car, my dad has always kindled and supported my intellectual curiosity without ever pressuring me in any direction, always making learning fun and enjoyable. I am so lucky to have my amazing mother as a role model; thanks to her, I've never had any doubt that it is possible to excel as a woman in a male-dominated field, that is possible to be a present, loving parent and also pursue a rewarding, challenging career. I couldn't have asked for more supportive parents, who taught me to think critically and to trust myself. I also have to thank my brother Sean, who has been my buddy for life. While we are so different in so many ways, there is no one with whom it is easier for me to hang out, and I'm really grateful to have him in my life. And, of course, I have to thank my non-human family members: Susie, Roxie and Hugo each offered me unconditional love and emotional support, in their own ways, as I worked my way through grad school.

Finally, I'd like to thank my family of Zach, Malcolm and Simon. I have had many emotional ups and downs throughout grad school, and Zach has always been there for me. Whether its helping me talk through my postdoc proposal over beers, working out hierarchical models on cocktail napkins, raising our ridiculous cats, helping me up literal mountains, or reassuring my most stressed, miserable self, I can always count on him. I can't thank him enough for his support, and can't emphasize enough how much he helped me get through this past particularly challenging year. I'm very lucky to have him on my team.

Paper Acknowledgments

The text of this dissertation includes reprints of several published papers led by Emily Cunningham.

The contents of Chapter 2 were published in the Astrophysical Journal, in Cunningham et al. 2016, ApJ, 820, 18, in collaboration with Alis Deason, Puragra Guhathakurta, Constance Rockosi, Roeland van der Marel, Elisa Toloba, Karoline Gilbert, S. Tony Sohn, and Claire Dorman. For this paper, I led the spectroscopic observations, data reduction, data analysis, and the preparation of all figures and the manuscript.

The contents of Chapter 3 were also published in the Astrophysical Journal, in Cunningham et al. 2019, 876, 124. This work was done in collaboration with Alis Deason, Puragra Guhathakurta, Constance Rockosi, Zachary Jennings, Evan Kirby, Elisa Toloba and Guillermo Barro. For this paper, I led the target selection, spectroscopic observations, data reduction, data analysis, the analysis of the simulations, and the preparation of all figures and the manuscript.

The contents of Chapter 4 are in press in the Astrophysical Journal at the time of thesis submission, and are published on the arXiv (Cunningham et al. 2018, arXiv:1810.12201). This work was done in collaboration with Alis Deason, Robyn Sanderson, Jay Anderson, S. Tony Sohn, Puragra Guhathakurta, Constance Rockosi, Roeland van der Marel, Sarah Loebman, and Andrew Wetzel. For this paper, I led the proper motion measurement effort as well as the development of the new measurement method, the 3D kinematic modeling, and the preparation of all figures and the manuscript.

I would like to thank all of my co-authors for their valuable contributions to these works, and for their permission to include these papers in this dissertation.

Acknowledgment of Financial Support

I would like to thank my sources of financial support over the course of my PhD. I was supported by a NSF Graduate Research Fellowship, and ARCS Foundation Fellowship, as well as NSF Grant AST-1616540. Partial support for this work was provided by NASA through grants for program AR-13272 from the Space Telescope Science Institute (STScI), which is operated by the Association of Universities for Research in Astronomy (AURA), Inc., under NASA contract NAS5-26555. I recognize and acknowledge the significant cultural role and reverence that the summit of Mauna Kea has always had within the indigenous Hawaiian community. We are most fortunate to have the opportunity to conduct observations from this mountain.

Chapter 1

Introduction

The Milky Way (MW) galaxy is a spiral galaxy that is embedded in halo of dark matter. According to the Lambda Cold Dark Matter (ΛCDM) paradigm, the MW assembled its dark matter halo over its evolution by consuming smaller dark matter halos. Some of these smaller dark matter halos contain smaller galaxies, known as dwarf galaxies. The remnants from these consumed galaxies are strewn throughout the MW's dark matter halo in a diffuse envelope known as the stellar halo. Because the orbital timescales of these halo stars are so long compared to the age of the Galaxy, the velocities of these stars retain a link to their initial conditions. Therefore, by mapping the positions and velocities of halo stars (i.e., the phase-space structure of the halo), we can learn about the MW's formation and evolution.

Our position within the MW offers provides us with the unique opportunity to study a galaxy star by star from within. However, the MW halo is very difficult to observe, because it is so diffuse and because of our position within the Galaxy. Until recently, our knowledge of the structure of the MW's stellar halo has been hindered by observational constraints. Our knowledge of the motions of halo stars has been limited to one component (the line-of-sight, LOS, velocity) of only bright, rare tracer populations (e.g., red giant branch [RGB] or blue horizontal branch [BHB] stars). For the first time, we finally have knowledge of the tangential motion of these stars, thanks to the second data release from the *Gaia* mission (*Gaia* DR2; Gaia Collaboration et al. 2018b). While evolved stars make excellent tracers due to their bright apparent magnitudes, there are disadvantages to working exclusively with giants. First of all, giants are rare: the vast majority of stars in the stellar halo are main sequence stars. Secondly, different types of evolved stars do not necessarily trace the same age and metallicity populations in the halo. For example, RR Lyrae are found more commonly in metal-poor populations, whereas M-giants are found in metal-rich populations (e.g., Price-Whelan et al. 2015). On the other hand, all stellar populations have MS stars. With its limiting magnitude of $G \sim 20$, *Gaia* will not provide proper motions (PMs) for MS stars farther than $D \sim 15$ kpc in the halo.

We designed the Halo Assembly in Lambda-CDM: Observations in 7 Dimensions (HALO7D) project to address these limitations. The HALO7D dataset consists of Keck II/DEIMOS spectroscopy and Hubble Space Telescope (*HST*) measured PMs measured of distant ($D \sim$ 10–100 kpc) MW main-sequence turnoff (MSTO) stars. *HST* is currently the only instrument capable of measuring PMs of distant ($D > 15$ kpc) halo MS stars, thanks to its stability, high spatial resolution, and well studied geometric distortion solutions and point-spread functions (PSFs; e.g., Anderson & King 2006).

Because of our faint limiting magnitude $v < 24.5$, this dataset is unique even in the era of *Gaia*.

1.1 Outline

This dissertation is organized as follows. In Chapter 2, we present the HALO7D pilot study, in which we present the first sample of stars to have measured 3D kinematics outside of the solar neighborhood. We use this sample of 13 stars, located along the line of sight towards M31, to estimate the halo velocity anisotropy β. We find β along this line of sight to be consistent with isotropy, lower than solar neighborhood β measurements by 2σ ($\beta_{SN} \sim 0.5 - 0.7$). The potential decrease in β with Galactocentric radius is inconsistent with theoretical predictions, though consistent with recent observational studies, and may indicate the presence of large, shell-type structure (or structures) at $r \sim 25$ kpc.

In Chapter 3, we present the spectroscopic component of HALO7D. We discuss target selection, observations, and survey properties. We also present a novel method for measuring the LOS velocities of stars from multiple spectroscopic observations using Bayesian hierarchical modeling. We present the LOS velocity distributions of the four HALO7D fields, and estimate their means and dispersions. All of the LOS distributions are dominated by the "hot halo": none of our fields are dominated by substructure that is kinematically cold in the LOS velocity component. Our estimates of the LOS velocity dispersions are consistent across the different fields, and these estimates are consistent with studies using other types of tracers. To complement our observations, we perform mock HALO7D surveys using the synthetic survey software Galaxia to "observe" the Bullock & Johnston (2005) accreted stellar halos. Based on these simulated datasets, the consistent LOS velocity distributions across the four HALO7D fields indicates that the HALO7D sample is dominated by stars from the same massive (or few relatively massive) accretion event(s).

In Chapter 4, we present the HALO7D 3D kinematic sample. We discuss our novel PM measurement methodology, developed to efficiently measure PMs from many *HST* pointings that are not necessarily well-aligned across epochs. This methodology makes use of a Bayesian mixture modeling approach for creating the stationary reference frame of distant galaxies. Using the 3D kinematic HALO7D sample, we estimate the parameters of the halo velocity ellipsoid, $\langle v_\phi \rangle, \sigma_r, \sigma_\phi, \sigma_\theta$, and the velocity anisotropy β. Using the full HALO7D sample, we find $\beta = 0.68^{+0.04}_{-0.05}$ at $\langle r \rangle = 23$ kpc. We also estimate the ellipsoid parameters for our sample split into three apparent magnitude bins; the posterior medians for these estimates of β are consistent with one another. Finally, we estimate β in each of the individual HALO7D fields. We find that the velocity anisotropy β can vary from field to field, which suggests that the halo is not phase mixed at $\langle r \rangle = 23$ kpc. We interpret this finding in the context of simulations by exploring the properties of β in two of the stellar halos from the *Latte* suite of FIRE-2 simulations. We find that the two simulated halos both show spatial variation (of similar magnitude to the observed variation) across their skies, and that the two halos have very different β maps as a result of their different accretion histories.

In Chapter 5, we summarize our findings, and discuss several future directions for studying the MW stellar halo with HALO7D, *Gaia*, and simulations.

Chapter 2

Isotropic at the Break? 3D Kinematics of Milky Way Halo Stars in the Foreground of M31

2.1 Introduction

The Milky Way halo devours hundreds of lower mass dwarf galaxies over its lifetime. The stripped stellar material from this voracious eating habit is splayed out in a vast, diffuse stellar halo. The orbital timescales at these large distances (< 10 kpc) are very long, and the halo stars retain a memory of their initial conditions. Thus, by studying the phase space distribution of halo stars, we are privy to a unique window into the past accretion history of our Galaxy.

Global kinematic properties, such as the velocity anisotropy (i.e., the relative pressure between tangential and radial velocity components), can provide important insight into the formation of the stellar halo (see Binney & Tremaine 2008). The exact merger and dissipation history of a spheroid can strongly affect its velocity anisotropy profile (e.g., Naab et al. 2006; Deason et al. 2013a). Local studies, limited to heliocentric distances $D \lesssim 10$ kpc, have measured the full 3D kinematics of halo stars. This has revealed a strongly radially biased velocity anisotropy with $\beta = 1 - \sigma^2_{\rm tan}/\sigma^2_{\rm rad} \approx 0.5 - 0.7$ (e.g., Smith et al. 2009; Bond et al. 2010), in seemingly good agreement with the predictions of simulations (e.g., Bullock & Johnston 2005; Cooper et al. 2010).

In Deason et al. (2013b) (hereafter D13), we exploited the long time-baselines and exquisite photometry of deep, multi-epoch *HST* fields to measure the *proper motions* (PMs) of main sequence turn-off (MSTO) stars in the distant Milky Way halo. Our pilot program used 5–7 year baseline *HST*/ACS fields towards M31 to measure PMs of $N \sim 13$ halo stars in the foreground. Our PMs are extremely accurate, with random errors of ~ 5 km s^{-1}. These 13 halo stars provided the first direct bound on the tangential velocity moments of the halo in this extreme radial regime, and provide new insights into halo structure. From the PMs measured for 13 Milky Way halo stars at $18 \lesssim r \lesssim 30$ kpc in our M31 *HST* fields, D13 inferred approximate isotropy between radial and tangential motions: $\beta = 0.0^{+0.2}_{-0.4}$. This differs by 3σ from local measures of the velocity anisotropy, which find strongly radial orbits. This trend of decreasing radial anisotropy with galactocentric distance conflicts with numerical simulations, which predict an outward increase in radial anisotropy.

In D13, we had no line-of-sight (LOS) information for these stars: we relied on

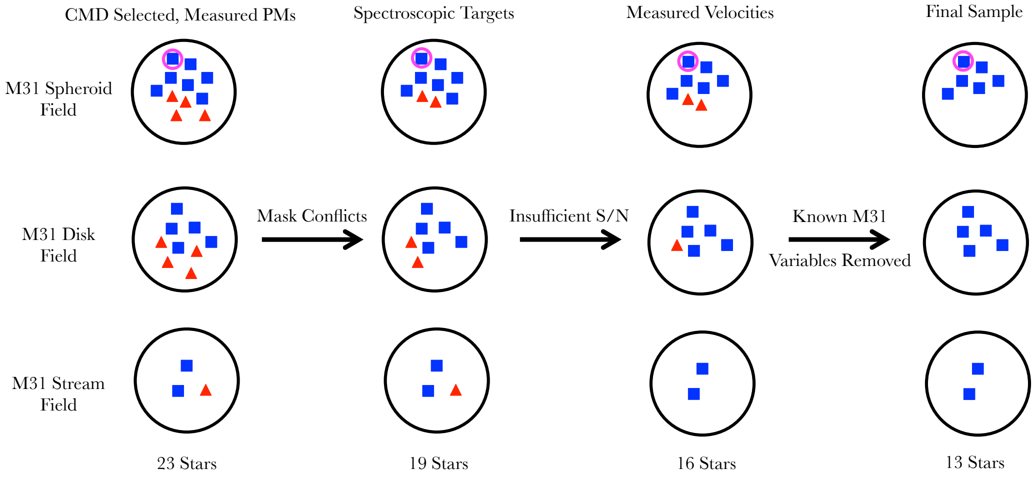

Figure 2.1: Visual representation of the evolution of our sample, from the 23 CMD-selected halo star candidates from D13 to our sample of 13 stars. The different symbols represent the classification of the stars based on their proper motions: red triangles are M31 star candidates, while blue squares are Milky Way stars (see Fig. 3 of D13). The pink circle denotes the object classified as a potential Milky Way disk star in D13.

the LOS velocities of other halo tracers (blue horizontal branch (BHB) stars, K giants) in different regions of the sky to form our argument. With spectroscopic information, we circumvent the need to rely on independent, and perhaps biased, tracers. In this paper, we present the LOS velocities for our halo star candidates, and use this 3D kinematic information to estimate the parameters of the velocity ellipsoid and the velocity anisotropy.

The paper is arranged as follows. In Section 2.2, we describe the target selection, proper motion measurements, spectroscopic observations and LOS velocity extraction. In Section 2.3, we describe our method for estimating the parameters of the velocity ellipsoid. Our results are presented in Section 2.4, and discussed in Section 2.5. We summarize our findings in Section 2.6.

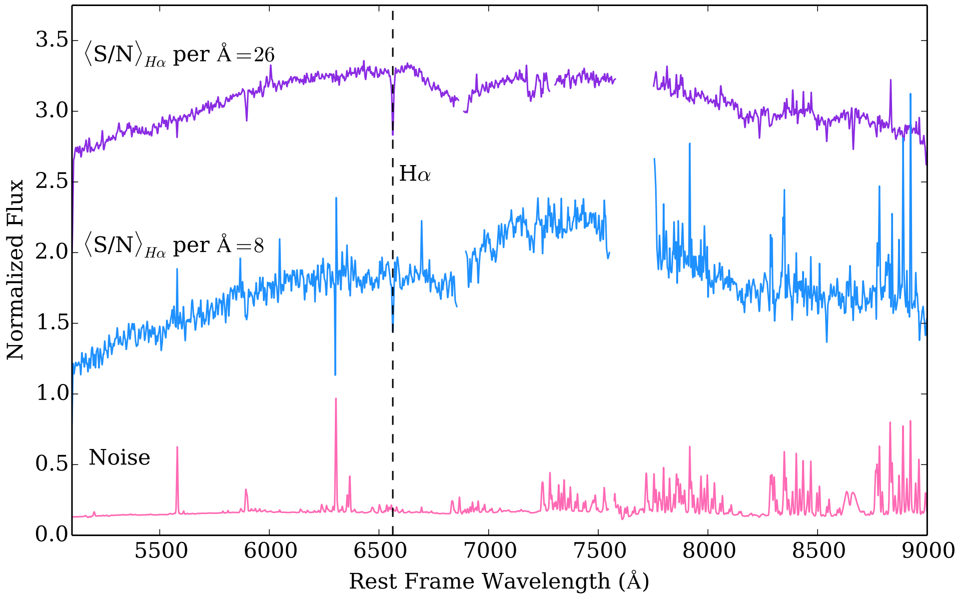

Figure 2.2: Two sample stellar spectra from our sample: one of our higher signal-to-noise spectra (S/N per Å=26) is shown in purple and a lower signal-to-noise spectrum (S/N per Å=8) is shown in blue. The noise spectrum for the lower signal-to-noise spectrum is shown on the bottom in pink. The spectra have been normalized, smoothed with a Gaussian kernel with $\sigma = 3$, and plotted with a vertical offset. We mask the chip gap in each spectrum, as well as the telluric A band at 7600 Åand the telluric B band at 6875 Å. The dip in the spectra between 6700 Åand 7200 Åis an instrumental feature. Spikes in the noise spectrum correspond to night sky emission lines. The signal to noise ratios are computed at Hα.

Field	RA (J2000)	DEC (J2000)	m_{F814W}	m_{F606W}	μ_l [mas yr^{-1}]	μ_b [mas yr^{-1}]	v_{LOS} (GSR) [km s^{-1}]	S/N per Å
M31 Spheroid								
	00:46:01.47	+40:41:35.53	21.86	21.45	-1.96 ± 0.04	-2.08 ± 0.04	$54 \pm 8^*$	19.4
	00:46:03.79	+40:41:22.81	22.53	22.19	1.36 ± 0.02	-1.33 ± 0.02	-68 ± 17	9.3
	00:46:03.67	+40:41:56.60	22.88	22.52	2.12 ± 0.03	-0.82 ± 0.02	-90 ± 19	9.4
	00:46:06.41	+40:42:15.07	22.53	22.06	1.45 ± 0.02	-0.90 ± 0.02	37 ± 16	10.5
	00:46:05.14	+40:43:37.19	21.82	21.47	3.91 ± 0.02	-1.59 ± 0.02	120 ± 6	20.9
	00:46:12.92	+40:41:22.51	22.92	22.61	1.88 ± 0.06	-2.83 ± 0.06	$43 \pm 15^*$	8.1
M31 Disk								
	00:49:08.91	+42:44:13.62	21.79	21.40	-0.59 ± 0.03	-1.50 ± 0.04	73 ± 4	30.0
	00:49:08.30	+42:44:50.44	22.12	21.66	$+1.03 \pm 0.04$	-0.78 ± 0.04	-42 ± 6	26.0
	00:49:13.50	+42:43:36.17	22.71	22.35	-0.71 ± 0.07	-0.67 ± 0.08	-117 ± 10	15.9
	00:49:13.38	+42:45:56.93	23.62	23.30	$+2.16 \pm 0.05$	-0.40 ± 0.06	142 ± 39	7.2
	00:49:13.69	+42:45:52.07	24.76	24.29	$+0.64 \pm 0.07$	$+0.58 \pm 0.06$	-175 ± 10	12.2
M31 Stream								
	00:44:26.44	+39:47:33.43	22.69	22.35	$+0.00 \pm 0.06$	-1.85 ± 0.06	-89 ± 7	18.3
	00:44:23.93	+39:46:26.25	23.83	23.46	-0.43 ± 0.05	-1.13 ± 0.07	16 ± 16	6.9

Table 2.1: The properties of the candidate halo stars with measured 3D kinematics used in this analysis. We give the right ascension (RA) and declination (DEC), *HST*/ACS STMAG magnitudes, PMs in Galactic coordinates and LOS velocity (in the Galactocentric frame). The RA, DEC and magnitudes come from Brown et al. (2009), and the proper motions derive from the study by Sohn et al. (2012). The LOS velocity measurements are described in Section 2.2.2. Potential TriAnd members are indicated by an asterisk. The signal to noise ratios are computed at Hα.

2.2 Dataset

2.2.1 HST Imaging: Proper Motions

A detailed description of the target selection can be found in D13, but we summarize the key points here. Our objects were selected from three *HST* observing programs: GO-9453, GO-10265 (PI: T.Brown), and GO-11684 (PI: R.P. van der Marel). The combination of these three programs provide deep, multi-epoch optical imaging of three fields in M31 (M31 Spheroid, M31 Disk and M31 Stream). These observations were used to measure the proper motion of M31 (Sohn et al. 2012), and during the course of this study, proper motion catalogs for *individual* stars in the three *HST* fields were created.

D13 selected Milky Way halo star candidates in color-magnitude space, using photometry from Brown et al. (2009): all stars fall within $m_{\rm F606W} - m_{\rm F814W} \sim -0.3$ and $21.5 \lesssim m_{\rm F814W} \lesssim 25.5$. In this region of the color-magnitude diagram (CMD) we expect minimal contamination from the Milky Way disk and M31's red giant branch (see Section 2.2 and Figure 1 of D13). Proper motions were then used to classify the objects as M31, Milky Way halo and Milky Way disk stars. The average uncertainty in the proper motion measurements is $\sigma_\mu \sim 0.05$ mas yr^{-1}.

2.2.2 Keck/DEIMOS Spectra

Spectroscopic Sample

Figure 2.1 demonstrates how our initial sample from D13 evolved into the sample used in this analysis. In D13, we presented proper motions for the 23 candidate halo stars selected from color-magnitude diagrams (CMDs): 11, 9, and 3 stars in the M31 Spheroid, M31 Disk and M31 Stream fields, respectively. Based on the proper motions, 13 of these stars were classified as Milky Way halo stars, 9 as M31 stars, and 1 as a potential Milky Way disk star (see Figure 3 of D13). The symbols in Figure 2.1 represent the proper motion classification: Milky Way halo star candidates are blue squares, M31 star candidates are red triangles, and the pink circle denotes the potential Milky Way disk star. We obtained spectra for 19 of the original 23 stars; we were not able to obtain spectra for all of the halo star candidates due to conflicts in the spectral direction on the DEIMOS slitmask. Three additional stars were too faint to measure velocities. After removing known variables in M31 (Brown et al. 2004; Jeffery et al. 2011), we were left with our final sample of 13 objects. It is worth noting that this is *not* the exact same sample of 13 stars used in the kinematic analysis of D13: one of the objects we used in D13 was very faint ($m_{\mathrm{F814W}} = 24.05$) and without strong spectral features, so we were unable to measure its velocity. We include the object classified as a potential disk star in D13 in our analysis (as its LOS velocity is consistent with halo kinematics).[1] The properties of our 13 stars are summarized in Table 2.1.

Observations

Observations were taken on September 28–30, 2014 on the Keck II telescope with the DEIMOS spectrograph (Faber et al. 2003). Over the course of the run, the seeing varied from $0.45'' - 0.9''$. We observed one slitmask in each of the three fields with the 600 line/mm grating. The central wavelength was 7200 Å, resulting in a

---

[1] As outlined in Section 2.2.2, we find that this star is likely a member of TriAnd.

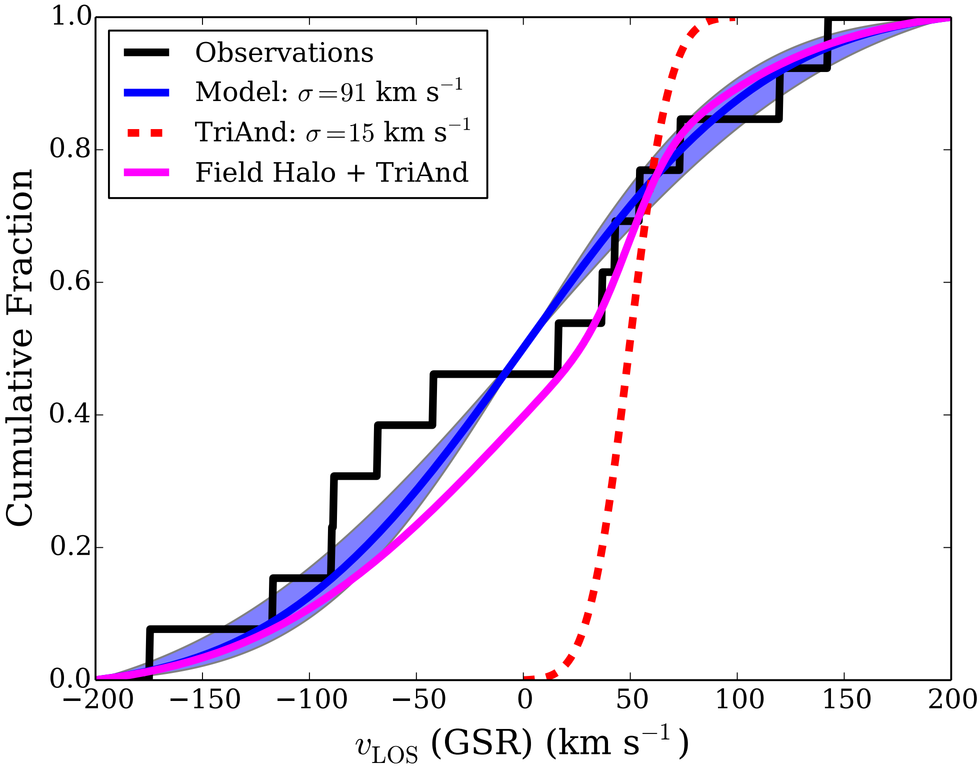

Figure 2.3: Cumulative histogram of LOS velocities (in the Galactocentric frame) of the 13 halo stars in our sample (black). The overplotted blue line shows the CDF for the most likely value for $\sigma_{\rm LOS}$ for the full sample (see Sec. 2.3), with the shaded blue region indicating the 68% confidence region. An approximate CDF for the Triangulum-Andromeda Stream (TriAnd) is shown in red ($v_0 \sim 50$ km s^{-1}, $\sigma \sim 15$ km s^{-1}). The pink line shows the CDF when the LOS velocity distribution is modelled as a double Gaussian, with TriAnd ($\sim 20\%$) and the field halo treated as separate components.

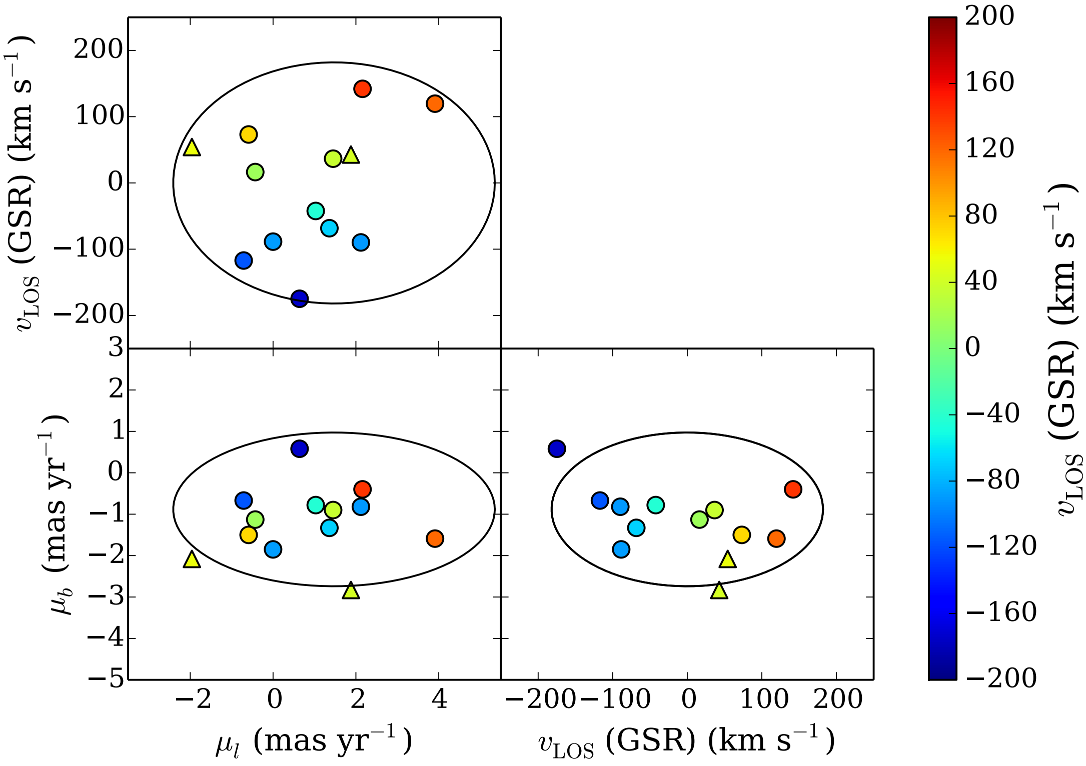

Figure 2.4: Projections of our 3D kinematic sample, color coded by LOS velocity in the GSR frame. Likely TriAnd members are indicated by triangles. The ellipses show the 2D projection of the 3D velocity ellipsoid; the mean sample distance (20 kpc; see Table 2.2, Section 2.3) was used to convert from km s^{-1} to mas yr^{-1} for the proper motion axes. Ellipses are drawn to enclose 2σ.

wavelength range of $\sim 4500-9300$ Å, where the exact wavelength range for each object varies depending on its position on the mask. The spectral resolution at $H\alpha$ (6563 Å) is $R \sim 2000$ (measured at the FWHM). In order to limit the flux losses due to atmospheric dispersion, we tilted our slits such that the position angle of the slit was consistent with the median parallactic angle of the observing block. The masks in the Spheroid and Disk fields were observed for a total of 5.9 hours, and the Stream field mask was observed for 5.3 hours. The slitmasks were then processed by a modified version of the *spec2d* pipeline developed by the DEEP2 team at UC Berkeley (Cooper et al. 2012). Two spectra from our sample are plotted in Figure 2.2; the top spectrum in Figure 2.2, shown in purple, has one of the higher signal-to-noise ratios of our sample (S/N per Å=26 at Hα), while the lower spectrum, shown in blue, is an example of one of our lower signal-to-noise objects (S/N per Å=8 at Hα). The noise spectrum from the lower signal-to-noise object is shown at the bottom of the figure in pink.

Velocity Measurements

Line-of-sight (LOS) velocities are measured from one-dimensional spectra using the Penalized Pixel-Fitting method (pPXF) of Cappellari & Emsellem (2004). The program determines the best fit composite stellar template for a given target using a penalized maximum likelihood approach. The 31 stellar templates employed in this analysis are described in detail in Toloba et al. (2016); the templates have high signal-to-noise ratios (100–800 Å^{-1}), and span a range of spectral types (from B1 to M8) and luminosity classes (from dwarfs to supergiants).

Errors in the raw velocity are determined through 1000 Monte Carlo simulations. In each simulation, we perturb the flux of the spectrum by adding noise to each pixel based on the uncertainty of the flux measurement in that pixel. The amount of noise added is drawn from a Gaussian distribution with width equal to the flux uncertainty. We then measure the velocity of each perturbed spectrum, and the error on the LOS velocity is taken to be the biweight standard deviation of the Gaussian distribution of velocities of perturbed spectra.

A-band telluric corrections are measured using the same method, and heliocentric LOS velocities are calculated by applying the A-band and heliocentric corrections to the raw velocities. The final uncertainty in the heliocentric LOS velocity is determined by adding in quadrature the errors on the raw velocity and the A-band correction.

Figure 2.3 shows a cumulative histogram of the LOS velocities for our sample of halo stars, in the frame of the Galactic Standard of Rest (GSR). Observed heliocentric velocities are converted to Galactocentric ones by assuming a circular speed of 240 km s^{-1} (e.g., Reid et al. 2009; McMillan 2011; Schönrich 2012) at the position of the sun ($R_0 = 8.5$ kpc) with a solar peculiar motion (U, V, W)=(11.1, 12.24, 7.25) km s^{-1} (Schönrich et al. 2010). Here, U is directed toward the Galactic center, V is positive in the direction of Galactic rotation and W is positive towards the North Galactic Pole.

In Figure 2.3, we see evidence for a “hot halo” population: there are no sharp increases where we expect to see contamination from the Milky Way Disk (along this line of sight, $\langle v_{\mathrm{disk}} \rangle \sim 145$ km s^{-1}) or M31 ($\langle v_{\mathrm{M31}} \rangle \sim -150$ km s^{-1}). The blue curve shows the cumulative distribution function (CDF) for the σ_{LOS} value with maximum posterior probability (see Section 2.3), with the shaded blue region indicating the 68%

confidence region. In contrast, as an example of substructure that is dynamically cold in LOS velocity, an approximate CDF for the Triangulum-Andromeda Stream (TriAnd; located along the line-of-sight towards M31) is shown in red ($v_0 \sim 50$ km s^{-1}, $\sigma \sim 15$ km s^{-1}; e.g., Deason et al. 2014; Sheffield et al. 2014). D13 suggested that the presence of a cold stream or TriAnd could be the reason for the relative increase in tangential pressure seen in this sample. However, our LOS velocity measurements confirm that this is not the case: the significant dispersion in the LOS velocity distribution demonstrates that our sample is not dominated by members of a cold stream nor by TriAnd.

While the LOS velocity distribution confirms that our sample isn't dominated by TriAnd, TriAnd members could still be biasing our measurement of the anisotropy. Given that our sample is in the same part of the sky and occupies the same region of CMD space as TriAnd (cf. Martin et al. 2014), we estimated the TriAnd contamination in our sample by fitting a double Gaussian to the LOS velocity distribution.[2] The resulting fit revealed that we expect 2-3 TriAnd stars in our sample, though the underlying hot halo LOS dispersion only changes by $\sim 5\%$ (see Table 2.2). The two stars that most likely belong to TriAnd based on their LOS velocities also happen to lie directly over the TriAnd overdensity as seen in CMDs (see Figure 1 of Martin et al. (2014)). The third star with the LOS velocity closest to that of TriAnd lies off the CMD overdensity. We therefore conclude that two of our stars are likely members of TriAnd. The CDF for the double Gaussian best-fit is shown in pink in Figure 2.3.

Figure 2.4 summarizes our 3 dimensional kinematic sample, showing the Galac-

---

[2] We computed the ratio of evidence (or Bayes factor) to compare the single and double Gaussian models, and found that neither model was strongly favored over the other.

tic proper motion components of the 13 halo stars color coded by LOS velocity. Our sample does not contain any members of M31, as all of these stars have proper motions too large to be associated with M31. As in Figure 2.3, Figure 2.4 shows no obvious clumpiness in any kinematic component, indicating that our sample is dominated by a “hot halo” population. However, it is intriguing that the two stars likely belonging to the TriAnd overdensity (shown as triangles in Figure 2.4) have relatively large proper motions. In the following sections, we consider the halo velocity ellipsoid both with and without the potential TriAnd stars.

Velocity Ellipsoid [km s^{-1}]				
Galactic coordinates				
Full Sample	$\langle v_{\rm LOS}^2\rangle^{1/2} = 91^{+27}_{-14}$	$\langle v_b^2\rangle^{1/2} = 88^{+28}_{-17}$	$\langle v_l^2\rangle^{1/2} = 138^{+43}_{-26}$	$\langle v_l\rangle = -67 \pm 37$
Excluding TriAnd	$\langle v_{\rm LOS}^2\rangle^{1/2} = 96^{+33}_{-15}$	$\langle v_b^2\rangle^{1/2} = 82^{+35}_{-16}$	$\langle v_l^2\rangle^{1/2} = 103^{+50}_{-17}$	$\langle v_l\rangle = -50^{+37}_{-40}$
Spherical polar coordinates				
Full Sample	$\langle v_r^2\rangle^{1/2} = 95^{+25}_{-14}$	$\langle v_\theta^2\rangle^{1/2} = 85^{+29}_{-17}$	$\langle v_\phi^2\rangle^{1/2} = 135^{+41}_{-20}$	$\langle v_\phi\rangle = 65 \pm 38$
Excluding TriAnd	$\langle v_r^2\rangle^{1/2} = 100^{+30}_{-15}$	$\langle v_\theta^2\rangle^{1/2} = 83^{+35}_{-15}$	$\langle v_\phi^2\rangle^{1/2} = 118^{+50}_{-21}$	$\langle v_\phi\rangle = 53 \pm 39$
Velocity Anisotropy				
Full Sample	$\beta = -0.3^{+0.4}_{-0.9}$	$\sqrt{\frac{\langle v_t^2\rangle}{\langle v_r^2\rangle}} = 1.6^{+0.5}_{-0.4}$	$\sqrt{\frac{\langle v_\phi^2\rangle}{\langle v_\theta^2\rangle}} = 1.4^{+0.6}_{-0.4}$	
Excluding TriAnd	$\beta = 0.1^{+0.4}_{-0.9}$	$\sqrt{\frac{\langle v_t^2\rangle}{\langle v_r^2\rangle}} = 1.4^{+0.6}_{-0.3}$	$\sqrt{\frac{\langle v_\phi^2\rangle}{\langle v_\theta^2\rangle}} = 1.3^{+0.6}_{-0.3}$	
Position				
	$l = 121^\circ$	$b = -21^\circ$	$\langle D\rangle = 20 \pm 1 \pm 7$ kpc	$\langle r\rangle = 25 \pm 1 \pm 7$ kpc

Table 2.2: Summary of our main results. We give the velocity ellipsoid in Galactic and spherical coordinate systems and the resulting velocity anisotropy, both for when we include all 13 stars and for when we exclude the 2 stars that are likely TriAnd members. We also give the approximate location of our three *HST* fields in the plane of the sky, as well as the average heliocentric and Galactocentric distances for our sample (which are unchanged to within 0.5 kpc when TriAnd members are excluded). For the latter quantities we list two uncertainties, the first being the error in the mean, and the second being the root-mean-square spread of the sample.

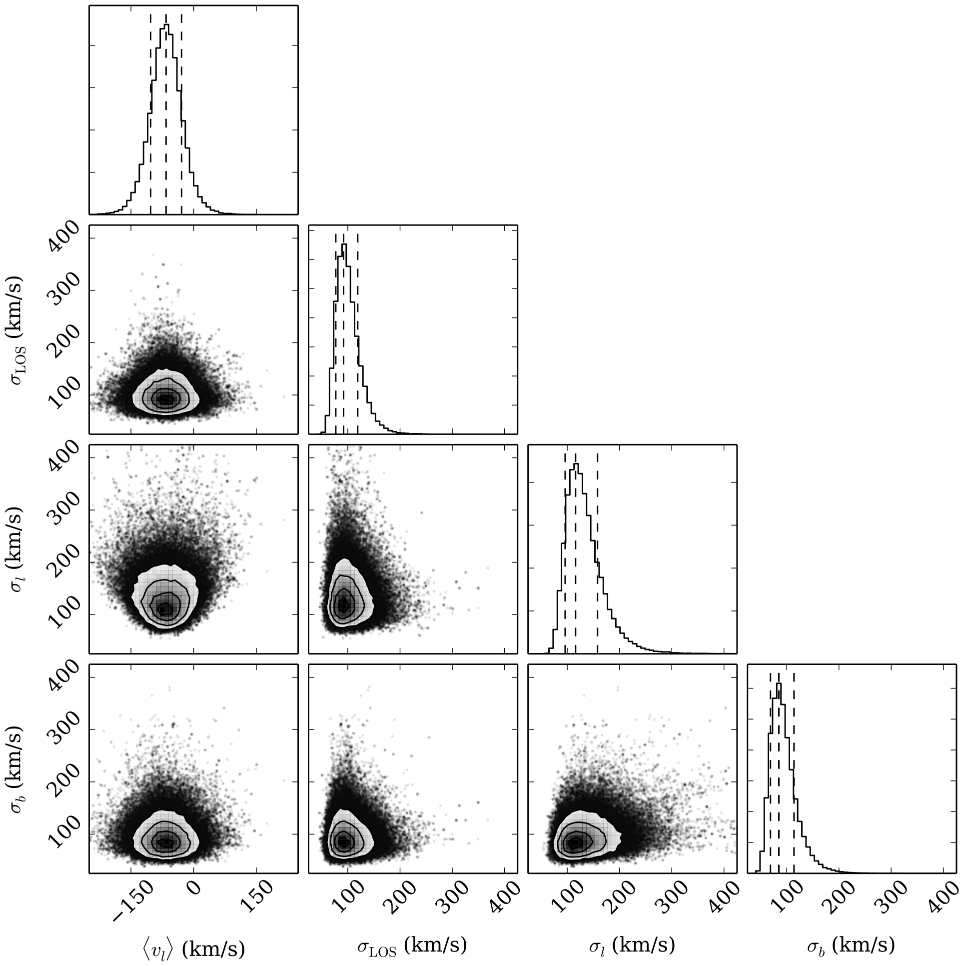

Figure 2.5: Projections of the posterior probability distribution for our four free parameters, when the full sample of 13 objects was used. Contours are shown at 0.5, 1, 1.5 and 2 σ, respectively. The top panel in each column shows the 1D marginalized PDF for each parameter, with peaks and 68 % confidence intervals indicated by dashed vertical lines. We acknowledge the use of triangle.py (Foreman-Mackey et al. 2014) to produce this figure.

2.3 Velocity Ellipsoid Parameter Estimation

We use a model of the halo probability distribution function (PDF) to estimate the parameters of the halo velocity ellipsoid ($\langle v_l \rangle$, $\langle v_b \rangle$, $\langle v_{\mathrm{LOS}} \rangle$, σ_l, σ_b, σ_{LOS}) from the observables ($m_{\mathrm{F814W}}, m_{\mathrm{F606W}} - m_{\mathrm{F814W}}, \mu_l, \mu_b, l, b, v_{\mathrm{LOS}}$). The method described is nearly identical to that in D13, though we have made modifications to incorporate the available LOS velocities. We summarize the key points here; see Section 3 of D13 for further details.

First, we determine the PDF for the heliocentric distance to each star. Continuous, double-Gaussian PDFs of absolute magnitude as a function of color were derived using IMF, metallicity, and age weighted VandenBerg et al. (2006) isochrones. We assume a Salpeter IMF, a Gaussian metallicity distribution with mean [Fe/H] $= -1.9$ and dispersion $\sigma = 0.5$ (e.g., Xue et al. 2008), and a Gaussian age distribution with mean $\langle T \rangle = 12$ Gyr and dispersion $\sigma = 2$ Gyr (e.g., Kalirai 2012). Possible systematics arising from these assumptions are explored in D13 (see Section 4.2). The resulting absolute magnitude PDF is given by:

$$\begin{aligned} G(M_{\mathrm{F814W}} | m_{\mathrm{F606W}} - \quad & m_{\mathrm{F814W}}) = G_1(A_1, M_1, \sigma_1, M_{\mathrm{F814W}}) \\ & + G_2(A_2, M_2, \sigma_2, M_{\mathrm{F814W}}), \end{aligned} \tag{2.1}$$

where $G(A, M, \sigma, x) = A \exp\left[-\left(x - M\right)^2 / (2\sigma^2)\right]$ and A, M and σ (amplitude, mean and sigma) are polynomial functions of $m_{\mathrm{F606W}} - m_{\mathrm{F814W}}$ color. See Section 3.1 and

Figures 5 and 6 in D13 for more detail. This absolute magnitude PDF is then translated into a distance PDF for each star in our sample using the distance modulus: $D = D(M_{\mathrm{F814W}}, m_{\mathrm{F814W}})$.

We then compute the velocity distribution function: $F_v = F_v(v_{\mathrm{LOS}}, D, \mu_l, \mu_b)$. We assume that the velocity distributions in both tangential and radial directions are Gaussian, with constant values of the ellipsoid parameters over the physical range spanned by our data. We convert observed heliocentric (v_l, v_b) velocities to the Galactocentric frame as outlined in Section 2.2.2. In the direction of M31, the velocity of the sun projects to: $(v_l, v_b) = (-139.5, 83.7)$. The 3-dimensional velocity probability distribution is given by:

$$\begin{aligned} F_v(v_l, v_b, v_{\mathrm{LOS}}) = & \frac{1}{(2\pi)^{3/2}\,\sigma_l\sigma_b\sigma_{\mathrm{LOS}}} \exp\left[-\frac{(v_l - \langle v_l\rangle)^2}{2\sigma_l^2}\right] \\ & \times \exp\left[-\frac{(v_b - \langle v_b\rangle)^2}{2\sigma_b^2}\right] \exp\left[-\frac{(v_{\mathrm{LOS}} - \langle v_{\mathrm{LOS}}\rangle)^2}{2\sigma_{\mathrm{LOS}}^2}\right] . \end{aligned} \tag{2.2}$$

The halo PDF at fixed $m_{\mathrm{F606W}} - m_{\mathrm{F814W}}$ color, in increments of absolute magnitude, apparent magnitude, Galactic PM, LOS velocity and solid angle (Ω), $F(y)$, where y is defined as $y = y(M_{\mathrm{F814W}}, m_{\mathrm{F814W}}, \mu_l, \mu_b, v_{\mathrm{LOS}}, \Omega)$, is given by:

$$F\,\Delta\mathbf{y} = F_v\,\rho\,D^5\,G\ \cos(b)\Delta\mathbf{y}. \tag{2.3}$$

Here, $\rho = \rho(D, l, b)$ is the density distribution of halo stars (we assume the broken power law profile derived by Deason et al. (2011)), $G = G(M_{\mathrm{F814W}}|m_{\mathrm{F606W}} -$

$m_{\rm F814W}$) is the absolute magnitude PDF in Eqn. 2.1 and

$$\Delta\mathbf{y} = \Delta M_{\rm F814W} \Delta m_{\rm F814W} \Delta\mu_l \Delta\mu_b \Delta v_{\rm LOS} \Delta\Omega \tag{2.4}$$

is the volume element.

We marginalize over absolute magnitude, and define the likelihood function:

$$L = \prod \bar{F}(\sigma_l, \sigma_b, \sigma_{\rm LOS}, v_{l,0}, v_{b,0}, v_{LOS,0}, \mathbf{x}), \tag{2.5}$$

where $\bar{F} = \int F \,{\rm d}M_{\rm F814W}$.

We sample the marginalized posterior probability distribution with `emcee` (Foreman-Mackey et al. 2013), a PYTHON implementation of the Goodman & Weare (2010) affine-invariant Markov chain Monte Carlo (MCMC) ensemble sampler. We set $\langle v_b \rangle = \langle v_{\rm LOS} \rangle = 0$, but allow for net motion in Galactic longitude, which approximates the net rotational velocity (v_ϕ) of the halo. We assume a flat prior on the mean velocity in galactic longitude $\langle v_l \rangle$ and a flat prior between 0 and 450 km s^{-1} on the dispersions. Projections of our posterior probability are shown in Figure 2.5.

2.4 Results

Figure 2.5 shows projections of the samples of the posterior, with marginalized one-dimensional PDFs for each parameter shown in the top panel of each column. We find the following values for the velocity ellipsoid parameters, with 68% confidence limits: $\langle v_l \rangle = -66^{+37}_{-37}$ km s^{-1}, $\sigma_{\rm LOS} = 91^{+27}_{-14}$ km s^{-1}, $\sigma_l = 117^{+42}_{-19}$ km s^{-1}, and

$\sigma_b = 88^{+28}_{-17}$ km s^{-1}. Here we have quoted the peaks of the 1D marginalized PDFs, and the limits enclose 68% of the points on either side of the peak.

We convert our velocity ellipsoid quantities to spherical polar coordinates using a Monte Carlo method. Our galactocentric polar coordinate system is defined such that the sun is located on the negative x axis, and the polar angle ϕ is the angle from the negative x axis to the positive y axis ($l = 90$ degrees), such that ϕ is positive in the direction of Galactic rotation. To make the conversion from v_l, v_b, $v_{\rm LOS}$ to v_r, v_θ, v_ϕ, we generate a random sample of $\sim 25,000$ stars drawn from the halo density distribution (Deason et al. 2011):

$$\rho(r_q) \propto \begin{cases} r_q^{-\alpha_{\rm in}} & r_q \leq r_b, \\ r_q^{-\alpha_{\rm out}} & r_q > r_b. \end{cases} \tag{2.6}$$

where $r_q = x^2 + y^2 + z^2 q^{-2}$, $q = 0.59$ is the halo flattening parameter, $r_b = 27$ kpc, $\alpha_{\rm in} = 2.3$, and $\alpha_{\rm out} = 4.6$. The stars are placed along the line-of-sight and have heliocentric distances ranging from 10 to 100 kpc. The stars are assigned a velocity distribution based on a random selection from our MCMC samples. Each star's velocity components v_r, v_θ, v_ϕ are calculated from the generated positions and v_l, v_b, $v_{\rm LOS}$ velocities. The second moments in spherical polar coordinates are computed from the resulting Galactocentric velocity distributions.

By repeating this process 10^5 times, we compute PDFs for the second moments for the galactocentric velocity ellipsoid parameters. The uncertainties on these parameters are computed in the same way as the heliocentric velocity ellipsoid parameters: the limits enclose 68% of the points on either side of the peak. Our results are summarized in Table 2.2. Using the PDFs for the galactocentric second moments, we compute the PDF for the anisotropy parameter (Binney & Tremaine 2008):

$$\beta = 1 - \frac{\langle v_\theta^2 \rangle + \langle v_\phi^2 \rangle}{2\langle v_r^2 \rangle}. \tag{2.7}$$

We find $\beta = -0.3^{+0.4}_{-0.9}$, where we again quote the peak of the PDF and limits that enclose 68% of the points on either side of the peak. If we repeat this analysis excluding the two likely TriAnd members, we find $\beta = 0.1^{+0.4}_{-0.9}$. Both of these values are consistent with the value found in D13 ($\beta_{D13} = 0.0^{+0.2}_{-0.4}$), though our new values have larger error bars because we measured the LOS velocity distribution directly. Our values for the ellipsoid parameters in this case are also quoted in Table 2.2.

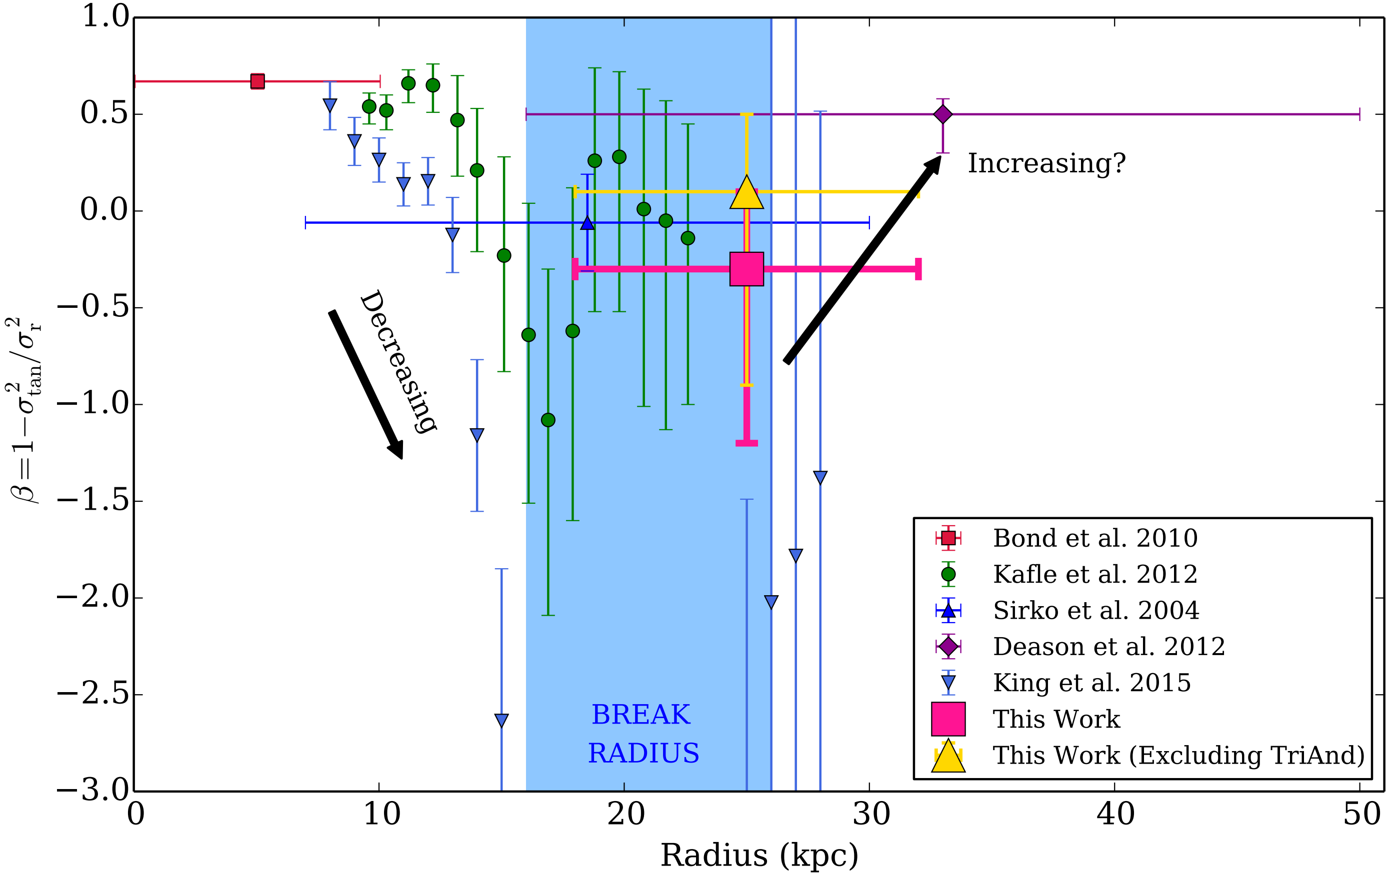

Figure 2.6: Radial velocity anisotropy profile. The “break radius” of the Milky Way stellar halo is shown by the blue shaded region ($16 \lesssim r/\mathrm{kpc} \lesssim 26$; Deason et al. 2011). Our measurement of β, from the 3D kinematics of $N = 13$ stars in the radial range $18 \lesssim r/\mathrm{kpc} \lesssim 32$, is shown in pink. Solar neighborhood measurements, using full 3D velocity information, find a radially biased β, shown in red (Bond et al. 2010; Smith et al. 2009). The remaining error bars show estimates of β for distant ($D \gtrsim 10$ kpc) halo stars using LOS velocity distributions (Sirko et al. 2004; Kafle et al. 2012; Deason et al. 2012; King et al. 2015).

2.5 Discussion

Our value of β is consistent with isotropy, and lower than local measurements by at least 1σ, which find a radially biased anisotropy ($\beta = 0.5 - 0.7$). The significant dispersion in the observed LOS velocity distribution (Figure 2.3) rules out the possibility that our sample is dominated by cold substructure.

Figure 2.6 shows the radial anisotropy profile of the Milky Way stellar halo. Our measurement is consistent with the observed "dip" in the anisotropy profile, seen in multiple studies that measured the velocity anisotropy of distant halo stars along different lines of sight using only LOS velocity distributions (Sirko et al. 2004; Kafle et al. 2012; Deason et al. 2012; King et al. 2015). This dip is also coincident with the observed break in the halo density profile around 16 kpc $\lesssim r \lesssim$ 26 kpc (Deason et al. 2011; Sesar et al. 2011; Watkins et al. 2009). In this section, we discuss some of the possible explanations of this result.

2.5.1 A Galactic Shell

In D13, we argued that the presence of global substructure, such as a shell (or multiple shells), is one explanation for both the steep fall-off in stellar density beyond the break radius and the decrease in anisotropy at that radius. Deason et al. (2013a) argued that a break in the Milky Way stellar density profile could be created by the build-up of stars at apocenter from either one relatively massive accretion event or several, synchronous accretion events. In this scenario, we would expect the stars to have an increase in tangential motion relative to radial motion at the turnaround radius, and thus a more isotropic β, just as we observe. This picture is consistent with what we find for likely TriAnd members: TriAnd is a large, cloud-like overdensity of stars likely at apocenter (Johnston et al. 2012), and including TriAnd in our sample makes β more tangentially biased. Chemical abundances for these stars may help to characterize the progenitor (or progenitors) of this shell (see Section 2.5.3).

Several of these cloud-like overdensities, such as TriAnd, the Virgo overdensity (VOD), the Hercules-Aquila overdensity (HerAq) and the Eridanus-Phoenix overdensity (EriPhe) are all located at approximately 20 kpc. Li et al. (2015) recently suggested that EriPhe, HerAq and the VOD could all be associated, and potentially fell in to the Milky Way as a group; TriAnd could also be a member of this group. A group infall event could explain the presence of all these overdensities at $\sim$ 20 kpc, the observed break in the density profile and the relative increase in tangential motion at this radius.

2.5.2 Dual Stellar Halo: In-Situ Star Formation

The break in the stellar density profile could also be an indication that the Milky Way has a "dual stellar halo," containing populations of different origins (Carollo et al. 2007; Carollo et al. 2010; Beers et al. 2012). Simulations predict that the stellar halo is composed both of accreted stars and stars that form *in-situ* (e.g., Zolotov et al. 2009; Font et al. 2011; McCarthy et al. 2012; Tissera et al. 2012; Cooper et al. 2015). In-situ stars have two flavors. The first are stars that form in the halo itself from gas accreted from the IGM or satellites; it remains unknown to what extent these populations and their properties are a result of the choice of hydrodynamics scheme. Secondly, stars can form in the disk of the Milky Way and then be kicked up into the halo due to merger events (these stars are sometimes called "heated disk stars"). In simulations, these stars can comprise a significant fraction of the stellar population (and sometimes even dominate) within $r \lesssim 30$ kpc. It's possible that our observed isotropy is a kinematic signature of a heated disk population. McCarthy et al. (2012) showed that these in-situ stars can have significant prograde rotation and therefore increased tangential pressure support from angular momentum, and we find a significant signal of prograde rotation ($\langle v_\phi \rangle \sim 70$ km s^{-1}). However, this scenario does not explain why measurements of the velocity anisotropy in the solar neighborhood find radially biased orbits, in the region of the halo where we would expect even more heated disk stars. Distinguishing between accreted and in-situ populations with kinematics alone remains challenging, and model predictions remain unclear. To better determine if our objects were accreted or formed in-situ, we need chemical abundances (see Section 2.5.3).

2.5.3 Future Work

Chemistry

Chemical information is key for disentangling the Milky Way's accretion history. Iron abundances of accreted populations are related to the masses of the dwarf progenitors (e.g., see Johnston et al. 2008; Kirby et al. 2013). If our 13 stars are accreted halo stars, measuring iron abundances may help to determine whether a single accretion event or several are responsible for the shell-type structure we observe, and we can use the abundances to estimate the mass(es) of the progenitor(s).

The chemical information in our stellar spectra is also our best hope of determining the relative contributions of different stellar halo formation mechanisms. Stars that form in the disk of the Milky Way in simulations are found to have a higher average [Fe/H] than accreted stars (Font et al. 2011; Tissera et al. 2012; Cooper et al. 2015). In addition, Zolotov et al. (2010) showed that in-situ stars are alpha-enriched relative to accreted stars at a given [Fe/H] at the high [Fe/H] end of the metallicity distribution function. These results are due to the fact that in-situ stars form in a deeper potential well than the accreted population. Several studies have used abundances in an effort to disentangle these populations locally (e.g. Nissen & Schuster (2010), with F and G main sequence stars within 335 pc; Sheffield et al. (2012) with M Giants out to 10 kpc). However, no such studies exist using main sequence stars outside the solar neighborhood. By measuring the iron and alpha abundances of distant main sequence halo stars, we can begin to assess the relative importance of different physical processes leading to the formation of the Milky Way's stellar halo.

HALO7D

In order to better understand the global halo properties, we need more than $N \sim 13$ stars! Through the HALO7D observing program (begun in Spring 2014), we will obtain deep (8–24 hour integrations) spectra of hundreds of distant MSTO halo stars *with measured HST proper motions* using Keck II/DEIMOS. We will target $N \sim 350$ stars in the four CANDELS fields (Grogin et al. 2011; Koekemoer et al. 2011): GOODS-N, GOODS-S, COSMOS, and EGS. All four of these fields are characterized by deep, multi-epoch *HST* imaging, and cover a total area of approximately 1000 square arcminutes. With this dataset, we will:

1. Measure LOS velocities of all stars, as well as [Fe/H] and [α/Fe] for those stars with sufficient signal to noise.
2. Measure the velocity anisotropy along four new lines of sight.
3. Measure the anisotropy as a function of galactocentric distance exclusively with stars that have 3D kinematic information.
4. Use chemical abundances to disentangle the Milky Way's accretion history and determine the relative contributions of stellar halo formation mechanisms.

HALO7D is an ongoing observational program with results forthcoming (Cunningham et al., in prep).

2.6 Conclusions

We present line-of-sight (LOS) velocities for $N = 13$ Milky Way halo stars with measured *HST* proper motions (PMs). Our sample is the first sample of halo stars with measured 3D kinematics outside of the solar neighborhood. The LOS velocities were measured from deep (5-6 hour) integrations on Keck II/DEIMOS. We combine the LOS velocity measurements with the proper motions to estimate the parameters of the velocity ellipsoid using an MCMC ensemble sampler. We find the velocity distribution in Galactic longitude l to have a mean $\langle v_l \rangle = -67^{+37}_{-37}$ km s^{-1} and a dispersion $\sigma_l = 117^{+42}_{-20}$ km s^{-1}. We find the dispersions in Galactic latitude b and the LOS to be $\sigma_b = 88^{+28}_{-17}$ km s^{-1} and $\sigma_{\mathrm{LOS}} = 91^{+27}_{-14}$ km s^{-1}, respectively.

Using our estimates of the ellipsoid parameters, we find $\beta = -0.3^{+0.4}_{-0.9}$, consistent with isotropy and with the result from D13, but lower than solar neighborhood measurements, which find a radially biased β, by at least 1σ. If we exclude likely TriAnd members from our sample, we find $\beta = 0.1^{+0.4}_{-0.9}$. These values are also consistent with other observational studies (using only LOS velocities) that have found a decrease in β around the observed break radius in the Milky Way density profile (16 kpc $\lesssim r \lesssim$ 26 kpc). These two findings in tandem suggest the presence of a shell-type structure in the halo at this radius, potentially formed by several destroyed dwarfs with similar apocenters. It is also possible that we are observing a population dominated by *in-situ* stars rather than an accreted population.

We need more observations and chemical information for distant halo stars to better understand the origin of the Milky Way stellar halo and its accretion history. We will achieve this with the HALO7D observing program, which will increase our sample of stars with 3D kinematics by a factor of ~ 30. The velocities and abundances measured from these observations will vastly improve our understanding of the Galaxy's accretion history and the origin of the stellar halo.

Chapter 3

HALO7D I: The Line of Sight Velocities of Distant Main Sequence Stars in the Milky Way Halo

3.1 Introduction

When a dwarf galaxy falls in to the Milky Way (MW) potential and is tidally disrupted, its stars become members of the MW stellar halo. The orbital timescales of these stars are long compared to the age of the Galaxy; thus, long after these debris have lost their spatial association, they remain linked by their kinematic (and chemical) properties. Six dimensional (6D) phase-space information and chemical abundances of halo stars can therefore be used to unravel the accretion events that have contributed to the mass assembly of the MW. With the Halo Assembly in Lambda-CDM: Observations in 7 Dimensions (HALO7D) survey, we are measuring 3D kinematic information and chemical abundances (as well as constraints on 3D positions) for distant main sequence (MS) MW halo stars.

The HALO7D dataset consists of Keck II/DEIMOS spectroscopy and Hubble Space Telescope (*HST*) measured PMs measured of distant ($D \sim$ 10–100 kpc) MW main sequence turnoff (MSTO) stars. In this paper, the first in the HALO7D series, we present the spectroscopic component of this dataset. In a companion paper (Cunningham et al. 2018b, in preparation; hereafter, Paper II), we present the PM dataset and analysis of the 3D kinematic sample.

This paper is organized as follows. In section 3.1.1, we motivate the HALO7D survey, and place our survey in context with other MW halo studies. In section 3.1.2, we introduce Velociraptor, our hierarchical Bayesian method for measuring the LOS velocities for our faint targets. In Section 3.2, we describe the HALO7D fields, target selection and observations. In Section 3.3, we present the details of the Velociraptor method. In Section 3.4, we present the LOS velocity distributions for the four HALO7D fields, estimate their velocity dispersions, and compare our results with those derived from other tracers. In Section 3.5, we compare our resulting LOS velocity distributions with predictions from simulations. We summarize our findings in Section 3.6.

3.1.1 HALO7D: A Deep, Pencil Beam Complement to Gaia

Our current picture of the stellar halo has largely been shaped by its giant population. Giants and evolved stars, such as red giant branch (RGB) stars, blue horizontal branch (BHB) stars, and RR Lyrae variables, have many advantages as halo tracers, particularly because of their bright absolute magnitudes. Giants have enabled the mapping of the stellar halo out to great distances: Slater et al. (2016) used K-giants to measure the density profile out to 80 kpc, Hernitschek et al. (2018) measured the density profile of the MW stellar halo out to 150 kpc with RR Lyrae from Pan-STARRS1, and Deason et al. (2018a) used BHBs in the Hyper Suprime-Cam survey to measure the density profile out to ∼200 kpc. These tracers have also revealed a wealth of substructure in the distant halo (see Sesar et al. 2017 and Conroy et al. 2018 as some recent examples).

Until recently, our kinematic knowledge of the stellar halo beyond $D \sim 10$ kpc has been limited to one component of motion (the line-of-sight (LOS) velocity) for these bright tracers. While progress has been made on measuring the LOS velocity dispersion profile (e.g., Xue et al. 2008 with SDSS BHBs; Cohen et al. 2017 with RR Lyrae), there has been little knowledge of the tangential motion of these stars until this year — from the second *Gaia* mission data release (Gaia Collaboration 2018). Increasing our knowledge of the tangential motions of stars, in order to better map our Galaxy and understand its formation and structure, is the primary science goal of the *Gaia* mission (Perryman et al. 2001).

Giants and evolved stars represent the upper echelon of a stellar population.

While they are bright, they are very rare: MS stars are the dominant population in every stellar population. In addition, it is difficult (perhaps impossible) to uniformly select giants across all age and metallicity populations in the halo. For example, RR Lyrae only can evolve in metal poor populations, while M giants are only found in metal rich populations (see Price-Whelan et al. 2015 for a discussion on how the relative numbers of RR Lyrae and M Giants in a population can be used to constrain the metallicity of a progenitor). While MS stars are fainter than giants, they are also more numerous, and all populations, regardless of age or metallicity, contain MS stars.

While MS stars are ideal tracers thanks to their presence in all stellar populations, they are challenging to observe, because they are faint. Because of its limiting magnitude of $G \sim 20$, beyond $D \sim 15$ kpc in the halo, *Gaia* will not provide proper motions for distant halo MS stars. The only instrument presently capable of measuring the proper motions (PMs) of distant ($D > 20$ kpc) MW MS stars is *HST*. *HST* is a powerful instrument for precision astrometry, due to its stability, high spatial resolution and well-studied geometric distortions and point spread functions (PSFs) (e.g., Anderson & King 2006). Multi-epoch *HST* imaging has been exploited to make extremely accurate PM measurements of resolved stellar systems in the Local Group (LG), including the Magellanic Clouds (Kallivayalil et al. 2006a, 2006b, 2013), MW globular clusters (Sohn et al. 2018), MW dwarfs Draco and Sculptor (Sohn et al. 2017), and M31 (Sohn et al. 2012).

The first individual MW stars with measured *HST* PMs were published by Deason et al. (2013b) (hereafter D13). These faint stars ($21 < m_{F606W} < 24.5$) had their PMs measured serendipitously during the Sohn et al. (2012) M31 PM study. The third component of the motion for these stars, the LOS velocity, was measured by Cunningham et al. (2016), using the DEIMOS spectrograph on the Keck II telescope, making this sample of 13 stars the first sample of stars with measured 3D kinematics outside the solar neighborhood. D13 and C16 confirmed that we can measure kinematic properties of distant MS stars with these two world class telescopes. However, these studies were limited to only 13 stars across three *HST* pointings. More stars and lines of sight through the halo are required to use MS star kinematics to investigate the formation of the Galaxy.

The HALO7D survey aims to address the current lack of distant MS stars with measured 3D kinematics. This dataset is unique even in the era of *Gaia*, measuring 6D phase space information for MS stars as faint as $m_{F606W} \sim 24.5$. In order to obtain spectra of these stars with sufficient signal-to-noise (S/N) for LOS velocity and abundance measurements, deep spectroscopy with a large telescope is required. HALO7D complements the *HST* proper motions with deep spectra (8–24 hour integrations) observed with Keck II/DEIMOS. However, spectra of individual stars at this depth is unprecedented, and new techniques are required to make measurements with these data. For the interested reader, we motivate this new technique in the next subsection; for those interested in the survey details, please skip ahead to Section 3.2.

3.1.2 Need for Velociraptor: Challenges of Deep Slit Spectroscopy

HST can measure proper motions for exceedingly faint stars. The deep spectroscopy required to observe to these same magnitudes presents a significant challenge!

To achieve the depth of our survey, targets were observed over multiple nights, and, in some cases, over years. In order to combine our different spectroscopic observations of a given target into a single measurement of the star's velocity, we required a new approach that took into account the fact that different observations of the same star will have different velocities. The VELOCIRAPTOR software employs Bayesian hierarchical modeling in order to combine multiple, often noisy, observations of a star, each with different zero-point offsets, into a single posterior probability distribution for the star's velocity.

One origin of zero-point offset is slit miscentering. Because stars are point sources, they do not fill the full width of the slit during observations (thanks to the exquisite seeing on Mauna Kea). If the star is not perfectly centered in the slit, the wavelength solution for the object is slightly offset from the wavelength solution given by the calibration of arc lamps (e.g., Sohn et al. 2007; Simon & Geha 2007). This wavelength solution difference corresponds to an apparent velocity shift that can be measured from the velocity of the telluric A-band absorption feature. Velocities of telluric features should be 0 km s^{-1} if the wavelength solution is correct. We refer to this velocity offset as the *A-band correction* (v_{Aband})[3], and it is subtracted from the raw velocity (v_{raw}, the velocity of stellar absorption features in an *observed* spectrum), along with the heliocentric correction (v_{helio}, due to the Earth's motion around the Sun) to

---

[3] While it has not been treated as such in the literature, we note that the effect of this slit miscentering is closer to a wavelength shift than a velocity shift. In this work, we measure stellar velocities using both the Hα and Ca triplet absorption features. Because these absorption features are approximately equidistant in wavelength from the telluric A-band feature, we can safely treat the A-band correction as a velocity offset, as others have done in the past. However, if one were to measure velocities using *only* Hα or Ca, or if one were to also use lines farther in the blue or red, then it is better to treat the A-band correction as a wavelength shift.

yield the corrected velocity in the heliocentric frame:

$$v = v_{\text{raw}} - v_{\text{Aband}} - v_{\text{helio}}. \tag{3.1}$$

The offset in the slit can be due to astrometric errors or slight mask misalignment. As such, the A-band correction varies from object to object on a given mask, and varies from observation to observation of the same object. For a star observed through a 1 arcsecond slit with Keck II/DEIMOS, configured with the 600ZD grating, v_{Aband} can be up to $\approx \pm 60$ km s^{-1}. Given that the velocity dispersion of the halo is on the order of 100 km s^{-1}, and that velocity dispersions of streams and dwarfs can be less than 10 km s^{-1}, it is essential to take into account this velocity offset before combining spectra from different observations. In addition, when the spectra in question are noisy, the measurements of v_{Aband} are also noisy, and their uncertainties need to be incorporated in to the ultimate measure of the corrected velocity uncertainty. In order to best leverage our signal to get accurate velocity measurements of our stars, along with correct uncertainties, we employ Bayesian hierarchical modeling to measure the velocities of individual spectroscopic measurements and their differing zero-point offsets *simultaneously.*

For the details of the model implemented by VELOCIRAPTOR, we refer the reader to Section 3.3. Fake data testing and further details are discussed in the Appendix.

3.2 Data: The HALO7D Keck Program

In this section, we describe the properties of HALO7D Keck II/DEIMOS spectroscopic program. We describe our choice of survey fields in Section 3.2.1; our target selection procedure is outlined in Section 3.2.2; and observations are described in Section 3.2.4. The extragalactic "piggy-back" programs are briefly described in 3.2.3. In Section 3.2.5, we describe how we selected the halo star candidates used for dynamical modeling from the spectroscopically observed targets.

3.2.1 Survey Fields

For a 3D kinematic study of distant halo stars, we aimed to survey high latitude fields that were characterized by many, contiguous, multi-epoch *HST* pointings. Deep, multi-epoch imaging is required in order to measure proper motions of distant main sequence stars (D13), while the large field of view of DEIMOS enables efficient spectroscopic follow-up of many contiguous *HST* pointings. The fields targeted by the Cosmic Assembly Near-infrared Deep Extragalactic Legacy Survey (CANDELS; Grogin et al. 2011; Koekemoer et al. 2011; PIs: S. Faber, H. Ferguson) were therefore a natural choice. The *HST* footprints of the CANDELS fields were designed with spectroscopic follow-up in mind, and they have been observed by *HST* many times over the course of *HST*'s operation.

HALO7D surveyed four out of the five CANDELS fields: EGS, COSMOS, GOODS-N, and GOODS-S (the fifth field, UDS, only has one epoch of *HST* imaging). The coordinates of the four HALO7D fields are listed in Table 3.1; their relatively high latitudes, resulting in minimal foreground contamination from MW disk stars, makes them ideal for both extragalactic and MW halo studies. Figure 3.1 shows the footprints of the four HALO7D fields. Tiling patterns for one epoch of HST imaging are shown in grey; the HALO7D Keck/DEIMOS mask pointings are shown in purple (see Section 2.3).

3.2.2 Halo Star Candidate Selection

Halo star candidates were selected using u, F606W (broad V filter) and F814W (broad I filter) photometry, from the catalogs listed in Table 3.1. Star candidates were identified based on image morphology (using the SExtractor parameter `class_star`; Bertin & Arnouts 1996) measured in WFC3 F160W images (the images used for source detection for CANDELS; see catalog references in Table 3.1 and references therein for more details on source detection and photometry). In order to select as many stars as possible, we used the fairly generous stellarity cut of `class_star`>0.5 (a more typical stellarity threshold for a study interested in including as many galaxies as possible would require that stars have `class_star`>0.98). We also excluded all targets which have non-zero measured redshifts.

To determine our selection boxes for optimally selecting halo stars, we used the Besançon Galaxy Model (Robin et al. 2003). Figure 3.2 shows color magnitude diagrams (CMDs) generated from the Besançon model from a 1-square degree field centered on the coordinates of the EGS field. Green points are disk members and halo stars are shown in magenta. In order to target as many halo star candidates as possible with minimal disk contamination, we targeted faint, blue stars. Figure 3.2 shows the

Field	R.A. (J2000)	Dec (J2000)	l (deg)	b (deg)	Area (arcmin^2)	Catalog Refs.
COSMOS	10:00:28	+02:12:21	236.8	42.1	288	Nayyeri et al. (2017), Muzzin et al. (2013)†
GOODS-N	12:36:44	+62:14:24	125.9	54.8	166	Barro et al. (2018; in prep)
GOODS-S	03:32:30	−27:48:11	223.6	-54.4	160	Guo et al. (2013)
EGS	14:15:29	+52:08:19	96.4	60.4	384	Stefanon et al. (2017), Barro et al. (2011)†

Table 3.1: Coordinates of the four CANDELS fields studied in HALO7D. These fields were chosen for their deep, multi-epoch *HST* photometry. The listed field area corresponds to the field area covered with multi-epoch imaging. Catalogs indicated with daggers were used in the secondary target selection; see Section 3.2.2.

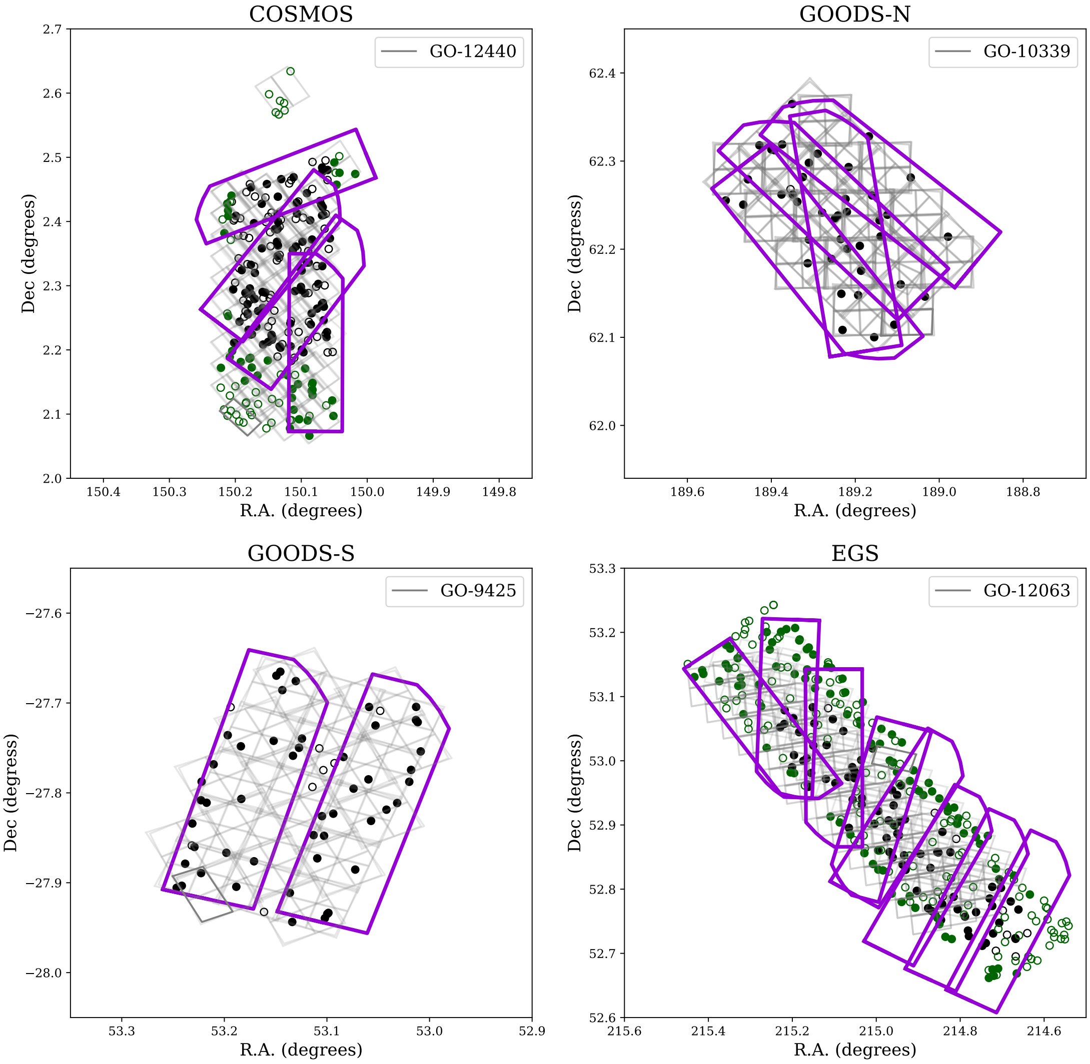

Figure 3.1: Footprints of the four HALO7D fields. Black points are the positions of halo star candidates selected from the CANDELS catalogs. Filled points were observed with DEIMOS, while empty circles denote halo star candidates that were not observed. Green points in EGS and COSMOS are halo star candidates selected from the IRAC and UltraVista catalogs, respectively. DEIMOS mask outlines are drawn in purple. Grey squares indicate one epoch of HST imaging that is used for measuring PMs of these same stars (see Paper II).

HALO7D selection boxes in blue; our highest priority selection boxes are shown with solid lines, and the dashed line indicates our lower priority selection box.

Figure 3.3 shows the CMDs for the four HALO7D fields, and our selection boxes in blue. Magnitudes in the F606W and F814W bands are in the STMAG system. The *u* band photometry is from ground based imaging, and magnitudes are in the AB system. In COSMOS and EGS, we used CFHT *u* band photometry; in GOODS-S, we used CTIO *U* band photometry; and in GOODS-N, we used KPNO *U* band photometry (see references in Table 1).

Targets were assigned a priority for selection on a scale from 1-4 (with 4 being highest priority) based on the selection boxes:

- Priority 4: Target falls in both solid selection boxes.
- Priority 3: Target falls in one of the solid boxes.
- Priority 2: Target falls in both dashed boxes.
- Priority 1: Target falls in one of the dashed boxes.

Additional Target Selection

In the COSMOS and EGS fields, the CANDELS catalogs (developed from WFC3) did not overlap the full area with multi-epoch ACS imaging. To select targets in these regions (which contain stars that can have measured PMs), we used additional catalogs. In EGS, we used the ACS F606W/F814W fluxes published in the Barro et al. (2011) photometric catalog, and used the same prioritization scheme as described above. Sources in this catalog were identified using IRAC 3.6+4.5 μm imaging. Stars were identified by combining eight stellarity criteria based on photometric and morphological properties; see section 3.1 of Barro et al. (2011) for more detail. We included all targets with a total sum of stellarity criteria greater than 2, meaning that it was classified as star-like by at least two of the eight stellarity criteria (greater than 3 would be typical, but we again made our selection generous in the interest of not excluding stars).

In COSMOS, we selected targets from the K_s selected catalog of the COSMOS/UltraVISTA field from Muzzin et al. (2013). Stars are classified in this catalog using $u^* - J, J - K_s$ colors; see section 3.3 and Figure 3 in Muzzin et al. (2013). To select HALO7D targets, we used the CFHT u band and the Subaru V band fluxes (Capak et al. 2007), using the same selection box as in the top panels of Figure 3.3.

3.2.3 Extragalactic Targets

Because the stellar halo is so diffuse, we typically placed only ~ 25 halo star candidates on a given DEIMOS mask. DEIMOS masks can contain up to ~ 150 slits: this provided an opportunity to obtain deep spectra for extragalactic targets as well as Galactic targets. These data have been used to study galactic winds in $z \sim 1$ (Yesuf et al. 2017, Yesuf et al., in prep); quiescent galaxies at $z \sim 0.7$ (Conroy et al., in prep); internal galaxy kinematics (Wang et al., in prep); and dwarf galaxies (Guo et al., in prep).

3.2.4 Observations

Spectra were obtained on the Keck II telescope with the DEIMOS spectrograph (Faber et al. 2003). Observations took place over the course of three years, beginning

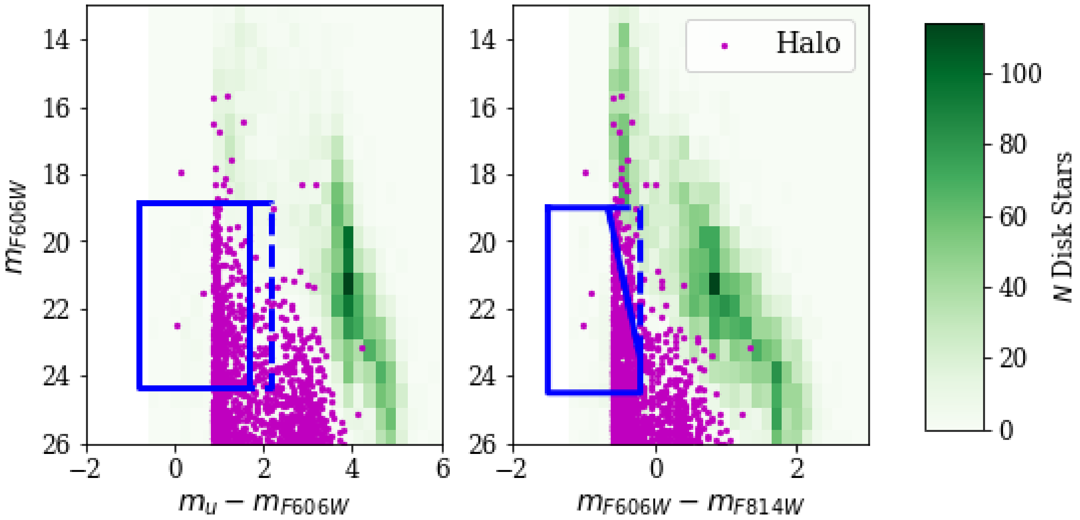

Figure 3.2: CMDs for stars in the Besançon Galaxy Model, in a 1 deg^2 field of view centered on the coordinates of EGS. The green density maps show the CMD locations of the disk stars, with the number of disk stars in each CMD bin indicated by the colorbar. Halo stars are shown in magenta; only one out of five halo stars are shown for clarity. The HALO7D selection boxes are shown in blue. Stars were assigned priority based on their positions in these two CMDs: stars were assigned top priority if they fell within both solid selection boxes, and lowest priority if they fell into only one of the dashed selection boxes.

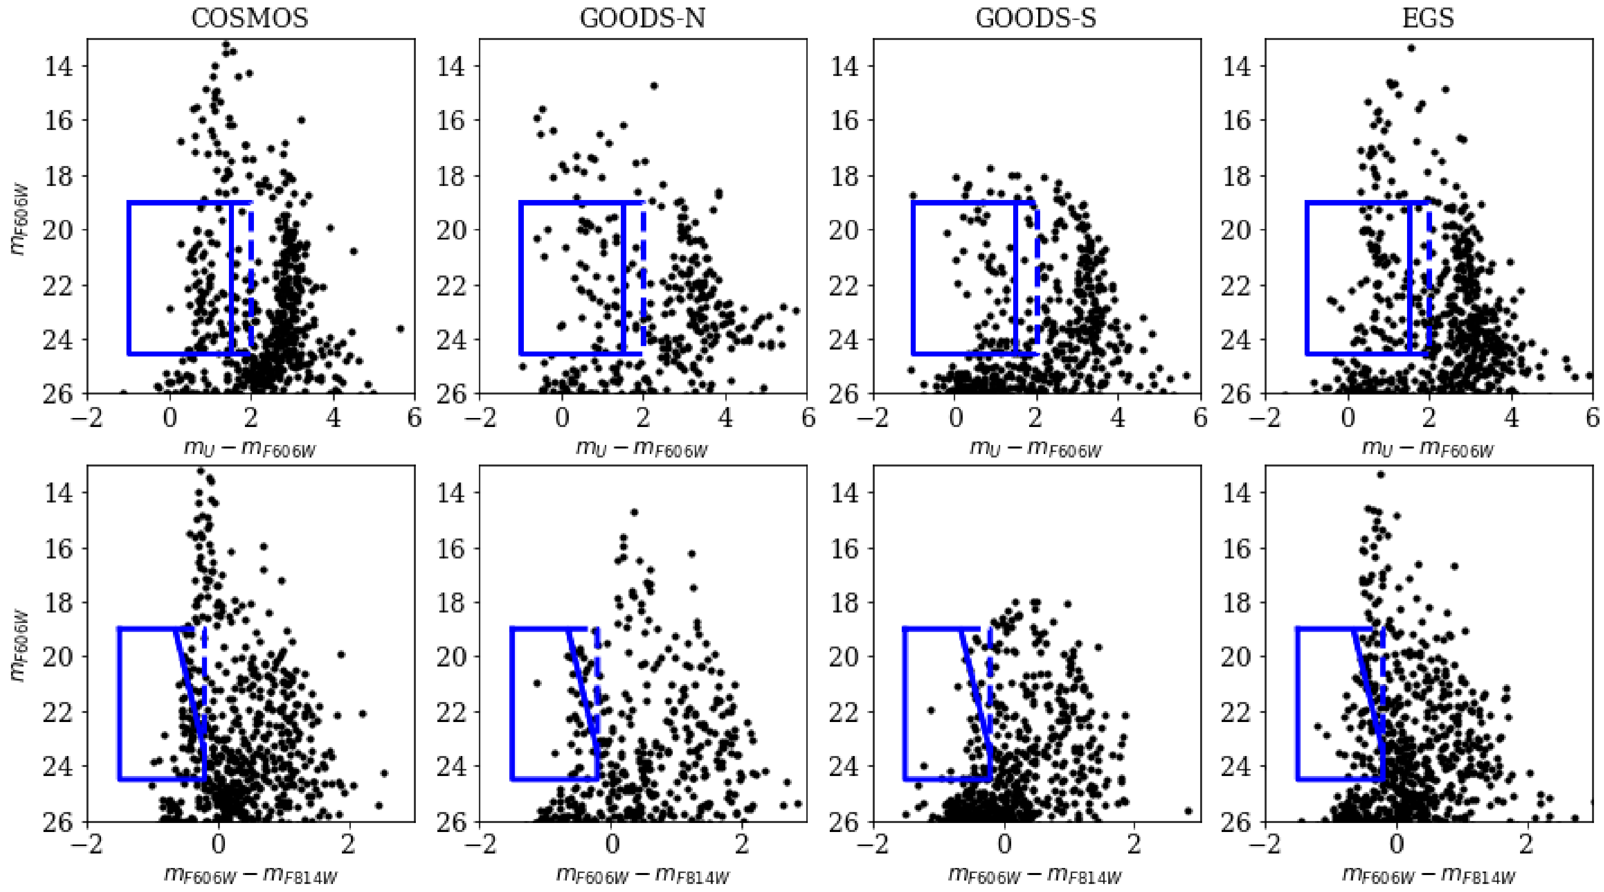

Figure 3.3: CMDs for stars in the four HALO7D fields. Selection boxes are shown in blue. Stars were assigned priority based on their positions in these two CMDs: stars were assigned top priority if they fell within both solid selection boxes, and lowest priority if they fell into only one of the dashed selection boxes. Magnitudes in F606W and F814W are computed in the STMAG system. We note that the bright stars in GOODS-S have been masked out in the catalog used for selection.

in March 2014 and ending in April 2017. While this program was intended to be completed over 19 nights in three semesters of observing, due to poor weather, observations extended through four spring semesters of observations plus several fall nights.

Observations were conducted with the same DEIMOS configuration as described in C16, and we summarize the key details here. For HALO7D observations, DEIMOS was configured with the 600 line/mm grating centered 7200 Å, resulting in a typical wavelength range of 5000-9500 Å. In the interest of limiting flux losses due to atmospheric dispersion, we divided our nights into observing "blocks" of 1-2 hours each, and tilted the slits on our masks so that their position angles were consistent with the median parallactic angle of the observing block. Our typical exposure time was 20 minutes.

Our goal was to expose each mask for 8 "effective" hours: we sought to achieve the signal to noise as a function of apparent magnitude predicted by the DEIMOS exposure time calculator for 8 hours of exposure (grey dashed lines in Figure 3.4). Signal to noise (computed at $H\alpha$) for each mask as a function of apparent magnitude is shown in Figure 3.4; in practice, we achieved a typical $\sim 5-6$ effective hours of exposure on most masks.

DEIMOS mask footprints are shown on top of HST pointings in each of the four fields in Figure 3.1. We observed eight masks in EGS, four masks in GOODS-N, two masks in GOODS-S, and four masks in COSMOS; properties of our observed masks are listed in Table 3.2. For one of our mask pointings in GOODS-N, we observed the same mask for twice as long as the other masks, but switched the list of extragalactic targets after an effective 8 hours was reached (GN3/GN4 masks have same pointings and MW target lists, but different extragalactic targets).

The slitmasks were then processed by the *spec2d* pipeline developed by the DEEP2 team at UC Berkeley (Cooper et al. 2012). Table 3.3 summarizes the progression of the HALO7D sample, from all CMD-identified halo star candidates to stars used for kinematic analysis in the subsequent sections. As seen in Figure 3.1, we weren't able to observe all of our CMD selected candidates; this is reflected in the difference between the columns 2 and 3 of Table 3.3. In addition, as with any observational program, we suffered the occasional loss due to errors in the reduction (e.g., the object is too near the edge of the mask, bad columns, etc.; column 6 of Table 3.3). There were also targets that did not achieve sufficient S/N in their spectra for a measured velocity (column 7 of Table 3.3). Finally, there were also contaminants to our sample of CMD-selected MS stars that we identified spectroscopically; we discuss these contaminants in the next subsection.

3.2.5 Spectroscopically Confirmed Contaminants

Following observations, all successfully reduced spectra were visually inspected in order to identify obvious contaminants in the MSTO sample. There were two sources of contamination that were identified spectroscopically and removed from the sample: extragalactic contaminants (column 4 of Table 3.3) and Galactic disk contaminants (column 5 of Table 3.3).

Extragalactic contaminants include quasars and emission line galaxies; our galaxy contamination rate was the highest amongst targets selected from the Barro

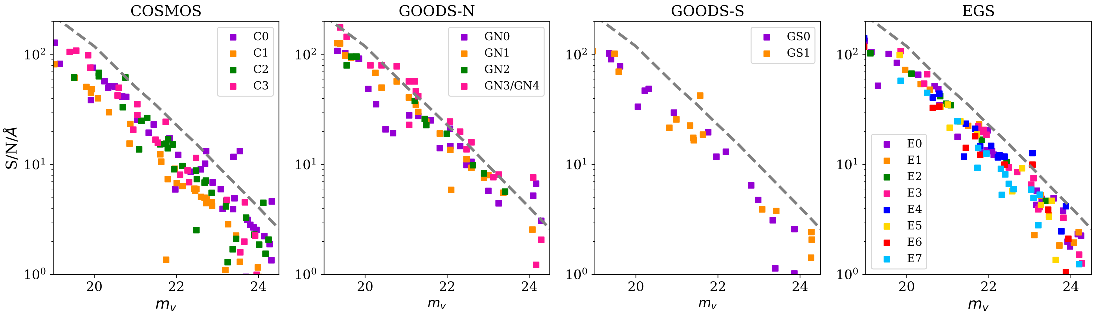

Figure 3.4: Signal to noise per angstrom for all HALO7D masks, as a function of V-band magnitude. Grey dashed lines indicate the predicted signal to noise with 8 hours of exposure time.

et al. (2011) catalog, for which we had the most generous stellarity cut.

Our spectroscopically confirmed Galactic disk contaminants are white dwarf (WD) stars and red disk stars. We identified two types of WDs in our sample; we found WDs with very broad Balmer features as well as WDs with strong continua but no absorption features (these objects have disk-like PMs; see Paper II). The red disk stars contain obvious titanium oxide features in their spectra. The red stars in our sample made it into our selection boxes because they are located on the sky close to blue galaxies, which resulted in blended colors for the ground-based u band photometry used for target selection. While we model contamination from blue disk main sequence stars in Section 3.4, we exclude the WD and red stars from our sample for dynamical modeling. In the following Section, we describe in detail how we measure LOS velocities for our target spectra. To skip straight to the results, we refer the reader to Section 3.4.

Field	Mask Name	R.A. (J2000)	Dec (J2000)	Mask P.A. (deg)	Semesters Observed	MW Targets	Extragalactic Targets
COSMOS							
	C0	10:00:36.50	+02:20:47.8	−38.2	2014A, 2015A	39	83
	C1	10:00:31.52	+02:16:14.6	−36.2	2017A	38	80
	C2	10:00:23.41	+02:11:54.4	−0.20	2017A	38	74
	C3	10:00:29.45	+02:26:16.0	111.8	2016A	23	81
GOODS-N							
	GN0	12:37:08.33	+62:12:44.6	−142.0	2015A, 2016A	23	95
	GN1	12:37:01.22	+62:14:05.3	46.1	2014A, 2015A	24	105
	GN2	12:36:38.82	+62:15:48.8	51.3	2016A	12	95
	GN3,GN4	12:36:58.73	+62:13:02.6	9.4	2014A, 2015A, 2016A	23, 21	104, 94
GOODS-S							
	GS0	03:32:18.81	−27:49:04.9	−17.3	2015B, 2016B	25	92, 88, 90
	GS1	03:32:47.28	−27:47:26.8	−15.6	2015B, 2016B	24	85, 85, 85
EGS							
	E0	14:20:15.98	+53:01:13.9	180.0	2014A	30	101
	E1	14:18:48.21	+52:45:18.4	−26.1	2014A	19	103
	E2	14:19:33.70	+52:49:40.4	−28.1	2014A	24	97
	E3	14:19:51.37	+52:55:04.1	−30.4	2015A	28	81
	E4	14:20:42.48	+53:05:50.4	177.8	2015A	34	80
	E5	14:19:10.83	+52:47:16.3	−25.8	2016A	22	81
	E6	14:21:03.34	+53:04:42.3	−144.7	2017A	29	87
	E7	14:19:49.96	+52:56:05.3	164.0	2017A	31	79

Table 3.2: Summary of the masks observed through HALO7D. In GOODS-N, the masks GN3 and GN4 contain the same MW targets, but different extragalactic targets. In GOODS-S, three sets of extragalactic targets were observed alongside each mask of MW targets.

Field	N_{CMD}	N_{Obs}	N_{Gal}	N_{WD}	N_{RedStars}	$N_{\mathrm{ReductionErrors}}$	$N_{\mathrm{LowS/N}}$	N Halo Star Candidates
COSMOS	101,67	87,36	5,2	1	2	4	21	88
GOODS-N	48	47	1	2	4	1	6	33
GOODS-S	57	49	8	2	1	2	11	25
EGS	96,135	84,94	6,28	7	7	8	25	97

Table 3.3: Summary of the progression of the HALO7D sample, from CMD selected targets to objects used in kinematic analysis. In COSMOS and EGS, we first indicate halo star candidate selected from the CANDELS catalogs, followed by candidates from the secondary catalogs for the first three columns of the table (see Section 3.2.2).

3.3 Hierarchical Bayesian LOS Velocities: Velociraptor

In this section, we describe in detail the model implemented by the VELOCIRAPTOR software. As explained in Section 3.1.2, different observations of the same star will have different raw velocities, due to the motion of the Earth around the Sun as well as slit miscentering. We demonstrate this effect in Figure 3.5, which shows the $H\alpha$ region and the telluric A-band region for two spectra of a relatively bright HALO7D target ($m_{F606W} = 19.1$) observed on different nights. The raw spectra are clearly not at the same velocity: while ~ 5 km/s of this velocity offset is due to the Earth's motion around the Sun, the remaining 50 km/s offset is entirely due to the misalignment of the object in the slit.

As such, applying these corrections prior to co-adding or stacking spectra is essential in order to accurately estimate the velocity of a star. However, because the A-band correction is measured from an absorption feature, if the spectrum is faint and noisy, the estimate of the A-band correction will also be noisy.

In order to address these challenges, we present the VELOCIRAPTOR technique. VELOCIRAPTOR implements a Bayesian hierarchal model, modeling the raw velocities and A-band corrections of all observations of a star simultaneously. Standard practice would be to stack spectra and then measure a velocity, usually using a cross correlation method (e.g., spec1d; Newman et al. 2013) or a maximum-likelihood method (such as the Penalized Pixel Fitting method of Cappellari & Emsellem 2004). However, stacking before measuring a velocity neglects the A-band corrections (and associated uncertainties) of different observations. Bayesian hierarchical modeling provides a natural, fully

Level	Parameters	Prior	Description
Spectral Reg.	θ_{Line}		
	v_{Line}	$p(v_{\text{Line}}) \propto \text{const}$	Velocity of region
	C	$C + 1 \sim \text{Gamma}(2,\ 2)$	Absorption line strength parameter
	b_l	$p(b_l) \propto \text{const}$	Legendre polynomial coefficients for continuum
Obs. with J Reg.	θ_{spec}		
	v_{raw}	$p(v_{\text{raw}}) \sim \text{Unif}[-600, 600]$	Velocity of stellar absorption regions (e.g.,$H\alpha$)
	v_{Aband}	$p(v_{\text{Aband}}) \sim \text{Unif}[-100, 100]$	Velocity of telluric region(s)
	$C_1, ..., C_J$	$C_j + 1 \sim \text{Gamma}(2,\ 2)$	Absorption line strength parameters
	$b_{l,1}, ..., b_{l,J}$	$p(b_{l,j}) \propto \text{const}$	Legendre polynomial coefficients for continuum
Star with K Obs.	θ		
	v	$p(v) \sim \text{Unif}[-600, 600]$	LOS velocity of the star
	σ_v	$p(\sigma_v) \sim \text{Inv} - \text{Gamma}(7, 72)$	Dispersion of measurements; systematic error
	$v_{\text{raw},k}, v_{\text{Aband},k}$	$v_{\text{corr}} = v_{\text{raw},k} - v_{\text{Aband},k} - v_{\text{helio},k}$ $v_{\text{corr},k} \sim \text{N}(v, \sigma_v^2)$	Raw and A-band velocities of individual spectra
	$C_{1,1}, ..., C_{J,K}$	$C_{j,k} + 1 \sim \text{Gamma}(2,\ 2)$	Absorption line strength parameters
	$b_{l,1,1}, ..., b_{l,J,K}$	$p(b_{l,j,k}) \propto \text{const}$	Legendre polynomial coefficients for continuum

Table 3.4: Summary of model parameters (and priors) for different levels of our hierarchical model. In the first level of the model, we model a region of the spectrum containing an absorption line feature, such as Hα, the telluric A-band region, or the CaT region. In the next level, we use multiple spectral regions to estimate the corrected velocity of a star from a single spectroscopic observation. Finally, in our hierarchical model, we incorporate multiple spectroscopic observations into our estimate of the corrected velocity of the star.

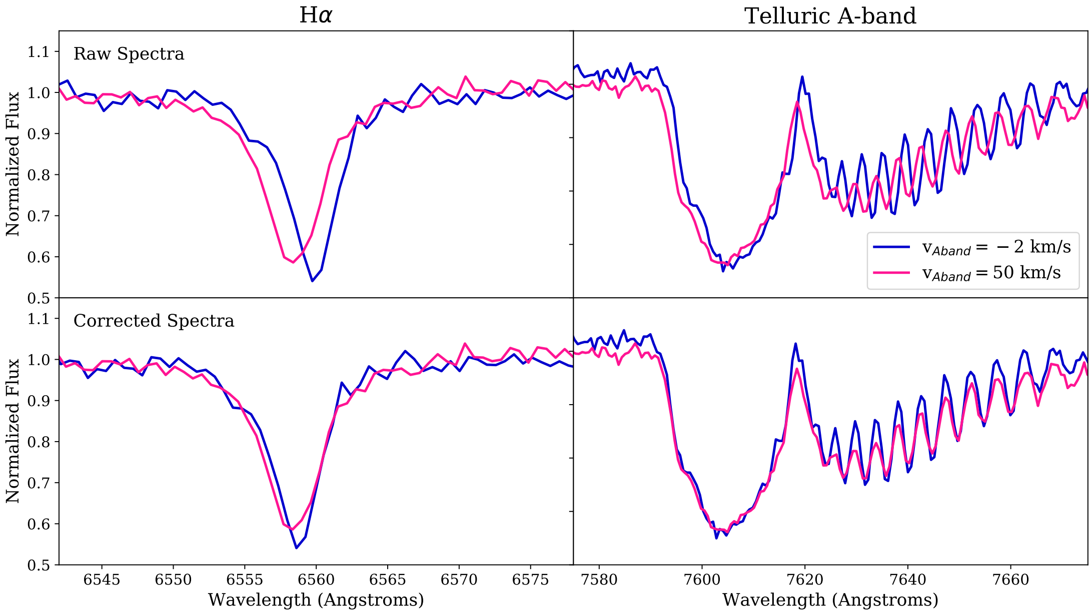

Figure 3.5: Illustration of the velocity offset caused by slit miscentering. Blue and pink lines show spectra of the same star taken during different observing runs; lefthand panels show the Hα region of the spectrum, and righthand panels show the telluric A-band region. Top panels show the raw spectra, in the observed frame, uncorrected for heliocentric motion as well as slit miscentering; the Hα lines and telluric absorption lines are clearly misaligned. The lower panels show the spectra with heliocentric and A-band corrections applied.

probabilistic framework for incorporating all available information in the spectra while properly accounting for uncertainties.

In Section 3.3.1, we define terminology used to describe our model. We then explain how we model individual spectroscopic observations in Section 3.3.2. Section 3.3.3 describes the hierarchical model employed to measure the velocity of a star from multiple observations. Details of fake data testing, including sample trace and corner plots, can be found in the Appendix.

3.3.1 Definitions

We begin by explicitly defining some terminology and symbols used in the description of our method.

Spectral Regions: a segment of a spectrum, a few hundred Ångstroms in wavelength, centered on an absorption feature. We use the letter j to denote a specific spectral region; for example, $F_j(\lambda)$ denotes the flux as a function of wavelength for spectral region j.

Pixels: Each spectral region contains pixels $i = 1, ..., I$. The value I_j denotes the total number of pixels in spectral region j.

$\lambda_i \to x_{ij}$: For evaluating polynomials, we rescale the wavelength array of a given spectral region j onto the range [-1,1]:

$$x_{ij}(\lambda_i) = \frac{2(\lambda_i - \lambda_{\mathrm{min},j})}{\lambda_{\mathrm{max},j} - \lambda_{\mathrm{min},j}} - 1. \tag{3.2}$$

So, x_{ij} denotes the value that λ_i takes when rescaled onto the range determined by the range of spectral region j: $[\lambda_{\mathrm{min},j}, \lambda_{\mathrm{max},j}]$. Using this definition, $x_{0,j} = x(\lambda_{\mathrm{min},j}) = -1$, and $x_{I,j} = x(\lambda_{\mathrm{max},j}) = 1$.

Observations: Each spectrum has $k = 1, ..., K$ observations. Therefore, each observation has its own $v_{\mathrm{raw},k}$, $v_{\mathrm{Aband},k}$, $v_{\mathrm{helio},k}$.

Distributions: We use standard statistical notation to express random variable distributions, which will include normally-distributed ($x \sim \mathrm{N}(\mu, \sigma^2)$) and Gamma-distributed ($x \sim \mathrm{Gamma}(a, b)$). We also make use of inverse-gamma distributions in this work; the inverse-gamma distribution is the conjugate prior for the variance of a normal distribution. If a random variable x is distributed according to an inverse-gamma distribution with parameters a, b (i.e., $x \sim \mathrm{Inv} - \mathrm{Gamma}(a, b)$), the corresponding probability density function is

$$p(x|a, b) = \frac{b^a}{\Gamma(a)} x^{-(a+1)} \exp\{-b/x\}, \tag{3.3}$$

where $\Gamma(a)$ denotes the gamma function.

3.3.2 Single-Mode: Modeling a Single Spectrum

We first present our Bayesian method of estimating the velocity of a star from a single spectrum. To estimate a stellar velocity, we use the spectral regions that contain the most velocity information. We typically use 3 regions: the region around Hα ($6500 - 6650$ Å) and the Calcium triplet region (CaT; $8450 - 8700$ Å) to estimate the raw velocity, and the telluric A-band region ($7500 - 7750$ Å) for the A-band correction.

To model a region of the spectrum, we first begin with a template. Our templates consist of bright velocity standards that were observed with a very similar configuration to our science spectra. The templates used in this analysis are described in detail by Toloba et al. (2016); they have high signal-to-noise ratios (100–800 Å^{-1}), and span a range of spectral types (from B1 to M8) and luminosity classes (from dwarfs to supergiants).

For the HALO7D targets, we use the template HD105546. While this star is a horizontal branch star, its color is consistent with the color range of our targets, and its spectrum has absorption in Hα and CaT. In order to estimate the raw velocity of our template star, we use a simple model of a polynomial with inverted Gaussians for the absorption lines. Because the template star was trailed through the slit during observation, it does not suffer from slit miscentering, so its A-band correction is 0 km/s. We verify that no additional correction to the wavelength solution is required by checking the consistency of the velocities measured at Hα and CaT.

We use the spectrum of HD105546, shifted to the rest frame, to estimate the velocity of $H\alpha$ and CaT regions of the HALO7D target. We use the same spectrum in the observed frame (i.e., unshifted) to estimate the A-band velocities of the HALO7D targets. We model the different regions of the target spectrum separately, while demanding that the velocities at $H\alpha$ and CaT be the same.

To model a spectral region, we allow the velocity, absorption line strength, and continuum level to vary. Our vector of parameters, which we denote as θ_{Line}, are the velocity v_{Line}, the absorption line strength C, and the Legendre polynomial coefficients b_l which control the continuum level. Given that we look at narrow spectral regions, we find $l = 1$ (i.e., a straight line with varying slope and intercept) sufficient to model the continuum.

As a function of scaled wavelength x, our model $M(x, \theta_{\mathrm{Line}})$ can be written as:

$$M_{\mathrm{Line}}(x, \theta_{\mathrm{Line}}) = \sum_l b_l P_l(x) \times \frac{T_{v_{\mathrm{Line}}}(x) + C}{1 + C}, \tag{3.4}$$

where P_l are the Legendre polynomials and $T_{v_{\mathrm{Line}}}(x)$ is the template flux, shifted to velocity v_{Line}.

The likelihood of the observed spectral flux F_{Line} given the model parameters is thus:

$$p(F_{\text{Line}}|\theta_{\text{Line}}) = \prod_{i=0}^{I} \mathrm{N}(F_{\text{Line}}(x_i)|M_{\text{Line}}(x_i, \theta_{\text{Line}}), \sigma_i^2), \tag{3.5}$$

where, for $i = 1, ..., I$ pixels, x_i is the rescaled wavelength value, $F_{\text{Line}}(x_i)$ is the flux at that rescaled wavelength, $M(x_i, \theta_{\text{Line}})$ is the model flux, and σ_i is the noise in that pixel (as returned by the spec2d pipeline).

We can write down the posterior probability distribution for our model parameter making use of Bayes' Theorem. Bayes' Theorem gives the probability of a vector of model parameters θ given a vector of of data y:

$$p(\theta|y) = \frac{p(y|\theta)p(\theta)}{p(y)}, \tag{3.6}$$

where $p(y|\theta)$ is the likelihood of the data given the parameters; $p(\theta)$ is the prior probability of the parameters; and $p(y)$ is the probability of the data (in practice, this term serves as a normalization). In order to sample from the posterior distribution for our model parameters θ_{Line}, we must specify their prior distributions: $p(\theta_{\text{Line}}|F_{\text{Line}}) \propto p(F_{\text{Line}}|\theta_{\text{Line}})p(\theta_{\text{Line}})$.

Our prior distributions are listed in Table 3.4; we generally assume reference (i.e., Jeffreys) priors. For the absorption line parameters C, we assign a Gamma distribution prior to the quantity $C + 1$. The Gamma distribution is defined over the range $x > 0$, and is a common prior choice for scale parameters. Because of our chosen parameterization, if $C < -1$ the line becomes an emission line. In addition, as C becomes large, the absorption line becomes indistinguishable from the continuum. We therefore assign a Gamma(2, 2) prior to $C + 1$ in order to constrain the possible allowed values for C.

The total posterior probability of the full vector of spectrum parameters θ_{spec} is given by the product of the posterior probabilities of the different lines used:

$$p(\theta_{\mathrm{spec}}|\mathcal{F}) \propto \prod_{j=1}^{J} p(F_{\mathrm{Line},j}|\theta_{\mathrm{Line}})p(\theta_{\mathrm{Line}}), \tag{3.7}$$

where $\theta_{spec} = (v_{\mathrm{raw}}, v_{\mathrm{Aband}}, C_1, .., C_J, b_{l,1}, ..., b_{l,J})$ is the full vector of parameters describing the spectrum. Here we are denoting $\mathcal{F} = \{F_{\mathrm{Line},1}, ..., F_{\mathrm{Line},J}\}$ as the set of fluxes over all spectral regions. When modeling three spectral regions, θ_{spec} contains 11 free parameters.

In order to sample from the posterior, we use `emcee` (Foreman-Mackey et al. 2013), a PYTHON implementation of the Goodman & Weare (2010) affine-invariant Markov chain Monte Carlo (MCMC) ensemble sampler. We first initialize our walkers by estimating the parameters one at a time. Results from extensive fake data testing, including sample trace and corner plots, can be found in the Appendix.

3.3.3 Hierarchical Modeling

In order to combine spectra from different observations, we employ Bayesian hierarchical modeling. While Bayes' Theorem (Equation 3.6) gives the probability of a vector of model parameters θ given a vector of of data y, it is often desirable for the parameters themselves to be drawn from a distribution, whose values we would like to estimate. These *hyperparameters* (φ), are incorporated into Bayes' Theorem as follows:

$$p(\theta, \varphi | y) \propto p(y|\theta, \varphi) p(\theta|\varphi) p(\varphi). \tag{3.8}$$

$p(\theta|\varphi)$ is the probability of the hyperparameters given parameters θ and $p(\varphi)$ is the hyperprior.

For our model for multiple observations of a star, we have two hyperparameters: v, which is the "true" velocity of the star, and σ_v, the dispersion of velocity measurements (this term serves to model additional uncertainty/noise not captured by the reduction pipeline). For the velocity of a star with K observations with spectra $\mathcal{F}_1, ..., \mathcal{F}_K$, the full posterior is given by

$$\begin{split} p(v, \sigma_v, \theta_{\text{spec},1}, ..., \theta_{\text{spec},K} | \mathcal{F}_1, ..., \mathcal{F}_K) = \\ p(v_1, ..., v_K | v, \sigma_v) \times \prod_{k=1}^{K} p(\theta_{\text{spec},k} | \mathcal{F}_k) p(v, \sigma_v), \end{split} \tag{3.9}$$

where $v_k = v_{\text{raw},k} - v_{\text{Aband},k} - v_{\text{helio},k}$ is the corrected velocity for observation k, and $p(v, \sigma_v)$ is the prior distribution on the hyperparameters.

For these measurements, we can consider the fact that we have substantial prior information about the extent to which these velocities should agree: we know that we are observing the same star with each observation. Therefore, it does not make sense for σ_v to be arbitrarily large, and a standard non-informative prior is not necessarily the best choice. We assign σ_v^2 to be drawn from an Inverse-gamma distribution with parameters $a = 7, b = 72$. This prior distribution has a mean of $b/(a-1) = 12$, a mode of $b/(a+1) = 8$, and variance $b^2/((a-1)^2(a-2)) = 28.8$. This prior distribution therefore assigns highest probability to σ_v in the range of $3-4$ km/s, but does allow for σ_v to take on higher values if demanded by the data.

Given the complexity of our model, we use `emcee` to sample from the posterior. All our model parameters and prior distributions, for each level of the model, are listed in Table 3.4.

A demonstration of this technique is shown in Figure 3.6, for a HALO7D target with $m_{F606W} = 22.0$. This target was observed seven times, in varying conditions, over the course of Spring 2015. Each of the histograms in the top panel represents the posterior distribution for the *corrected* velocity for each of these seven observations. The varying widths of these PDFs reflect the varying quality in observing conditions across the different nights of observing. These posterior samples were derived using `emcee` to sample the posterior distribution given in Equation 3.7. However, once we link the observations by combining them with the hierarchical model, the posterior distributions for the individual velocities converge (lower panel). The posterior distribution for the corrected velocity of the star is shown in black: this posterior incorporates all information, as well as sources of uncertainty, from the seven observations.

The resulting velocity uncertainties for the HALO7D targets as a function of m_{F606W} apparent magnitude are plotted in Figure 3.7. Velocity errors are computed as half the difference between the 84th and 16th posterior percentiles: $v_{err} = (v_{84} - v_{16})/2$. At the bright end of our sample, our velocity uncertainties are as low as 1-2 km/s;

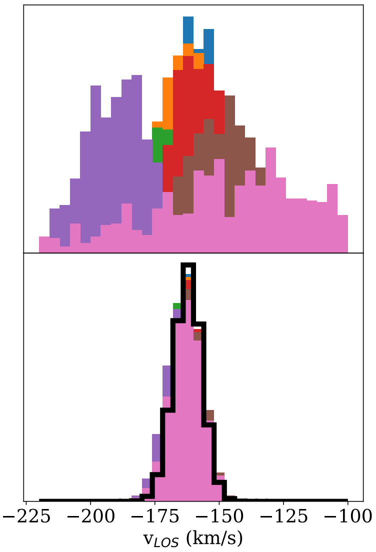

Figure 3.6: Histograms of posterior samples for the corrected velocity of a HALO7D target ($m_{F606W} = 22.0$) from seven observations. The top panel shows the posterior samples for the velocities when the spectra are modeled independently: note that these are the PDFs for $v = v_{raw} - v_{Aband} - v_{helio}$. The bottom panel shows the PDFs for the individual observations once they have been combined into the hierarchical model. The final PDF for the corrected velocity (thick black line), incorporating all observations, thus folds in all information and uncertainty from all observations of a star.

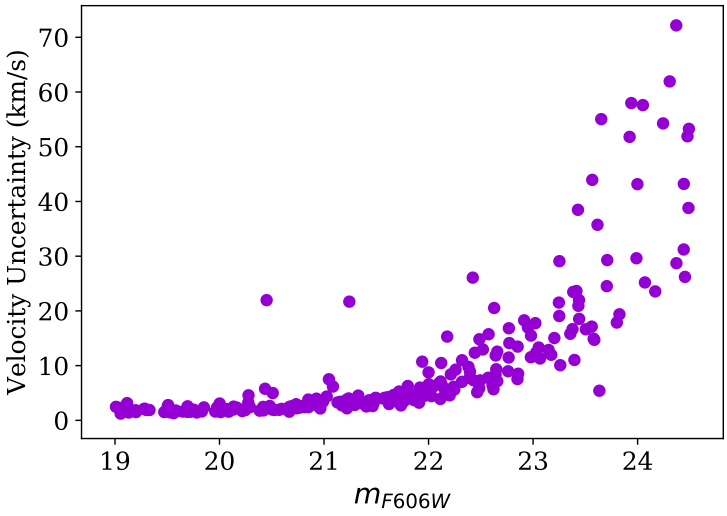

Figure 3.7: Velocity uncertainties for the HALO7D sample as a function of m_{F606W} apparent magnitude. Velocity uncertainties are the 16th and 84th percentiles of the posterior distributions.

velocity uncertainties remain below 10 km/s for stars brighter than $m_{F606W} = 22$. Stars at the faint end of our sample reach velocity uncertainties as high as ~ 50 km/s.

For more details on testing our method on fake data, including sample trace and corner plots, please see the Appendix.

3.4 Results

3.4.1 LOS Velocity Distributions

We use the velocities measured from VELOCIRAPTOR to study the LOS velocity distributions of the stellar halo. Heliocentric LOS are converted to the Galactocentric Standard of Rest (GSR) frame by assuming a circular speed of 240 km s^{-1} at the position of the Sun ($R_0 = 8.5$ kpc), with solar peculiar motion $(U, V, W) = (11.1, 12.24, 7.25)$ km

s^{-1} (Schönrich et al. 2010).

Figure 3.8 shows cumulative histograms for the LOS velocity distributions (in the GSR frame) for the four HALO7D fields. To capture the effects of our velocity uncertainties, we have plotted 100 realizations of the velocity cumulative distribution, each time drawing a new value for every velocity from its posterior. Therefore, the apparent thickness of a given step in the histogram is an indication of the uncertainty of that measurement. We also show traditional histograms of the LOS velocities in Figure 3.9.

Based on the histograms, we see that our samples across all four fields are dominated by a "hot halo" population; while their could be hints of substructure in these fields, we find that none of our fields are dominated by kinematically cold substructure, which would appear as narrow ($5 - 15$ km/s) peaks in the velocity distributions. We leave the discussion of the search and characterization of potential substructure in these fields to future work, where we will also utilize PMs and abundances.

To estimate the LOS velocity dispersion of the halo, σ_{LOS}, we model the LOS velocity distributions as a two component mixture model, with a halo component and a disk component. We model the halo velocity distribution as a normal distribution with unknown mean and variance: $v \sim \mathrm{N}(\langle v_{LOS} \rangle, \sigma^2_{LOS})$.

We model the disk velocity distribution along each line of sight as a skew-normal distribution, with skew parameter α, location parameter ζ, and scale parameter ω. The likelihood of an observed velocity v_i given disk parameters is given by:

$$p(v_i|\alpha,\zeta,\omega) = \frac{2}{\omega}\phi\left(\frac{v_i-\zeta}{\omega}\right)\Phi\left(\alpha\left(\frac{v_i-\zeta}{\omega}\right)\right), \tag{3.10}$$

where $\phi(x)$ is the standard normal PDF and $\Phi(x)$ is the standard normal CDF. We fix the parameters of the disk velocity distribution, but leave the fraction of disk contamination as a free parameter. We denote our disk PDF as $g_{\text{disk}}(v_i) = p(v_i|\alpha,\zeta,\omega)$.

To determine the parameters of our disk model, we use the Besançon Galaxy Model (Robin et al. 2003). We use synthetic catalogs of 1 square degree areas centered on the coordinates of our survey fields (the larger area gives us better statistics for our simulated disk and halo populations). We then determine the velocity distribution of the (non-WD) disk contaminants within the HALO7D selection boxes, and fit a skew normal to this distribution. The resulting parameters for the disk distributions can be found in Table 3.5; they are also plotted in Figures 3.8 and 3.9.

For this mixture model of disk and halo, the likelihood of a given LOS velocity observation v_i, with error $\sigma_{v,i}$, is given by

$$\begin{aligned} p(v_i|f_{\text{disk}},\langle v_{LOS}\rangle,\sigma_{LOS},g_{\text{disk}}) &= f_{\text{disk}}\times g_{\text{disk}}(v_i) \\ &+(1-f_{\text{disk}})\times N(v_i|\langle v_{LOS}\rangle,\sigma_{LOS}^2+\sigma_{v,i}^2). \end{aligned} \tag{3.11}$$

In order to sample from the posterior distribution for our model parameters, we first must assign prior distributions. We assign a standard uniform $[0,1]$ prior on the fraction of disk contamination, and we assign the Jeffreys prior to the mean and dispersion for the halo LOS velocity distribution ($p(\langle v_{LOS}\rangle,\sigma_{LOS}) \propto 1/\sigma_{LOS}$).

Our posterior is thus:

$$
\begin{aligned}
p(\langle v_{LOS}\rangle, \sigma_{LOS}, f_{\mathrm{disk}}|v) \propto\ & p(\langle v_{LOS}\rangle, \sigma_{LOS})p(f_{\mathrm{disk}}) \\
& \times \prod_{i=1}^{N_{stars}} p(v_i|\langle v_{LOS}\rangle, \sigma_{LOS}, f_{\mathrm{disk}}).
\end{aligned}
\tag{3.12}
$$

We use `emcee` to sample from this posterior. We used 500 walkers, ran the sampler for 1000 steps, and discarded the first 800 steps as burn-in. Median posterior values, along with error bars from the 16/84 percentiles, are quoted for the three model parameters in Table 3.5.

Posterior draws are overplotted on the histograms in Figures 3.8 and 3.9. Each pink line in Figure 3.8 is the CDF corresponding to a draw from the posterior for $\langle v_{LOS}\rangle$ and σ_{LOS}. In Figure 3.9, the amplitude of the disk PDFs reflects the uncertainty in the disk contribution. Thicker pink and green lines indicate the distributions corresponding to the median posterior values.

Histograms of posterior samples for our three free parameters are shown in Figure 3.10. The left panel shows the posterior distributions for $\langle v_{LOS}\rangle$; all fields have mean LOS velocity consistent with 0 km/s. The middle panel of Figure 3.10 shows the posterior samples for σ_{LOS}; posterior PDFs for σ_{LOS} are consistent across the four fields. The widths of the individual PDFs vary according to the sample size in a given field, but the PDFs substantially overlap. In the COSMOS field, we estimate $\sigma_{LOS} = 123^{+12}_{-11}$ km/s; in GOODS-N, $\sigma_{LOS} = 110^{+16}_{-13}$ km/s; for GOODS-S, $\sigma_{LOS} = 122^{+30}_{-21}$ km/s; and in EGS, we find $\sigma_{LOS} = 118^{+11}_{-9}$ km/s.

The right panel of Figure 3.10 shows posterior samples for f_{disk} in the four fields. We note that no color or distance information is incorporated into our estimates of the disk contamination, and that this estimate is based on LOS velocities alone. Our estimates of our disk contamination will be more accurate once PM and photometric information are incorporated. EGS and GOODS-N show $0 - 10\%$ disk contamination, consistent with the predictions from Besançon (see Table 3.5). COSMOS, our lowest latitude field, shows a slightly higher level of disk contamination ($\sim 25\%$, while predicted to be $\sim 11\%$). The posterior distribution for disk contamination in GOODS-S, our field with the smallest sample size, is very broad, with a posterior median of 30%, much higher than the 7% predicted by Besançon. We note that this high posterior median is largely due to the small sample size, and that the disk contamination in GOODS-S is poorly constrained based on LOS velocities alone.

3.4.2 Comparison with Other Tracers

In this section, we compare our results for the HALO7D LOS velocity distributions with other studies conducted over a similar distance range (though using different tracer populations). In order to compare our measurements of σ_{LOS} with other studies, we first need to estimate the distance range probed by our sample. We estimate these distance distributions in a similar method to that of D13 and C16, using weighted isochrones to derive the PDF for the absolute magnitude M_{F814W} of a star given its $m_{F606W} - m_{F814W}$ color.

We weight the VandenBerg et al. (2006) isochrones according to a Salpeter IMF, an age and a metallicity distribution typical of halo stars. We assume that the

Field	$\langle v_{LOS} \rangle$ (km/s)	σ_{LOS} (km/s)	f_{disk}	Predicted f_{disk}	α	ζ (km/s)	ω (km/s)	$\langle D \rangle$ (kpc)
COSMOS	13^{+23}_{-19}	123^{+12}_{-11}	$0.23^{+0.12}_{-0.11}$	0.11	0.6	−115	56	21
GOODS-N	6^{+20}_{-21}	110^{+16}_{-13}	$0.11^{+0.12}_{-0.06}$	0.07	0.0	66	46	20
GOODS-S	24^{+48}_{-35}	122^{+30}_{-21}	$0.34^{+0.22}_{-0.20}$	0.07	−1.1	−27	58	23
EGS	9^{+15}_{-14}	118^{+11}_{-9}	$0.13^{+0.10}_{-0.08}$	0.07	−1.1	121	58	21

Table 3.5: Summary of results from the modeling of the LOS velocity distributions. Posterior medians are quoted, with error bars giving the 16/84 percentile credible regions. We list the predicted fraction of disk contamination from Besançon; because we have included no color or apparent magnitude information into our estimate of the disk contamination, we expect our f_{disk} estimates are higher than the actual disk contamination in our sample. We list the parameters for the disk model for each field: the skew α, the location ζ, and the scale ω. These values were derived by fitting a skew normal distribution to the velocities of disk stars that fall in the HALO7D selection box in the Besançon Galaxy Model. Finally, we list the average distance to each field, as computed in Equation 3.13.

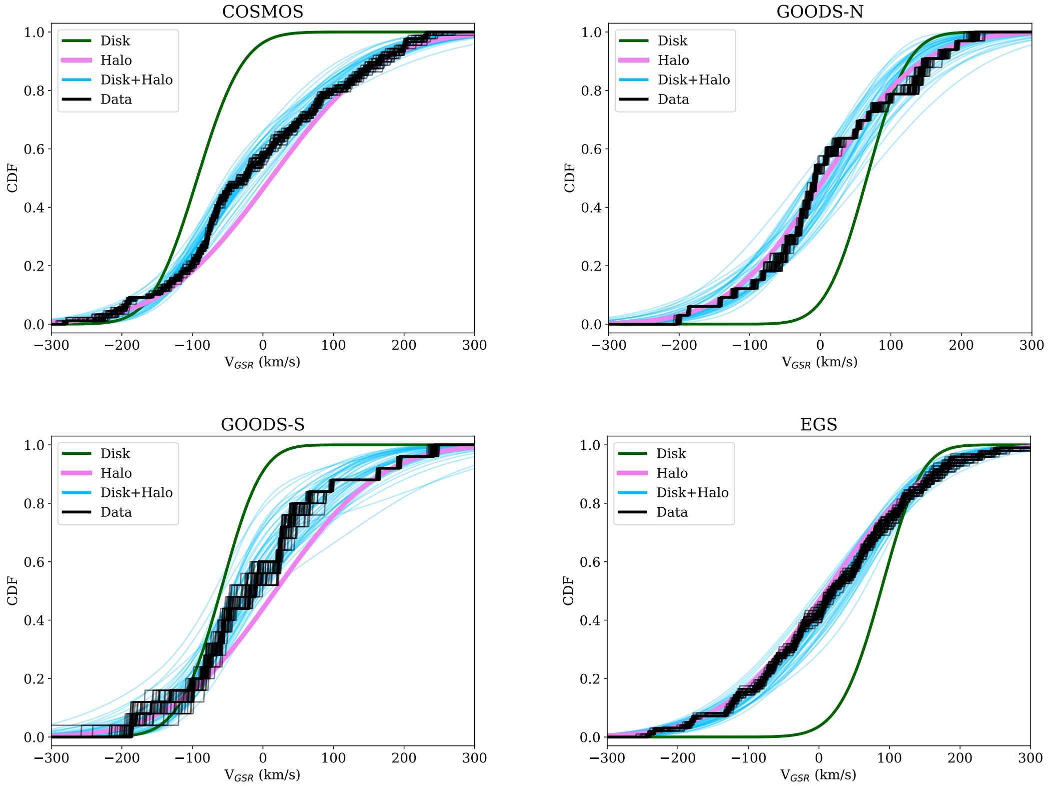

Figure 3.8: Cumulative LOS velocity histograms in the four HALO7D fields. Velocities are given with respect to the Galactic Standard of Rest (GSR). Black lines indicate the CDFs for the data: for each of the 100 black lines, velocity values were drawn from the posterior distributions for the measurements. The width of each step thus demonstrates the velocity uncertainty for that data point. The CDF for the disk model is shown in green, and the CDF for the halo model, using the posterior median values for $\langle v_{LOS} \rangle, \sigma_{LOS}$ is shown in pink. Blue lines show the CDF for the mixture of the disk and halo populations, for 100 draws from the joint posterior distribution for $f_{\rm Disk}, \langle v_{LOS} \rangle, \sigma_{LOS}$.

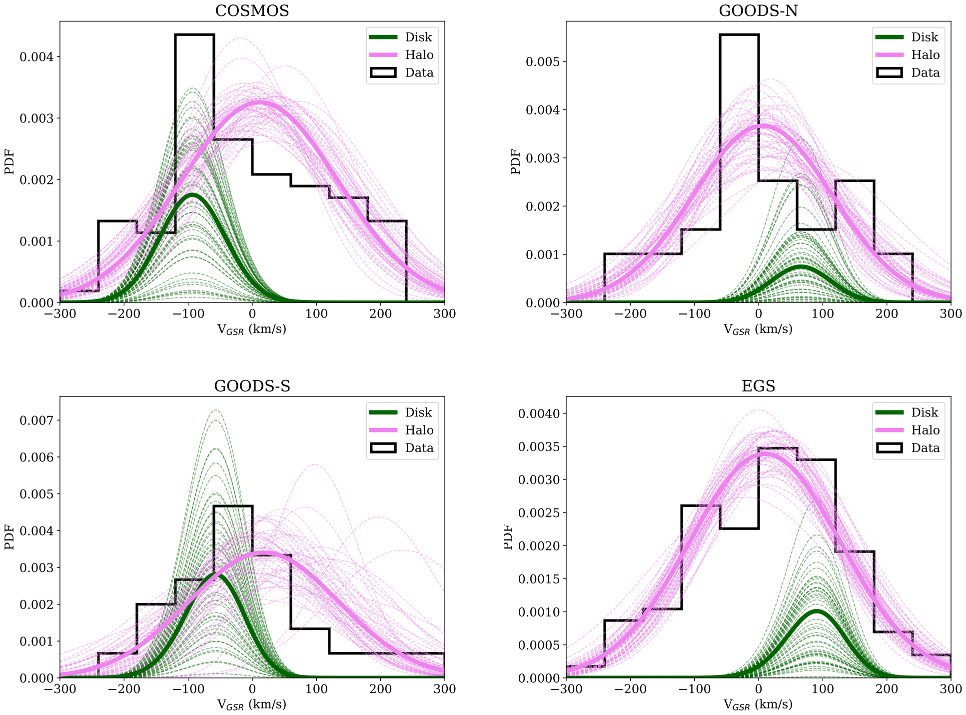

Figure 3.9: LOS velocity histograms in the four HALO7D fields. Shown in pink are the resulting velocity distributions from 50 draws from the posteriors for $\langle v_{LOS} \rangle$ and σ_{LOS}. The green line indicates the disk distribution. The parameters of the disk velocity distribution are fixed; only the fraction to the total contribution is allowed to vary. Bold lines show the corresponding distributions for the median posterior values of f_{disk}, $\langle v_{LOS} \rangle$ and σ_{LOS}.

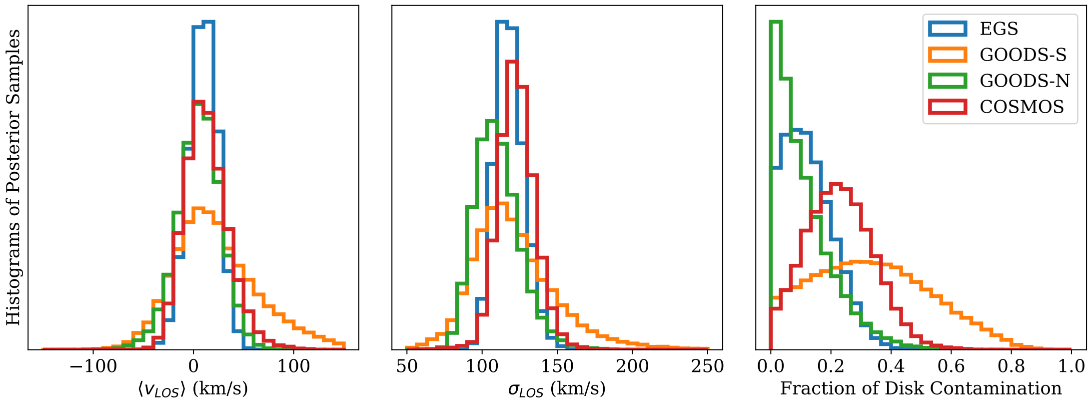

Figure 3.10: Posterior samples for the mean LOS velocity (right), LOS velocity dispersion (middle) and the fraction of disk contamination (right) for all four fields. All fields have mean LOS velocity consistent with 0 km/s, and dispersions consistent with one another.

age and metallicity distributions are Gaussian, with $\langle T \rangle = 12$ Gyr, $\sigma_T = 2$ Gyr (e.g., Kalirai 2012), $\langle[\mathrm{Fe/H}]\rangle = -1.9$ and $\sigma_{\mathrm{Fe/H}} = 0.5$ (e.g., Xue et al. 2008). We model the resulting weighted CMD with a kernel density estimate (KDE), using a kernel bandwidth of 0.025. The resulting PDFs for M_{F814W}, for six different colors, are shown in Figure 3.11.

Using the PDF for absolute magnitude as a function of color in conjunction with the halo density profile (Deason et al. 2011), we derive the PDF for the log distance distribution to our sample:

$$
\begin{aligned}
p(\log D | m_{F814W}, m_{F606W}, \rho) \propto\ & p(\log D|\rho) \times \\
& \sum_{n=1}^{N_{obj}} p(M_{F814W}(\log D) | m_{F606W,n} - m_{F814W,n}),
\end{aligned}
\tag{3.13}
$$

where $p(\log D|\rho)$ is the probability of $\log D$ given the Deason et al. (2011) density profile, and $p(M_{F814W}(\log D)|m_{F606W,n} - m_{F814W,n})$ is the probability of object n having absolute magnitude $M_{F814W}(\log D)$ given its color $m_{F606W,n} - m_{F814W,n}$.[4] We then estimate the mean distance to each field $\langle D \rangle = \int D \times p(\log D) d \log D$. Each of the four fields has an average distance $\langle D \rangle \sim 20$ kpc. Figure 3.12 shows the cumulative logarithmic distance PDFs of our samples across the four fields, and average distances to each field are listed in Table 3.5. We note that no kinematic information is incorporated into our distance estimate, and that information from 3D kinematics will improve our distance estimates in subsequent work.

Figure 3.13 shows the LOS dispersions of the four HALO7D fields plotted as a function of mean galactocentric radius (where we have converted heliocentric distance $\langle D \rangle$ to Galactocentric radius $\langle r \rangle$). Points indicate the median of the σ_{LOS} posterior samples, and error bars indicate the 16 and 84 posterior percentiles. We compare our results to the velocity dispersion profiles measured in other studies. The figure also shows measured velocity dispersion profiles from SDSS blue horizontal branch (BHB) stars (Xue et al. 2008, black dashed line); BHB and blue straggler (BS) stars from the Hypervelocity Star Survey (Brown et al. 2010; grey dashed line); and the SEGUE K-giant survey (Xue et al. 2016; connected black dots). The measured LOS velocity dispersions in the HALO7D fields using MSTO stars are consistent with other studies that have measured the LOS velocity dispersion profile over our distance range.

---

[4] In the COSMOS field, in the area where we used additional catalogs for selection, we used the Subaru B and V photometric measurements from Capak et al. (2007) and converted these to STMAG F606W as directed by Sirianni et al. (2005). We used the F814W magnitudes as published in the Leauthaud et al. (2007) catalog.

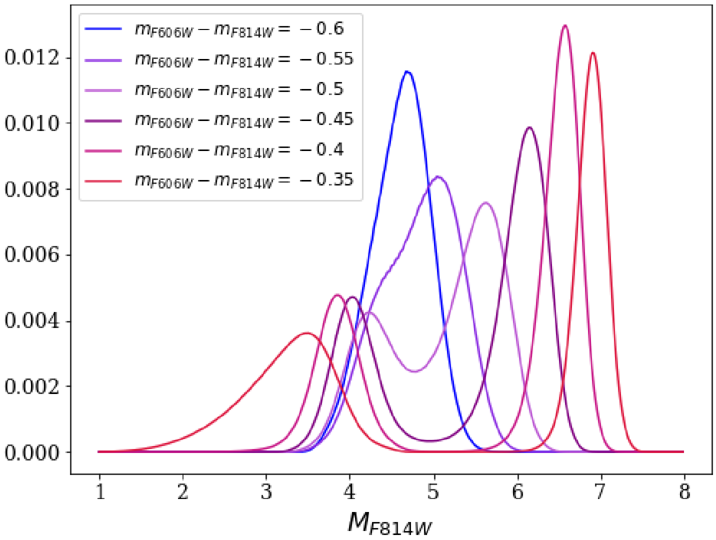

Figure 3.11: Normalized PDFs for absolute magnitude for six different choices of $m_{F606W} - m_{F814W}$ color. These PDFs are derived from the KDE constructed from the VandenBerg et al. (2006) isochrones, weighted by a Salpeter IMF and the approximate age and metallicity distributions of the halo.

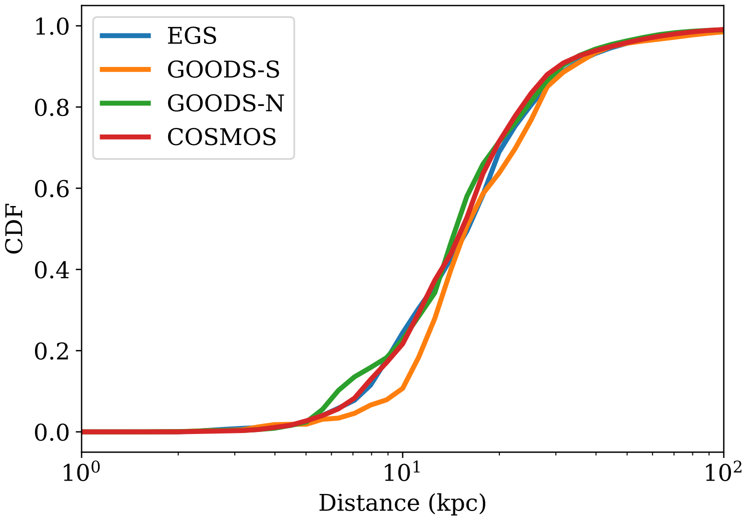

Figure 3.12: Cumulative distance distributions for the four HALO7D fields. Distance distributions are computed as given by Equation 3.13, using colors and assuming a MW stellar density profile (Deason et al. 2011). All fields have $\langle D \rangle \sim 20$ kpc.

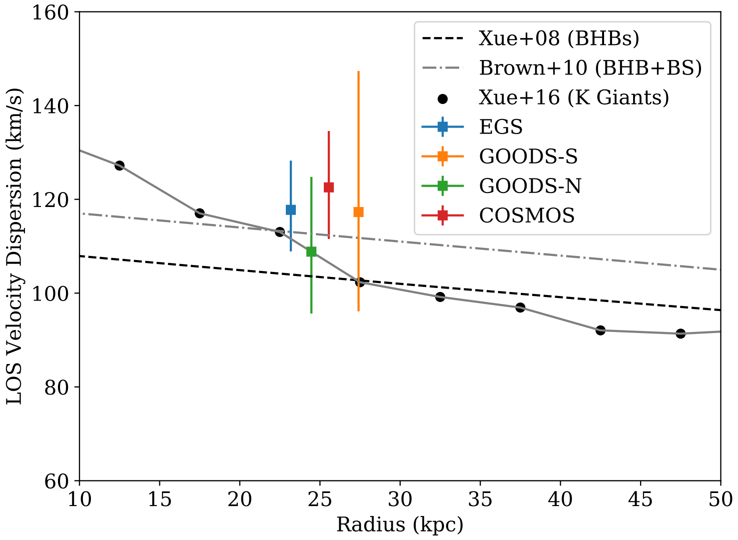

Figure 3.13: LOS velocity dispersions of the four HALO7D fields, plotted as a function of mean Galactocentric radius. Vertical errorbars show the 16-84% quantiles of the marginalized posterior. We compare our LOS dispersions with results from other studies: the black dashed line indicates best-fit LOS dispersion profile from Xue et al. (2008), measured from BHBs in SDSS. The grey dashed line indicates the best-fit dispersion profile from Brown et al. (2010) study, using BHB and BS stars as tracers. The black connected points show the resulting dispersion profile from the SEGUE K-giant survey Xue et al. 2016. The HALO7D dispersions are consistent with predictions from other tracers.

3.5 Comparison with Simulations

In all four lines of sight, we see the “hot halo” population; none of our fields appear to be obviously dominated by substructure that is cold in LOS velocity. In addition, the measured LOS velocity dispersions across the four fields are all consistent with one another, and are consistent with measurements made using other tracers. In this section, we interpret our LOS velocity distributions and the results from our velocity dispersion analysis in the context of simulations.

We perform mock HALO7D surveys on the eleven Bullock & Johnston (2005) halos (hereafter BJ05), using the synthetic survey software Galaxia (Sharma et al. 2011). The publicly available Bullock & Johnston simulations are high resolution N-body simulations of accreted dwarf galaxies onto a Milky Way-like parent galaxy. The parent galaxy has a time-dependent analytical potential with halo, disk and bulge components. Because there is no stellar disk in these simulations, there is no “in-situ” stellar halo component in these galaxies. Galaxia can accept N-body simulations as input, and it generates synthetic catalogs with smooth, continuous distributions of stars over any given volume.

We first use Galaxia to observe one square degree areas centered on the coordinates of the four HALO7D fields in all eleven BJ05 halos. We create catalogs of synthetic fields that are larger than our survey fields, because we are interested in exploring the underlying LOS velocity distributions along these lines of sight; the larger area provides us with more samples from these distributions. We then selected stars that fell within the HALO7D CMD selection boxes, in order to study the velocity distributions over the

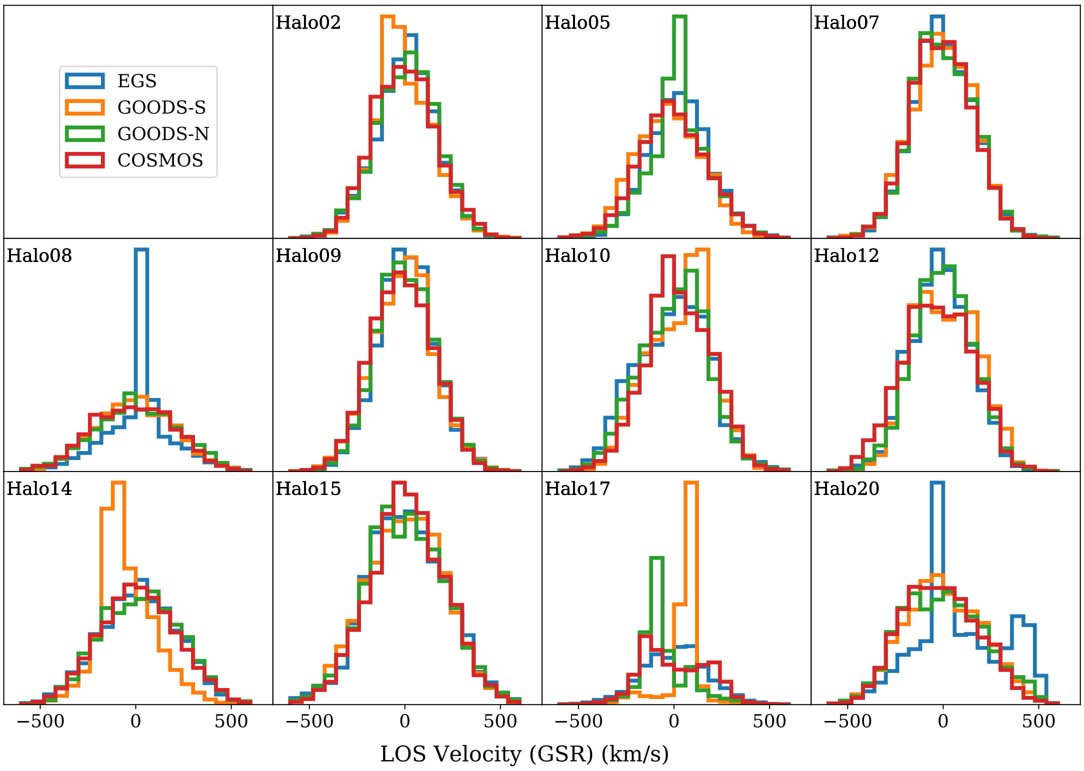

Figure 3.14: LOS velocity distributions from one square degree "observations" of the BJ05 accreted stellar halos generated with Galaxia. Different colored histograms denote the observations in the different HALO7D fields. Seven out of the eleven BJ05 halos show consistent velocity distributions across the four fields. Three halos show consistency across three fields with one field dominated by substructure. Halo17 shows four distinct LOS distributions across the four fields.

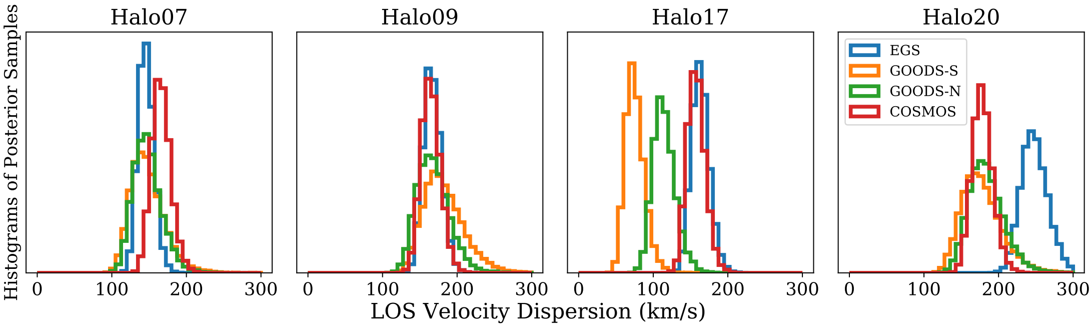

Figure 3.15: Examples of posterior samples from our dispersion analysis on our mock datasets. Each panel shows histograms of posterior samples for the 1D marginalized posterior distributions for the velocity dispersion (same as second panel of Figure 3.10) for a single mock HALO7D dataset. Each colored histogram represents a different HALO7D field. The first two panels show examples of consistent estimates of the LOS velocity dispersion for Halo07 and Halo09. The third panel shows an example of posterior estimates for Halo17: here, the cold peaks in GOODS-N and GOODS-S have caused the estimates of σ_{LOS} to be lower in these fields than in the other two fields. The fourth panel shows results from a mock dataset analysis for Halo20, in which the kinematically cold substructure at 400 km s^{-1}) and at 0 km s^{-1} result in a higher dispersion estimate in the EGS field.

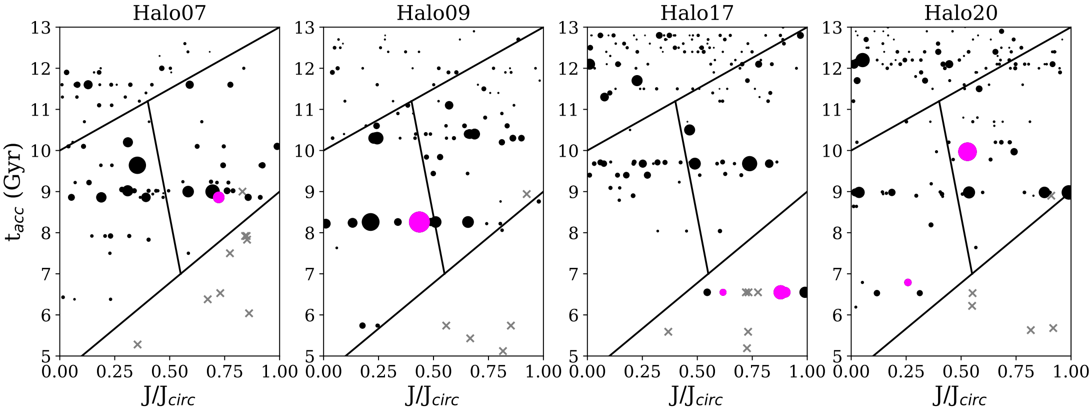

Figure 3.16: Age vs orbit circularity for the accretion events making up Halo07, Halo09, Halo17 and Halo20. Points are scaled by accretion event mass. Grey crosses indicate still-bound satellites. Colored points indicate the “dominant” satellites in the mock HALO7D samples. For Halo07 and Halo09, the same satellite is dominant across all four fields. In Halo17, one satellite dominates two fields while the other two fields are dominated by distinct satellites. For Halo20, three out of the four HALO7D fields are dominated by the most massive satellite, while the EGS field is dominated by a low mass, recent accretion event. The recent accretion events in Halo17 and Halo20 are responsible for the variable velocity distributions, while early, massive accretion results in consistent velocity distributions for Halo07 and Halo09.

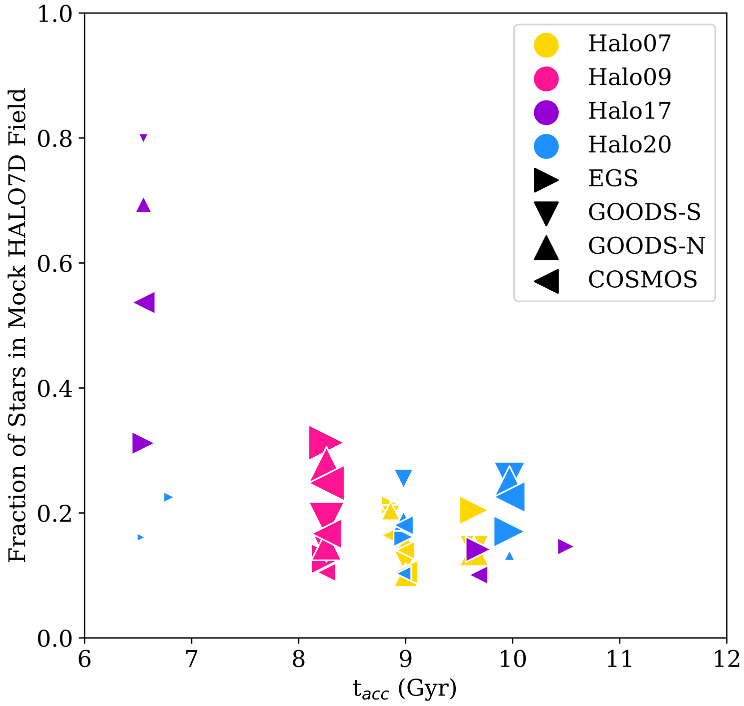

Figure 3.17: Fraction of stars contributed to a given line of sight as a function of accretion time, for Halo07 (gold), Halo09 (magenta), Halo17 (purple) and Halo20 (light blue). Point sizes are scaled by the mass of the accretion event. Only events that contribute $>10\%$ of the stars are shown. Different shape orientations denote different sightlines. Halo09's most massive satellite dominates the mock HALO7D sample along all sightlines, contributing 20-40% of the stars. The second most massive accreted satellite, which was accreted at the same time, also contributes 10-20% along all four sightlines. In Halo07, several relatively massive accretion events that were accreted around the same time all contribute between 10-20% of the stars in the four sightlines. These two accretion histories give rise to consistent velocity distributions along the different sightlines. In contrast, Halo17 has experienced recent accretion of low mass satellites. These recent accretion events dominate the four HALO7D sightlines, and create cold peaks in the LOS velocity distributions. In the case of Halo20, the most massive satellite contributes strongly to the samples across all four sightlines; however, as with Halo17, two lower mass, recent accretion events contribute $>15\%$ of the stars along the EGS sightline, creating two cold peaks in the velocity distribution.

same distance range as our survey.

Figure 3.14 shows the LOS velocity distributions for the four HALO7D fields in the eleven BJ05 halos. Each panel represents a different BJ05 halo, while each colored histogram shows the LOS velocity distribution for a given HALO7D field. In seven out of the eleven halos, the four LOS velocity distributions are all "hot", and appear to be consistent with one another. In three of the remaining halos, three of the fields have consistent LOS velocity distributions with one field having a strong cold peak (Halo08, Halo14, and Halo20). The only halo with four different velocity distributions across the four lines of sight is Halo17. Therefore, even if our survey had many more stars and no measurement uncertainties, based on the BJ05 velocity distributions, it would not be surprising to find that none of our four fields were dominated by kinematically cold substructure (especially given the fact that we expect more coherent substructure in the BJ05 halos than in the MW because of their assumption of a smoothly growing potential; see Section 2.5 of BJ05).

However, our survey is limited by small numbers (especially in GOODS-N and GOODS-S), and large measurement uncertainties at faint apparent magnitudes. Therefore, to make a valid comparison to our observations, we need to incorporate these effects. In order to assess how these different LOS velocity distributions respond to our observations and dispersion analysis procedure, we generate mock HALO7D datasets. We draw numbers of samples from the true underlying LOS velocity distributions corresponding to the sample sizes of our observations (e.g., we draw 88 stars from the EGS distribution, 25 stars from the GOODS-S distribution, etc.). We then assign velocity uncertainties to each star, based on a fit of the relation between apparent magnitude and LOS velocity error shown in Figure 3.7. We generate 30 sets of mock HALO7D observations for each BJ05 halo.

3.5.1 Dispersion Estimates

As can be seen in the middle panel of Figure 3.10, the posterior distributions for the LOS velocity dispersions in the four HALO7D fields substantially overlap, and the posterior median for each distribution lies within the 95% credible region of the other 3 distributions. For the purposes of this analysis, we define this result as having "consistent" estimates for the dispersion across all four fields. A specified posterior distribution is "consistent" with another posterior distribution if its posterior median lies within the 95% credible region of the other posterior distribution.

For each mock catalog, we evaluate the posterior distributions for the mean and dispersion along each sightline just as in Section 3.4.1 (though here we omit the disk model). Four examples of the resulting 1D marginalized posterior distributions for the velocity dispersions are shown in Figure 3.15. We find that the result of four consistent (as defined by our metric) posterior distributions for the velocity dispersions is relatively rare: in only 38 out of the 330 mock datasets did we find consistency across the four distributions. Out of 30 mock datasets per halo, the four distributions were consistent 5 times for Halo02, 7 times for Halo07, 3 times for Halo8, 6 times for Halo09, 2 times for Halo10, 8 times for Halo12, 7 times for Halo15. The first two panels of Figure 3.15 show examples of consistent velocity dispersion estimates (in Halo07 and Halo09). For Halo05, Halo14, Halo17, and Halo20, none of the mock HALO7D datasets yielded consistent dispersion estimates across all four fields. Example velocity dispersion posterior distributions for Halo17 and Halo20 are shown in the righthand panel of Figure 3.15.

Unsurprisingly, the halos that never achieved consistency across the four fields all contain cold peaks in their velocity distributions. It is important to note that these peaks do not necessarily have the same affects on the dispersion estimates. In the cases of Halo14 and Halo17, the cold peaks near 0 km s^{-1} decreased the dispersion estimates in these fields relative to those without strong peaks (as in the third panel of Figure 3.15). However, in the case of Halo20, which has a peak at 0 km s^{-1} as well as a peak at 400 km s^{-1} in its EGS field, the dispersion estimate in this field was consistently higher than the estimates in the other fields (as in the far righthand panel of Figure 3.15).

In spite of having a strong cold peak at 0 km s^{-1} in the EGS field, Halo08's posterior distributions were consistent for 3 out of the 30 mock datasets. This is because the stars responsible for the cold peak are all located at a far and common distance of $\sim$ 90 kpc; as a result, the velocity errors are sufficiently large for stars at these distances that the dispersion estimate is not strongly affected by the presence of the cold peak. In contrast, the two most dominant accretion events affecting Halo20's velocity distribution in EGS have mean distances of 16 kpc (responsible for the peak at 400 km s^{-1}) and 10 kpc (responsible for the peak at 0 km s^{-1}). With our large LOS velocity errors for stars at this distance, these stars do not strongly affect the resulting dispersion estimates. Based on our dispersion analysis alone, we cannot make claims about the presence or lack thereof of substructure in our sample, and leave the detection and characterization of potential substructure in HALO7D to future work.

3.5.2 Accretion Histories

We now look at four halos more closely: Halo07 and Halo09, which both achieved our metric for consistency numerous times, and Halo17 and Halo20, whose mock datasets never met our metric for consistency across the four fields. The accretion histories of these halos are shown in Figure 3.16, in the space of accretion time vs. $J/J_{\rm circ}$, with point sizes scaled by the mass of the accretion event. Lines indicate the regions of this plane dominated by the different morphological types discussed in Johnston et al. (2008); early, more radial accretion events tend to have mixed morphologies, whereas recent, circular accretion events tend to have great circle morphologies (see their Figure 3).

Accretion events that contribute the highest numbers of stars to one of the mock HALO7D fields are shown as pink points. In the case of Halo07 and Halo09, debris from the same accretion event contributes the most stars in all four of the HALO7D fields (hence only one accretion event in each halo is highlighted in pink). In contrast, in Halo17, debris from the same accretion event dominates the mock COSMOS and EGS distributions, but two different (and less massive) events contribute the most stars to the GOODS-N and GOODS-S fields. For Halo20, three out of the four HALO7D fields are all dominated by debris from the same, massive early accretion event; however, in the COSMOS field, the largest fraction of stars comes from a recent, low mass accretion event. The total fractions of stars contributed by accretion events that contribute $> 10\%$ of the stars along a given sightline are shown in Figure 3.17; point sizes are again scaled by mass, with the different triangle orientations indicating the different sightlines.

Neither Halo07 nor Halo09 have velocity distributions visibly dominated by kinematically cold substructure. In the case of Halo09, the most massive accretion event experienced by the halo ($m = 2.14 \times 10^{11} M_{\odot}$) contributes the most to the mock HALO7D samples along all four lines of sight (contributing $20 - 40$ %). The second most massive accretion event, accreted at the same time, also contributes $10 - 20\%$ of the stars in each line of sight (pink points in Figure 3.17). These massive, coincident accretion events are responsible for the hot velocity distributions that had consistent estimates of their dispersions. Debris from the most massive satellite accreted by the halo also dominated all four lines of sight in Halo02, Halo05, and Halo12.

Halo07 also experienced several coincident accretion events, though the events occurred earlier in its accretion history ($t_{\mathrm{acc}} \sim 9$ Gyr) and were on average lower mass than the events experienced by Halo09. The three most massive of these events ($m = 0.5 - 1.5 \times 10^{11} M_{\odot}$) each contribute $10-20\%$ of the stars in each of the four fields (gold points in Figure 3.17). We therefore see that accretion from several relatively early, coincident accretion events, even of lower mass, can also result in consistent estimates of the velocity dispersion across multiple lines of sight.

In contrast, Halo17's velocity distributions are not consistent across the four mock HALO7D fields. Figure 3.16 shows that Halo17's accretion history is characterized by several recent accretion events ($t_{\mathrm{acc}} \sim 6.5$ Gyr) on fairly circular orbits; Figure 3.17 shows that these recent accretion events strongly dominate the HALO7D sightlines, with the most massive recent event dominating two sightlines ($f_{\mathrm{sample}} \sim 30 - 50\%$) and the other two sightlines are strongly dominated by two different less massive events ($f_{\text{sample}} = 70 - 80\%$). The recent accretion experienced by this halo results in cold LOS velocity distributions along two of the four HALO7D sightlines. As a result, when we estimate the dispersion of mock observations in these fields, the cold peaks result in lower estimates of the dispersion in GOODS-N and GOODS-S compared to COSMOS and EGS.

In the case of Halo20, the most massive accretion event in its accretion history ($m = 1.65 \times 10^{11} M_\odot$, $t_{\text{acc}} \sim 10$ Gyr) contributes the most stars to the GOODS-N, GOODS-S and COSMOS fields. However, in Halo20's EGS field, two lower mass ($m \sim 1 - 2 \times 10^{10} M_\odot$), recent ($t_{\text{acc}} \sim 6.5$ Gyr) accretion events create distinct cold peaks in the LOS velocity distributions (represented by the small blue triangles in the lower left-hand corner of Figure 3.17).

To summarize, early, massive accretion events, as well as several, early, synchronous accretion events, give rise to consistent, hot velocity distributions along different halo sightlines, which lead to consistent estimates of the velocity dispersion across multiple lines of sight. Recent accretion can lead to sightlines dominated by kinematically cold substructure, even when these accretion events are low mass. Depending on the velocity of the substructure, as well as the number of accretion events causing peaks in the velocity distribution, this can lead to lower estimates of σ_{LOS} (as is in the case for Halo17) or higher estimates of σ_{LOS} (as for Halo20) along these sightlines.

While it is challenging to distinguish between the accretion histories of Halo09 and Halo07 with kinematics alone, we note that the chemical abundances will be different for these two scenarios. An accretion history like Halo09's should give rise to a higher average [Fe/H] than Halo07's accretion history, because of the mass-metallicity relation (e.g., Kirby et al. 2013).

In their analysis of the BJ05 halo density profiles, Deason et al. (2013a) found that halos with early, massive accretion events had breaks in their density profiles (like the density profile of the MW; e.g., Watkins et al. 2009, Deason et al. 2011, Sesar et al. 2011), whereas galaxies with prolonged accretion epochs had single power-law density profiles (like M31; e.g., Gilbert et al. 2012). Recent results from *Gaia* have discovered the remnant of an early, massive accretion event, known as the "Gaia-Sausage" (Belokurov et al. 2018) or "Gaia-Enceladus" (Helmi et al. 2018), which is both relatively metal rich and strongly radially biased in its orbital distribution. Deason et al. (2018b) find that the apocenters of these "Sausage" stars are at $r \sim 20$ kpc, coincident with the MW's break radius; this is also the approximate mean distance to our sample. Studying the LOS velocity distributions of the simulated BJ05 halos, we find that the consistent LOS velocity distributions of the HALO7D fields provides yet another piece of evidence that the MW likely experienced a massive, early accretion event. Proper motion information and abundances will help us to determine if our sample is dominated by Gaia-Sausage stars.

3.6 Conclusions

In this paper, we presented the spectroscopic component of the HALO7D survey; a Keck II/DEIMOS spectroscopic survey of distant, MSTO MW halo stars in the

CANDELS fields. We described the survey observing strategy, mask layouts, and target selection. We also presented a new method of measuring velocities from stellar spectra from multiple observations, utilizing Bayesian hierarchical modeling. We used the measured LOS velocities to estimate the parameters of the LOS velocity distributions in the four HALO7D fields.

We summarize our conclusions as follows:

1. When performing slit spectroscopy of point sources, it is essential to consider the apparent velocity shift due to slit miscentering when measuring velocities from individual spectra or when combining multiple spectroscopic observations. The hierarchical Bayesian approach presented in this work (implemented in VELOCIRAPTOR) allows for the parameters of individual observations to be modeled simultaneously, leveraging the available signal while properly propagating uncertainties.

2. All four HALO7D fields are dominated by the "hot" halo population, and have consistent LOS velocity distributions. The estimates of the velocity dispersions are consistent with estimates derived using other tracer populations.

3. We performed mock HALO7D observations using the synthetic survey software Galaxia to observe the Bullock & Johnston (2005) halos. We found that an early, massive accretion event, or several early events, can result in consistent estimates of the velocity dispersion along the different sightlines. This consistency in the velocity dispersion estimates arises because the same satellite (or the same few satellites) dominate the halo population along all sightlines. The consistent esti-

mates of σ_{LOS} in HALO7D therefore could indicate that the MW experienced an early, massive accretion event (or perhaps several events), whose stars are dominating the samples of all four fields. However, we emphasize that our dispersion analysis alone does not confirm or deny the presence of substructure: we intend to study substructure in HALO7D in future work when we can make use of PM and abundance information.

This paper is the first in the HALO7D series; our spectroscopy and the multi-epoch *HST* imaging will enable us to measure proper motions and abundances for these same stars. HALO7D is a deep complement to the *Gaia* mission: these stars in this dataset will be the faintest stars with measured 3D kinematics until LSST. With upcoming proper motions and abundances, we can continue to use the HALO7D dataset to improve our understanding of the Galaxy's formation.

Acknowledgments

Over the course of this work, E.C. was supported by a National Science Foundation Graduate Research Fellowship, as well as NSF grant AST-1616540. Partial support for this work was provided by NASA through grants for program AR-13272 from the Space Telescope Science Institute (STScI), which is operated by the Association of Universities for Research in Astronomy (AURA), Inc., under NASA contract NAS5-26555. E.C. expresses her profound gratitude to Alexander Rudy, a conversation with whom sparked the radial velocity measurement technique outlined in this work. E.C. also thanks Tony Sohn for his help with Figure 3.1. A.D. is supported by a Royal Society

University Research Fellowship. A.D. also acknowledges support from the STFC grant ST/P000451/1. P.G. and E.T. acknowledge support from the NSF grants AST-1010039 and AST-1412504. We thank the outstanding team at Keck Observatory for assisting us in our observations. This research made use of Astropy, a community-developed core Python package for Astronomy (Astropy Collaboration, 2013). This work has made use of the Rainbow Cosmological Surveys Database, which is operated by the Universidad Complutense de Madrid (UCM), partnered with the University of California Observatories at Santa Cruz (UCO/Lick,UCSC). We recognize and acknowledge the significant cultural role and reverence that the summit of Mauna Kea has always had within the indigenous Hawaiian community. We are most fortunate to have the opportunity to conduct observations from this mountain.

Chapter 4

HALO7D II: The Halo Velocity Ellipsoid and Velocity Anisotropy with Distant Main Sequence Stars

4.1 Introduction

The Milky Way (MW) stellar halo's kinematic structure contains key clues about the Galaxy's formation and mass assembly. According to the Lambda Cold Dark Matter (ΛCDM) paradigm for the evolution of the universe, the MW has built up its halo of dark matter over cosmic time by accreting smaller dark matter halos, some of which host dwarf galaxies. The remnants of these accreted dwarfs are found in the Milky

Way's stellar halo, and the velocities of these stars retain a link to their initial conditions because of their long dynamical times. The HALO7D project aims to investigate the MW's formation by studying the chemical and phase-space structure of the stellar halo's distant, main sequence (MS) stars.

One kinematic quantity that has long been of interest in MW formation studies is the velocity anisotropy β (Binney & Tremaine 2008), which provides a measure of the relative energy in tangential and radial orbits:

$$\beta = 1 - \frac{\langle v_\phi^2 \rangle + \langle v_\theta^2 \rangle}{2\langle v_r^2 \rangle}. \tag{4.1}$$

Systems with $\beta = 1$ are on completely radial orbits, while a population of stars on perfectly circular orbits has $\beta = -\infty$.

The velocity anisotropy parameter β plays a key role in the spherical Jeans (1915) equation:

$$M_{\text{Jeans}}(< r) = -\frac{r\sigma_r^2}{G}\left(\frac{\mathrm{d}\ln\rho}{\mathrm{d}\ln r} + \frac{\mathrm{d}\ln\sigma_r^2}{\mathrm{d}\ln r} + 2\beta\right). \tag{4.2}$$

Jeans modeling has been used to estimate the mass of the Galaxy in many studies (e.g., Dehnen et al. 2006, Gnedin et al. 2010, Watkins et al. 2009, Deason et al. 2012, Eadie et al. 2017, Sohn et al. 2018, Watkins et al. 2019 and references therein). However, estimates of the MW's mass have long been plagued by the mass-anisotropy degeneracy, owing to the lack of constraints on the tangential velocity distributions. It has only recently become possible to directly measure the tangential motion of kinematic tracers outside of the solar neighborhood. Previous studies have estimated β from line-of-sight (LOS) velocities alone (e.g., Sirko et al. 2004; Kafle et al. 2012; Deason et al. 2012, King et al. 2015), taking advantage of the fact that, because of our position within the Galaxy, the LOS velocity distribution contains information about the tangential velocity distributions. However, as pointed out by Hattori et al. (2017), studies of stars beyond $r \sim 15$ kpc with only LOS data (where $v_{LOS} \approx v_r$) result in systematic underestimates of β.

Fortunately, measuring tangential properties of tracers is now possible, thanks to the Hubble Space Telescope (*HST*) and the *Gaia* mission. The first estimate of β outside the solar neighborhood using directly measured 3D kinematics was presented by Cunningham et al. (2016), hereafter C16, using 13 MS stars with PMs measured from *HST* and radial velocities measured from Keck spectra. We found $\beta = -0.3^{+0.4}_{-0.9}$, consistent with isotropy and lower than solar neighborhood estimates, which find a radially biased $\beta \sim 0.5 - 0.7$ (Smith et al. 2009, Bond et al. 2010). However, the uncertainties on this measurement were substantial (primarily due to the small sample size), and in order to better constrain β and the MW mass, more tracers are required.

Studies have recently used the PMs of globular clusters (GCs) as kinematic tracers to estimate β and the mass of the MW. Sohn et al. (2018) used their own *HST* PM measurements of 16 GCs to find $\beta = 0.609^{+0.130}_{-0.229}$ in the Galactocentric distance range of $R_{\rm GC}$ = 10–40 kpc, and a corresponding MW virial mass of $M_{\rm MW,virial} = 2.05^{0.97}_{-0.79} \times 10^{12} M_{\odot}$. Watkins et al. (2019) used PM determinations of 34 GCs in the range $R_{\rm GC}$ = 2.0–21.1 kpc based on *Gaia* DR2 (Gaia Collaboration et al., 2018a) and found $\beta = 0.48^{+0.15}_{-0.20}$ consistent with Sohn et al. (2018), and a corresponding virial mass of $M_{\mathrm{MW,virial}} = 1.41^{1.99}_{-0.52} \times 10^{12} M_{\odot}$.

While studies have sought to estimate a single value β in order to estimate the mass of the MW, studies of β can have additional power in constraining the MW's accretion history. For example, the anisotropy radial profile $\beta(r)$ can contain information about the Galaxy's assembly history. In Deason et al. (2013b) and C16, we argued that our isotropic measurement of β, which is lower than solar neighborhood measurements and also distant halo estimates (Deason et al. 2012), indicates a "dip" in the β profile, and that this dip could indicate the presence of a shell.

Loebman et al. (2018) provided theoretical perspective on this question, by studying the β profiles in three suites of simulations, including accretion-only and cosmological hydrodynamic simulations. They found that both types of simulations predict radially biased $\langle\beta\rangle \sim 0.7$ beyond 10 kpc. Only one of the 17 simulations studied had tangentially biased β over a large range of radii at $z = 0$; this extended β dip was the result of a major merger at $z \sim 1$. While the other 16 simulations had radially biased β at $z = 0$, Loebman et al. (2018) found that temporal dips in the β profile could arise. They found that recently accreted material can result in short-lived dips in β, while the passage of a massive satellite can induce a longer-lived dip in the β profile from the in-situ component of the stellar halo. This latter scenario could explain the observed "dip" along the line of sight towards M31, as recent studies of the Triangulum Andromeda overdensity have suggested that its origin may be the disk rather than an accreted satellite (Price-Whelan et al. 2015; Bergemann et al. 2018), and that the event that disturbed the orbits of these disk stars may be the passage of the Sagittarius dwarf (Laporte et al. 2018).

The anisotropy variation across different subpopulations in the halo can also be used to disentangle accretion events. Using 7D measurements from the *Gaia* DR1 and SDSS of local MS stars, Belokurov et al. (2018) found that the relatively metal-rich stars ([Fe/H] > -1.7) show strongly radially biased velocity anisotropy (i.e., "sausage" stars, named thus because of the elongated radial velocity distribution relative to the tangential velocity distribution), while the metal-poor stars display an isotropic velocity distribution. They argue that presence of this radially biased, relatively metal-rich population in the inner halo indicates that the MW experienced a relatively massive, early accretion event. Evidence for this scenario has been bolstered with results from *Gaia* DR2 (Helmi et al. 2018, Deason et al. 2018b). Lancaster et al. (2019) showed that the kinematics of the BHBs in *Gaia* DR2 can be modeled by a mixture of two populations: one strongly radially biased and one isotropic. Debris from a massive, radialized dwarf that dominates the inner halo, known as the *Gaia*-Sausage, *Gaia*-Enceladus, or Kraken, is speculated to be responsible for this signature.

Thanks to the *Gaia* mission, it is now possible to estimate the β of stars in the MW; however, even with *Gaia* DR2, uncertainties remain substantial at large radii, and, even in the the final data release, *Gaia* will provide PMs only for stars brighter than $G \sim 20$. As a result, *Gaia* will only provide PMs for MS stars out to $D \sim 15$ kpc in the halo. Beyond $D \sim 15$ kpc, studies of tangential motion of the stellar halo using *Gaia* PMs will be limited to giants and evolved stars (e.g., Bird et al. 2018, Lancaster et al.

2019). While giants make excellent tracers due to their bright apparent magnitudes, it is impossible to uniformly select giants from all age and metallicity populations in the halo. Giants are also rare; averaging over large areas of the sky (and thus potential inhomogeneities in the halo) is often required when estimating halo properties with giants.

The HALO7D project seeks to complement the *Gaia* mission by measuring 3D kinematics of distant MW halo MS stars. HALO7D includes both Keck spectroscopy and *HST* PMs for MW halo star candidates in the magnitude range $19 < m_{F606W} < 24.5$. This dataset provides a deep, densely sampled view of the garden variety stars of the MW halo. In the first HALO7D paper (Cunningham et al. 2019; hereafter Paper I), we presented the spectroscopic component of the HALO7D dataset. In this paper, the second in the HALO7D series, we introduce the proper motion component of HALO7D, and use our full 3D kinematic sample to study the halo velocity ellipsoid and anisotropy.

In this work, we seek to use the HALO7D dataset to estimate the parameters of the velocity ellipsoid, and velocity anisotropy, of distant halo MS stars. This paper is organized as follows. In Section 4.2, we describe the HALO7D dataset and present the HALO7D PM samples. In Section 4.3, we describe our methodology for estimating the halo velocity ellipsoid parameters from our observables. In Section 4.4, we present our resulting posterior distributions for ellipsoid parameters and velocity anisotropy. In Section 4.5, we compare our results to previous work and other studies. In Section 4.6, we investigate the spatial and radial variation of β for two halos from the *Latte* suite of simulations. We conclude in Section 4.7. Details on our computational method for deriving PM uncertainties are given in Appendix B.1; a description of how we tested our ellipsoid parameter model with fake data is given in Appendix C.

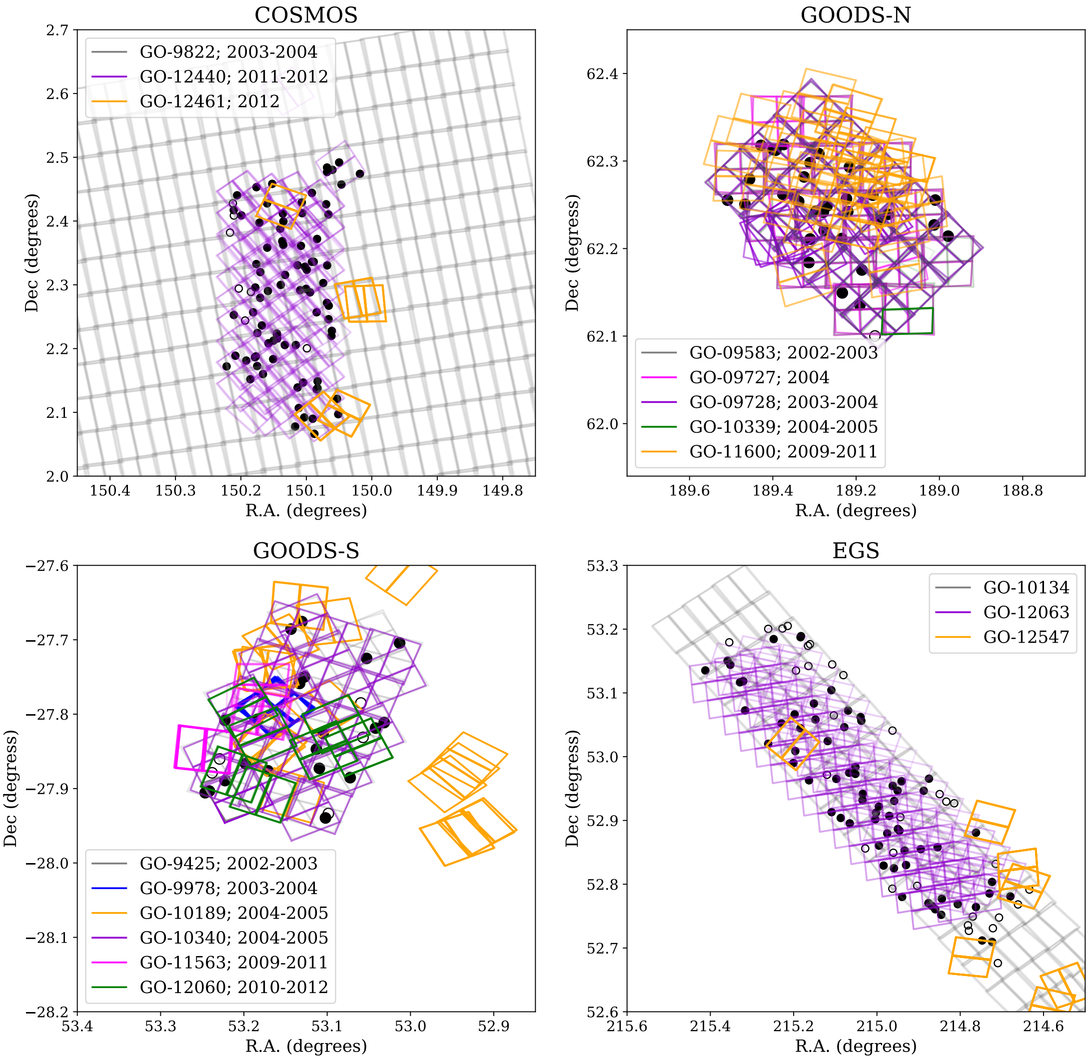

Figure 4.1: The multi-epoch *HST*/ACS footprints of the four HALO7D fields. Different colors indicate the positions of each ACS chip in the different *HST* programs used to measure PMs in this work. HALO7D spectroscopic targets are indicated by black points; filled points indicate targets for which we successfully measured a PM, whereas empty circles indicate targets for which we could not measure a PM.

4.2 Dataset

HALO7D consists of Keck/DEIMOS spectroscopy and *HST* measured PMs of MW MSTO stars in the EGS, COSMOS, GOODS-N and GOODS-S fields. Coordinates of the HALO7D fields are listed in Table 4.1. We begin this section by summarizing some of the key details on target selection, survey properties, and radial velocity measurements that are discussed in detail in Paper I; the remainder of this section is devoted to a discussion of the proper motion measurements.

4.2.1 Keck/DEIMOS Spectroscopy

The HALO7D spectroscopic program was described in detail in Paper, I, but we summarize the key details here.

Candidate halo stars were identified from color-magnitude diagrams. To minimize disk contamination, we selected blue, faint ($19 < m_{F606W} < 24.5$) objects with star-like morphologies. Stars were observed with Keck II/DEIMOS, configured with the 600ZD grating centered at 7200 Å, beginning in April 2014 with the final observations taking place in April 2017. We targeted each DEIMOS mask for a minimum of 8 hours of total integration time, and up to 24 hours.

The radial velocities for these stars were measured using a new Bayesian hierarchical method, called VELOCIRAPTOR. In order to build up sufficient signal to noise on our targets, stars were observed many times, sometimes over the course of years. Different observations of the same star will have different raw velocities; this is due to the motion of the Earth around the sun (the heliocentric correction) as well as offsets in wavelength solution due to slit miscentering (the A-band correction). We used a Bayesian hierarchical model in order to combine these different observations into a single estimate of the star's velocity. For further details on this technique, we refer the reader to Paper I.

4.2.2 HST Proper Motions

The measurement methodology in this work builds from existing *HST* PM measurement techniques, relying on the key concept that distant galaxies can be used to construct an absolute stationary reference frame (e.g., Mahmud & Anderson 2008). Sohn et al. (2012, 2013, 2017) present detailed descriptions of the state-of-the-art PM measurement techniques used to measure the PMs of Local Group systems with *HST* data. These techniques have been used to measure the PMs of M31 (Sohn et al. 2012), dwarf galaxies Leo I (Sohn et al. 2013), Draco and Sculptor (Sohn et al. 2017); MW GCs (Sohn et al. 2018); and several MW streams (Sohn et al. 2016). The PMs of individual MW halo stars measured with *HST* were first published by Deason et al. (2013b); subsequently, the PMs of individual stars belonging to MW streams were published by Sohn et al. (2015) and Sohn et al. (2016).

However, the previous Sohn PM studies have typically used only a few *HST* pointings in each study; in that work, they were able to carefully select galaxies by eye that are suitable for use in the reference frame. In order to measure PMs over the full area of the CANDELS fields, we required an approach that could identify "good" galaxies (with well-measured positions) and "bad" galaxies (with poorly measured positions) without relying on visual inspection. We therefore built upon existing PM

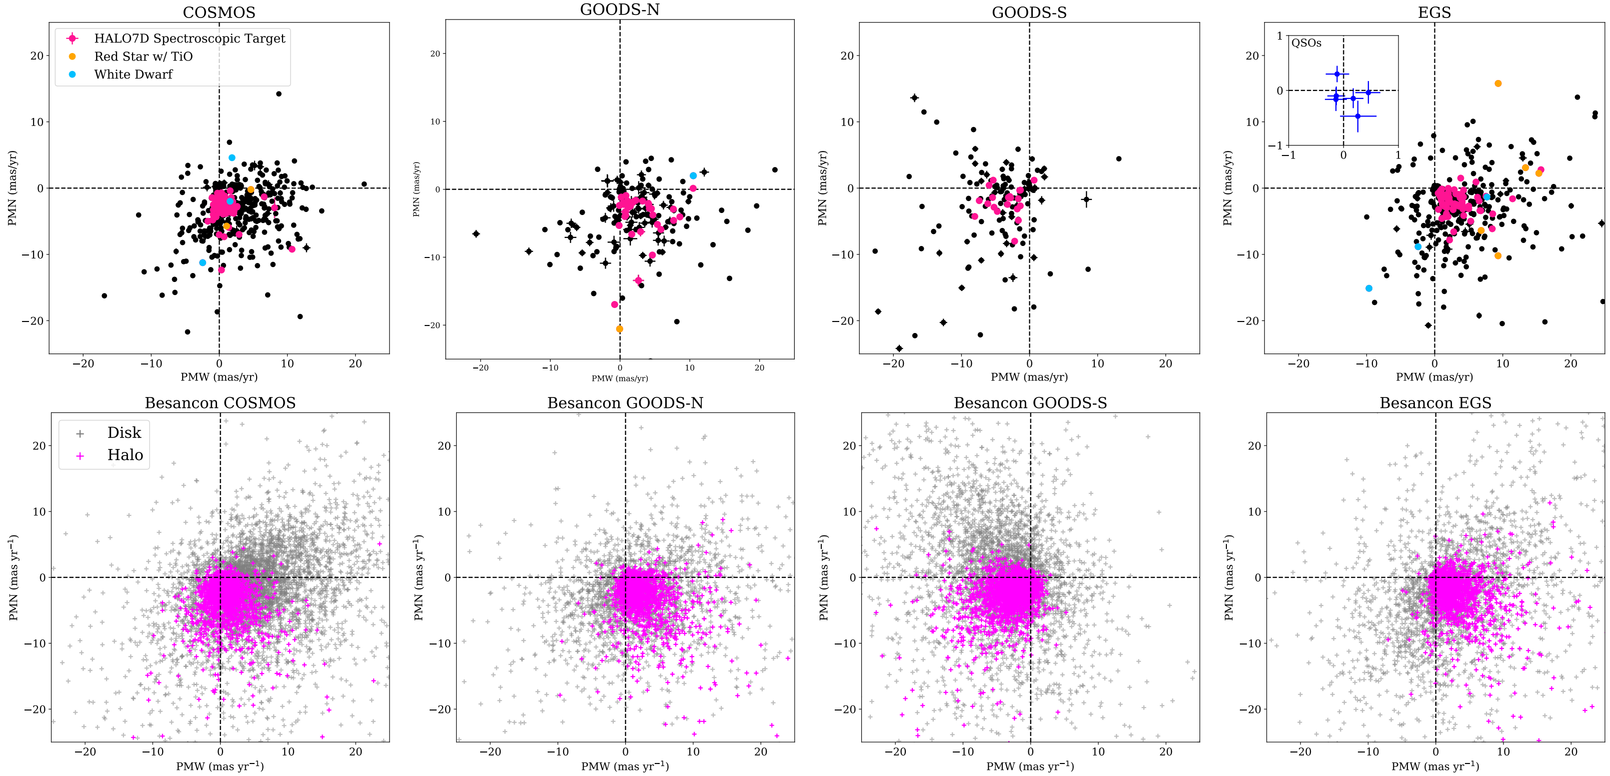

Figure 4.2: Top panels: Proper motion diagrams of the four HALO7D fields. Black points indicate stars that were in the *HST* images that were not HALO7D spectroscopic targets — primarily MW disk stars. HALO7D spectroscopic targets are shown in pink. Light blue points indicate spectroscopically confirmed WDs, while red MW disk star contaminants are shown in orange. In the EGS PM diagram, the inset shows PMs and errorbars for six spectroscopically confirmed quasars observed in the EGS field. Lower panels: PM diagrams from the Besançon Galaxy Model, from 1 square degree fields centered on the coordinates of the HALO7D fields. Grey crosses indicate Besançon disk stars, whereas magenta crosses are Besançon halo stars.

techniques in this work, implementing a Bayesian mixture model that identifies "good" and "bad" galaxies probabilistically and incorporates this uncertainty into the ultimate measurement of the PMs of the stars in the set of images.

Measuring Proper Motions

In order to measure PMs for the HALO7D targets, we first had to identify the *HST* programs and filters to use for PM measurements. The GOODS, COSMOS, and EGS fields have all been observed multiple times with various setups (detectors + filters). Among them, we selected data that provide astrometric quality sufficient for measuring absolute PMs of individual halo stars. Specifically, data used for our PM measurements meet the following conditions: (1) observations must be obtained with either ACS/WFC or WFC3/UVIS; (2) observations must be in one of the broad-band filters F606W, F775W, F814W, or F850LP; (3) time baseline of the multiple epochs must be at least 2 yr; (4) combined exposure time in the shallower epoch must be at least one orbit long; and (5) individual exposure time must be at least 300 s long. The details of the *HST* programs used for the PM measurements are listed in Table 4.2; the footprints from these programs are shown in Figure 4.1.

Once the *HST* programs and filters were chosen, the `*_flc.fits` images were downloaded from MAST. These images are corrected for imperfect charge transfer efficiency using the algorithms described in Anderson & Bedin (2010). The `*_flc.fits` images are processed by a custom-made FORTRAN routine called `flt2xym4rd`, which takes a list of RA, Dec positions for objects, identifies them in an `flc` image, and measures them with a library PSF (see Anderson & King 2006, AK06), determining for each a position, flux, and stellarity index. The routine then uses the WCS header of each exposure and the distortion solution in AK06 to convert the source positions into an RA-Dec frame. This routine is run on all the exposures that cover a particular field.

In this analysis, we measure the PMs on a star-by-star basis. For every target star, the first step is to identify all images that contain the star of interest. The single-exposure catalogs from the `flt2xym4rd` output are then fed into another custom-made routine, `xid2mat`, which takes the single-exposure catalogs in pairs and transforms one catalog into the frame of the other, using the galaxy positions as the basis for the transformation. This transformation makes an initial assessment of which galaxies have consistent positions between the two frames, though the ultimate weighting of the galaxies is done in a Bayesian fashion.

We then specify one image as the reference image: the reference image has the maximum amount of overlap with the other images across epochs containing the star of interest. All overlapping images are mapped onto the reference image frame with `xid2mat` using a six-parameter linear transformation:

$$\begin{pmatrix} A & B & x_t \\ C & D & y_t \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} u \\ v \\ 1 \end{pmatrix} = \begin{pmatrix} u_{ref} \\ v_{ref} \\ 1 \end{pmatrix}, \qquad (4.3)$$

where (u, v) are the vectors of distortion corrected positions of objects in one image and (u_{ref}, v_{ref}) are the vectors of positions in the reference image. The parameters x_t, y_t represent any linear translation offset between the two images, while parameters A, B, C, D incorporate scale, rotation, and off-axis linear camera distortion terms. The positions of stars are used to match frames within an epoch, and the positions of “good” galaxies are used to match images across epochs. For more details on why these transformations are required in comparing *HST* images, please see section 3.6.4 in Anderson & van der Marel (2010).

When the images across epochs have been matched via the linear transformation, the change in the positions of the stars across epochs provides an initial estimate of their PMs. In order to get full posterior probability distributions for the PMs, and incorporate all sources of uncertainty (such as uncertainty in star and galaxy positions, as well as which galaxies should be including in the stationary reference frame), we use a Bayesian mixture modeling approach. We leave as free parameters the positions of all stars and galaxies, the image transformation parameters, and the proper motions of all stars. We model the galaxies in our reference frame as being a mixture of “good” and “bad” galaxies (with poorly measured positions). Within an epoch, we use the positions of stars to precisely align the images.

Table 4.1 lists the resulting median PM errors in each of the HALO7D fields. Our PM errors are not a function of the magnitudes of our stars, but rather our ability to define the stationary reference frame for a given target. This is determined by how many images there are containing a given star, how much these images overlap across epochs, and how many “good” galaxies there are in the images. For a full description of the Bayesian model for this problem, as well as the details of the Gibbs sampling algorithm we used to sample from the full posterior, we refer the reader to Appendix B.1.

Field	l (deg)	b (deg)	$v_{l,\odot}$ (km s^{-1})	$v_{b,\odot}$ (km s^{-1})	Median PM Error (mas yr^{-1})
COSMOS	236.8	42.1	−126.0	148.2	0.15
GOODS-N	125.9	54.8	−153.8	−154.2	0.12
GOODS-S	223.6	−54.4	−171.5	−140.8	0.20
EGS	96.4	60.4	−38.7	−209.0	0.18

Table 4.1: Galactic coordinates, projection of the Sun's velocity in Galactic coordinates, and the median PM error (in Galactic coordinates) for the four HALO7D fields. Quoted median PM errors are the errors in a single component (e.g., $\mu_l \cos(b)$ or μ_b; we find both components of PM have the same median errorbars, to within 0.005 mas yr^{-1}, within a given field).

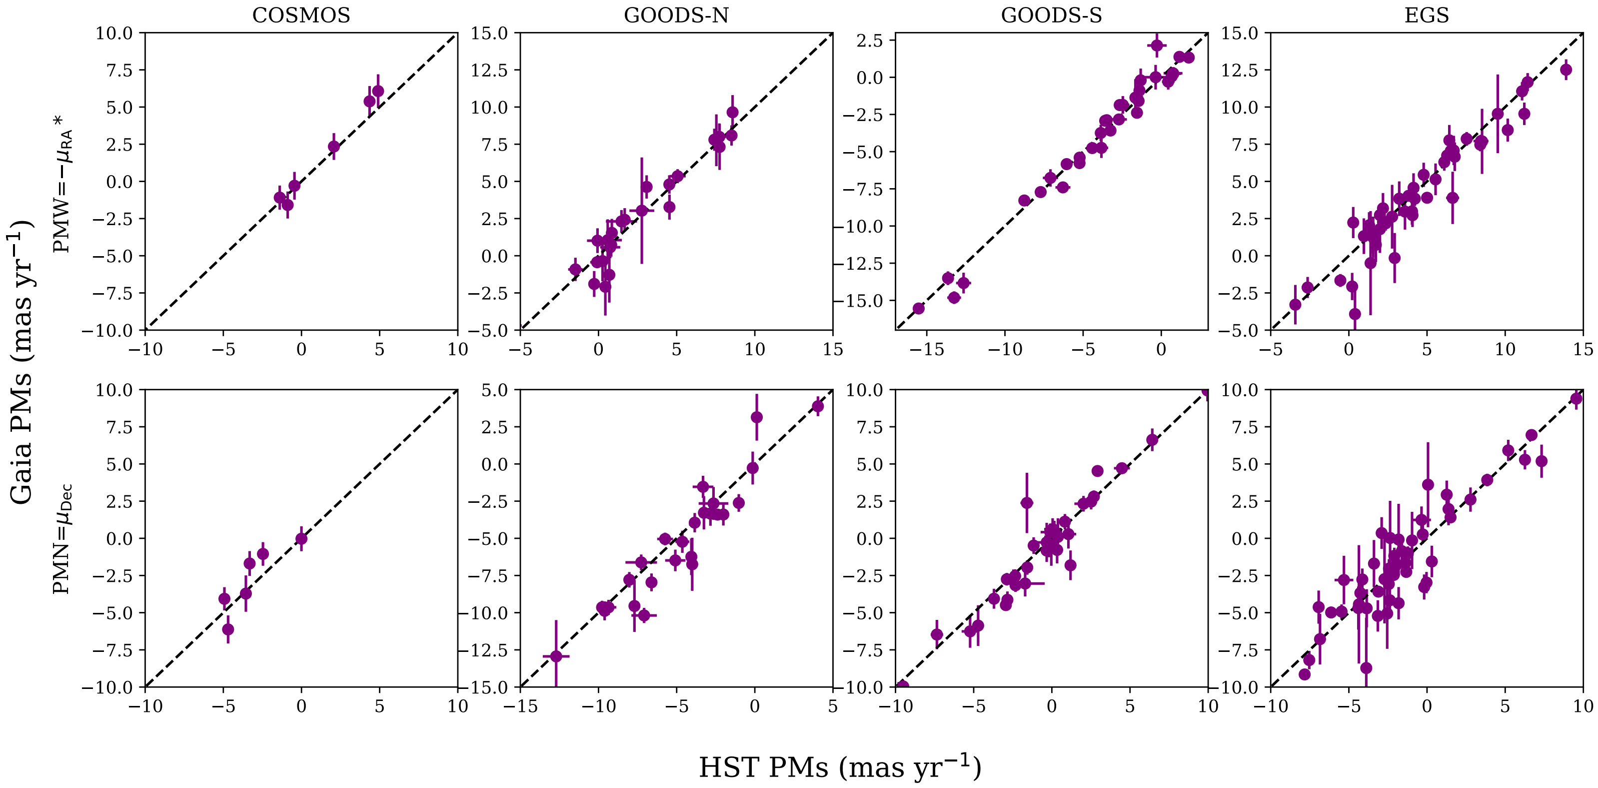

Figure 4.3: Comparison of the *HST* measured PMs for sources brighter than $v \sim 21$ with their *Gaia* counterparts, if available. Top panels show the comparison of PMW $= -\mu^*_{\mathrm{RA}}$; lower panels show the comparison of PMN $= \mu_{Dec}$. The two datasets show excellent agreement, with *HST* PMs generally having lower errorbars. For clarity, figures are zoomed in to the regions of highest target density; there are several high |PM| sources that are beyond the range of the figures (2/27 stars in GOODS-N; 6 high |PMW| and 4 high |PMN| out of 39 matched sources in GOODS-S; and 7 high |PMW| and 4 high |PMN| out of 56 sources in EGS). In addition, the number of targets that overlap across the two datasets is not necessarily a function of the area of the field, because of *Gaia*'s scanning law; for example, there are only 6 sources with reported PMs in COSMOS, while EGS, GOODS-S and GOODS-N have 56, 39, and 27 matches, respectively.

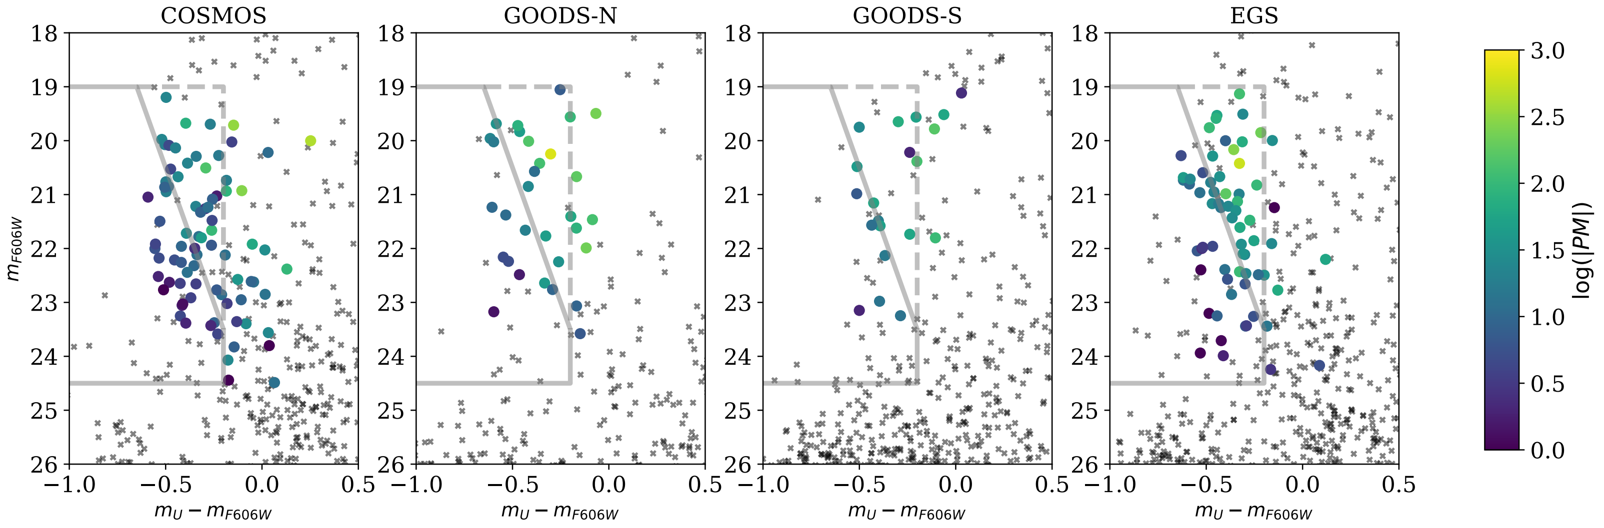

Figure 4.4: CMDs of the four HALO7D fields, in the STMAG bands F606W and F814W. Gray lines indicate the HALO7D selection boxes for these bandpasses; see Paper I for more details on target selection. Stars used in this analysis (main sequence stars with successful LOS velocity and PM measurements) are plotted as colored circles, with the colors corresponding to the logarithm of the magnitude of their proper motions ($|PM|^2 = (\mu_l \cos b)^2 + \mu_b^2$). For the CMDs over the full color range of these fields, please see Figure 3 of Paper I.

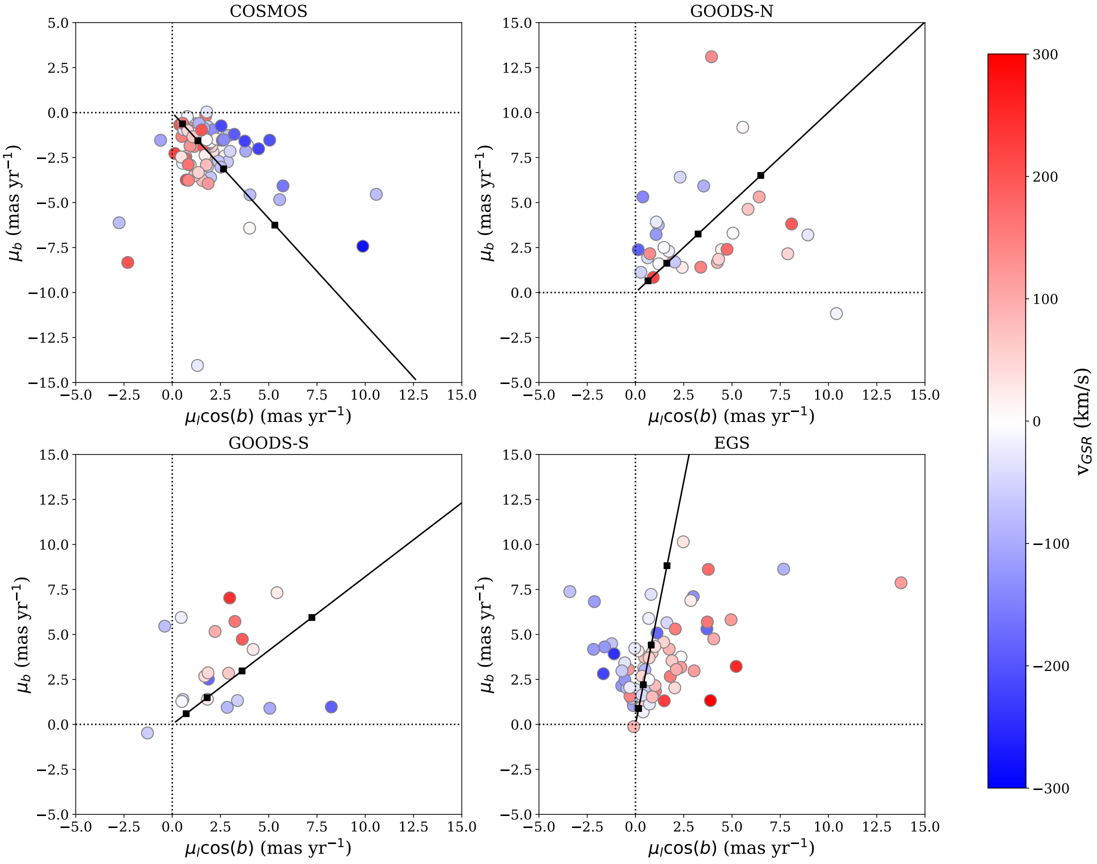

Figure 4.5: Proper motion diagrams of the four HALO7D fields, in Galactic coordinates, color coded by LOS velocity in the Galactocentric frame. Solid black lines indicate the solar reflex motion along each sightline; the squares indicate the implied mean PM along each line of sight for $D = 5, 10, 20, 50$ kpc (with mean PM at larger distances approaching $(\mu_l \cos(b),\, \mu_b) = (0, 0)$. Dotted lines indicate $\mu_l \cos(b) = 0$ mas yr^{-1}, $\mu_b = 0$ mas yr^{-1}.

Proper Motion Diagrams

Figure 4.1 shows positions of the HALO7D spectroscopic sample as black points; filled circles indicate targets for which we successfully measured PMs, and open circles are stars for which we couldn't measure a PM. As can be seen in Figure 4.1, the *HST* pointings from different epochs are not well aligned; this is because we are using archival data for *HST* programs that were not designed with astrometry in mind. As a result, some of the HALO7D targets only have one epoch of *HST* imaging. This usually arises when the target is on the edge of the field, or if the target falls in the ACS chip gap in one of the epochs.

PM diagrams for the four HALO7D fields are shown in the top panels of Figure 4.2. PMs are plotted in PMW $= -\mu_\alpha \cos(\delta)$, PMN $= \mu_\delta$. The PMs of HALO7D halo star candidates are shown in pink. Our PM method returns PMs and uncertainties for all point-like objects in the specified reference image that have multi-epoch coverage; PMs for objects that were not HALO7D spectroscopic targets are shown as black points. Most of these points are MW disk stars, though a few will be point-like distant galaxies. As explained in Paper I, our spectroscopically confirmed disk contaminants are white dwarfs (WDs) and red stars with titanium oxide absorption features. These disk contaminants are shown as light blue and orange points, respectively. The black, light blue and orange points occupy a larger area of PM space than the pink points; because they are mostly disk members, they are at closer distances than the HALO7D halo star candidates and thus have higher proper motions. The PM diagrams as predicted by the Besançon Galaxy Model (Robin et al. 2003), for one square degree fields centered on our field coordinates, are shown in the lower panels of Figure 4.2 for reference.

In addition, in the EGS field, we had six spectroscopically confirmed quasars for which we could also measure PMs. These PMs are shown in the inset of the upper lefthand panel of Figure 4.2; reassuringly, all quasar PMs are consistent with 0 mas yr^{-1}.

As an additional verification of our technique, we compare our measured PMs with those reported in the second data release of the *Gaia* mission (Gaia Collaboration et al. 2016, 2018b). The bright end of our sample corresponds to the faint end of the *Gaia* sample; though most of the halo star candidates used for dynamical modeling in this work do not have reported *Gaia* PMs, using the full catalog of stars (including stars not targeted for spectroscopy), we find a sizable sample of stars with reported PMs in both datasets.

Figure 4.3 shows the comparison of the *HST* PMs against the *Gaia* PMs, when available, for sources brighter than $v \sim 21$. Figures are zoomed in to highlight the regions of PM space where there is the highest target density in each field; there are more sources with measured PMs in both catalogs that lie outside the plotted range. We note that the number of sources available for comparison is not necessarily proportional to the area of the field, due to *Gaia* survey incompleteness at the faint end as a result of the *Gaia* scanning law (hence the few overlapping targets in COSMOS). The two datasets show excellent agreement, with the *Gaia* measurements generally having larger errorbars. We leave a more detailed comparison of the two datasets (including an exploration of their respective systematics) to future work.

Field	Program	P.I.	Filter	Dates
COSMOS	GO-9822	Scoville	F814W	12/2003–05/2004
	GO-12440	Faber	F814W	12/2011–02/2012
	GO-12461	Riess	F814W	02/2012 – 04/2012
GOODS-N	GO-9583	Giavalisco	F775W	11/2002–05/2003
	GO-9727	Perlmutter	F775W	04/2004–08/2004
	GO-9728	Riess	F775W	06/2003–09/2004
	GO-10339	Riess	F775W	10/2004–04/2005
	GO-11600	Weiner	F775W	9/2009–04/2011
GOODS-S	GO-9425	Giavalisco	F606W, F850LP	07/2002–02/2003
	GO-9978	Beckwith	F606W, F850LP	09/2003 – 01/2004
	GO-10189	Riess	F606W, F850LP	09/2004 – 08/2005
	GO-10340	Riess	F606W, F850LP	07/2004 – 09/2005
	GO-11563	Illingworth	F606W, F850LP	08/2009 – 02/2011
	GO-12060/1/2	Faber	F606W, F850LP	08/2010 – 02/2012
EGS	GO-10134	Davis	F814W	06/2004 – 03/2005
	GO-12063	Faber	F814W	04/2011 – 05/2013
	GO-12547	Cooper	F814W	10/2011 – 02/2013

Table 4.2: Summary of the *HST* programs used for the PM measurements in this paper.

Color magnitude diagrams (CMDs) for the four HALO7D fields are shown in Figure 4.4. Gray crosses indicate all star candidates, and HALO7D targets with successful PM measurements and LOS velocity measurements are shown as circles color coded by the magnitude of their PMs ($|PM|^2 = (\mu_l \cos b)^2 + \mu_b^2$). As expected, bluer, fainter stars tend to have lower PMs (as they are more distant), whereas brighter, redder stars have higher PMs (and are more nearby). For the reader interested in the full color range of the CMD, we refer them to Figure 3 of Paper I.

Our 3D kinematic sample is summarized in Figure 4.5; PM components in (l, b) are plotted against each other for the four HALO7D fields, color coded by LOS velocity as measured in Paper I. In EGS, we see an interesting covariance between μ_l and v_{LOS}; there appears to be a trend of increasing LOS velocity with increasing μ_l. However, we note that this is not a signature of rotation. EGS is located at a Galactic longitude $l = 96$ degrees; along this line of sight, $v_l \approx -V_X$, and $V_Y = v_{LOS}\cos(b) - v_b \sin(b)$. Given that $v_\phi = \frac{x}{R_p}V_Y - \frac{y}{R_p}V_X$, a covariance between V_X, V_Y arises naturally if we assume a Gaussian velocity distribution for v_ϕ. As we'll see in Section 4.4, the fact that v_{LOS} and μ_l increase together is consistent with zero net rotation along this line of sight.

4.3 Modeling The Halo Velocity Ellipsoid

We use our 3D kinematic sample to estimate the parameters of the halo velocity ellipsoid in spherical coordinates. In this work, we use only objects for which we have both a successful PM measurement and a successful LOS velocity measurement; we leave the analysis of stars with PM measurements but without LOS velocities to future work. Our method is very similar to the ones used in Cunningham et al. (2016) and Deason et al. (2013b), though in this work we have used notation and language consistent with a Bayesian construction of the problem.

For each star i located in field k, we have data $y_i = \{v_{LOS}, \mu_l, \mu_b\}$, with associated explanatory variables $x_i = \{m_{F606W,i}, m_{F814W,i}, l_k, b_k\}$. We model our sample as being drawn from a mixture of two distributions: the disk distribution (with fixed parameters), and the halo distribution.

The free parameters in our model are the absolute magnitudes (and, by extension, the distances) to each star $\mathbf{M} = \{M_{F814W,1}, ..., M_{F814W,N}\}$ (we denote the corresponding distances $\mathbf{D} = \{D_1, ..., D_N\}$); the fraction of disk contamination along a given line-of-sight $\mathbf{f} = \{f_{\mathrm{Disk},1}, ..., f_{\mathrm{Disk},k}\}$; and the halo velocity ellipsoid parameters $\theta_{\mathrm{Halo}} = \{\langle v_\phi \rangle, \sigma_r, \sigma_\phi, \sigma_\theta\}$.

4.3.1 Disk Model

For the disk model, we work in cylindrical coordinates (R_p, ϕ, z). We assume exponential density profiles in both R_p and z, with a disk scale length of $h_R = 3$ kpc and a disk scale height of $h_z = 1$ kpc.

For the disk velocity distributions, we assume distributions in R_p and z that are Gaussian with zero net motion, and have dispersions of $\sigma_{R_P} = 45$ km s^{-1} and $\sigma_z = 20$ km s^{-1}. For the tangential component, we assume that the rotational velocities are described by a skewed normal distribution with mean $\langle v_T \rangle = 242$ km s^{-1}, scale parameter 46.2 km s^{-1}, and shape parameter of -2. These parameters are derived based on the predicted marginalized velocity distributions from `galpy`[5] (Bovy 2015), using the quasi-isothermal distribution function discussed in Binney (2010) and Binney & McMillan (2011) and the `MWPotential2014` (see Bovy 2015 for details). While they are not free parameters in our model, for simplicity in notation, we denote the disk DF parameters as θ_{Disk}.

While this disk model is quite simple, we find that this model is effective at identifying stars in our sample that are disk-like (see Section 4.4.3). These stars are more likely to have higher proper motions, brighter apparent magnitudes, redder colors, and heliocentric LOS velocities closer to 0 km s^{-1}.

---

[5] http://github.com/jobovy/galpy

4.3.2 Halo Model

For the halo distribution, we work in spherical coordinates. We assume the broken halo density profile derived in Deason et al. (2011), with break radius $r_b = 27$ kpc and slopes $\alpha_{\text{in}} = 2.3$, and $\alpha_{\text{out}} = 4.6$. The probability that a star has a distance D_i given the density profile is given by:

$$p(D_i|\rho, l, b) \propto \rho(r_q(D_i, l, b)) \times D_i^2, \tag{4.4}$$

where the factor of D_i^2 arises from the spatial volume element in spherical coordinates.

We assume independent Gaussian velocity distributions for the three spherical components of motion, and assume $\langle v_r \rangle = \langle v_\theta \rangle = 0$ km s^{-1}. We define our vector of halo ellipsoid parameters to be $\theta_{\text{Halo}} = \{\langle v_\phi \rangle, \sigma_r, \sigma_\phi, \sigma_\theta\}$. We denote joint velocity PDF:

$$\begin{gathered} F_{v,\text{Halo}}(v_{r,i}, v_{\phi,i}, v_{\theta,i}) = \mathrm{N}(v_{r,i}|0, \sigma_r^2 + \sigma_{v_r,i}^2)\times \\ \mathrm{N}(v_{\phi,i}|\langle v_\phi \rangle, \sigma_\phi^2 + \sigma_{v_\phi,i}^2) \times \mathrm{N}(v_{\theta,i}|0, \sigma_\theta^2 + \sigma_{v_\theta,i}^2), \end{gathered} \tag{4.5}$$

where $v_{r,i}, v_{\phi,i}, v_{\theta,i}$ are the Galactocentric velocities corresponding to data y_i and distance D_i. The corresponding uncertainties on these velocities are denoted by $\sigma_{v_r,i}^2$, $\sigma_{v_\phi,i}^2$, $\sigma_{v_\theta,i}^2$. Proper motions in Galactic coordinates are converted to physical velocities using the fact that tangential velocity is proportional to distance: $v_T = 4.74047\mu D$, where μ is the proper motion in mas yr^{-1} and D is in kpc. Tangential velocities are converted to the Galactocentric frame by correcting for the projection of the Sun's velocity along a given line-of-sight. We convert (v_{LOS}, v_l, v_b) to spherical coordinates (v_r, v_ϕ, v_θ) by assuming a circular speed of 240 km s^{-1} at the position of the Sun ($R_0 = 8.5$ kpc), with solar peculiar motion $(U, V, W) = (11.1, 12.24, 7.25)$ km s^{-1} (Schönrich et al. 2010).

We note that in order to evaluate the probability of θ_{Halo} given our observables, we need to consider the Jacobian matrix from the coordinate transformation from the observed frame to the Galactocentric frame:

$$p(y_i|D_i, \theta_{\text{Halo}}) \propto F_{v,\text{Halo}} \times D_i^2 \cos b, \tag{4.6}$$

where the factor of $D_i^2 \cos b$ arises due to the change in variables.

4.3.3 Absolute Magnitudes

Finally, as in Deason et al. (2013b), Cunningham et al. (2016), and in Paper I, we additionally constrain the absolute magnitude to a given star $M_{F814W,i}$ using information on its $m_{F606W} - m_{F814W}$ color.

We weight VandenBerg et al. (2006) isochrones in the *HST* filters according to the approximate age and metallicity distributions of the MW halo. We then generate a KDE to get the probability distribution function $G(M_{F814W}|m_{F606W,i}, m_{F814W,i})$.

4.3.4 Full Posterior

We now summarize how we sample from the full posterior distribution, for parameters $\theta_{\text{Halo}}, \mathbf{M}, \mathbf{f}$ given observables $\mathbf{y}, \mathbf{x}$. We can write down the likelihood under this model for a star with data y_i, explanatory variables $x_i = \{m_{F606W,i}, m_{F814W,i}, l_k, b_k\}$, given our model parameters:

$$\begin{aligned} p(y_i|\theta_{\rm Halo}, M_{F814W,i}, f_{{\rm Disk},k}, x_i) = \\ p(M_{F814W,i}|m_{F606W,i}, m_{F814W,i}) \times D_i \\ \times \big[f_{{\rm Disk},k} \times p(y_i|\theta_{\rm Disk}, D_i) p(D_i|\rho_{\rm Disk}, l_k, b_k) + \\ (1 - f_{{\rm Disk},k}) \times p(y_i|\theta_{\rm Halo}, D_i) p(D_i|\rho_{\rm Halo}, l_k, b_k)\big], \end{aligned} \tag{4.7}$$

The extra factor of D_i arises due to the change of variables from absolute magnitude to distance: $M_{F814W} \propto \log(D)$. The full likelihood, using stars from $k = 1, ..., K$ fields containing $N_{*,k}$ stars, is given by the product of the likelihoods of each individual data point:

$$p(\mathbf{y}|\theta_{\rm Halo}, \mathbf{M}, \mathbf{f}, \mathbf{x}) = \prod_{k=0}^{K} \prod_{i=0}^{N_{*,k}} p(y_i|\theta_{\rm Halo}, M_{F814W,i}, f_{{\rm Disk},k}, x_i). \tag{4.8}$$

Likelihood in hand, we can write down the posterior distribution for our model parameters using Bayes Theorem:

$$p(\theta_{\rm Halo}, \mathbf{M}, \mathbf{f}|\mathbf{y}) \propto p(\mathbf{y}, \mathbf{M}|\theta_{\rm Halo}, \mathbf{f}) \times p(\theta_{\rm Halo}, \mathbf{f}), \tag{4.9}$$

where $p(\theta_{\rm Halo}, \mathbf{M}, \mathbf{f})$ is the prior distribution on model parameters. We assume standard reference priors on $\theta_{\rm Halo}$ (i.e., Jeffreys priors: $p(\langle v_\phi \rangle) \propto$ const and $p(\sigma) \propto 1/\sigma$ for all dispersions). We assume uniform priors on the $f_{\rm Disk}$ parameters ($p(f_{{\rm Disk},k}) = 1, f_{{\rm Disk},k}\ \epsilon\ [0, 1]$).

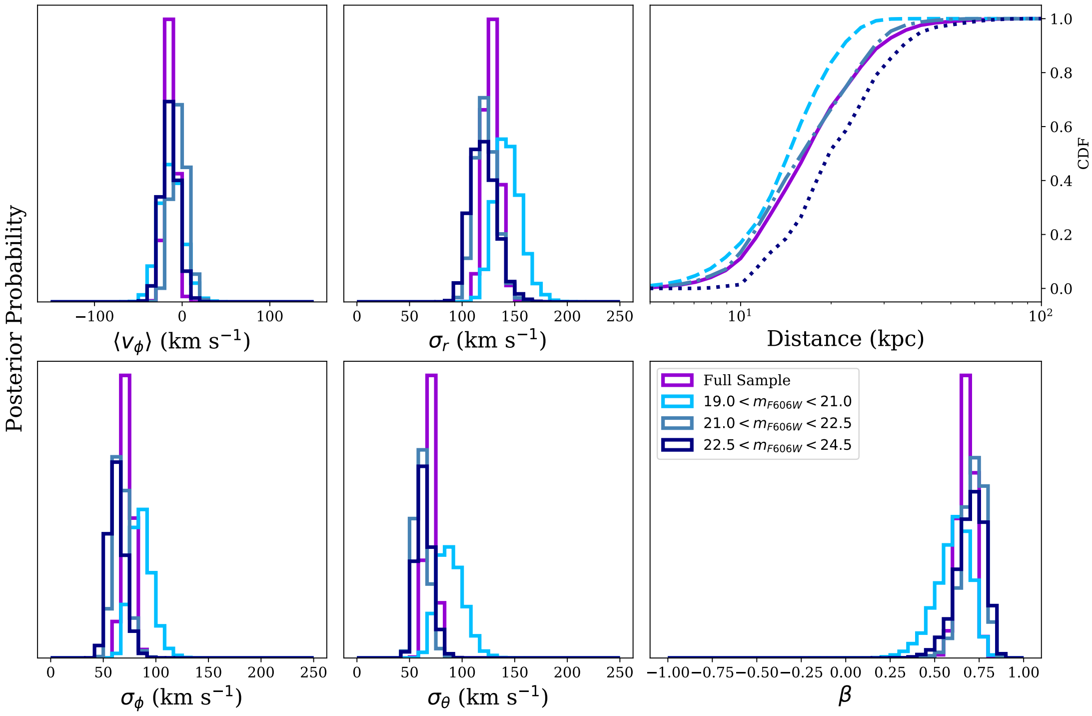

Figure 4.6: Summary of posterior results for spherically averaged samples. Left-hand panels: histograms of marginalized posterior samples for the four parameters of the halo velocity ellipsoid. Each of the estimates shown combines targets from all four survey fields. The estimates using the full HALO7D sample are shown in purple, while the blue histograms show the resulting estimates from three apparent magnitude bins. Upper righthand panel: CDFs of the distances of the full sample and the three apparent magnitude bins. Lower righthand panel: posterior distributions for β.

In order to sample for our model posterior parameters, we compute Equation 4.9 over a grid in absolute magnitude for every star. We then use `emcee` (Foreman-Mackey et al. 2013) to sample from our full posterior, marginalizing over the absolute magnitude of every star in each step of the chain. We test this modeling procedure on fake data; for details on how we generated fake data and tested our model, we refer the reader to Appendix C.

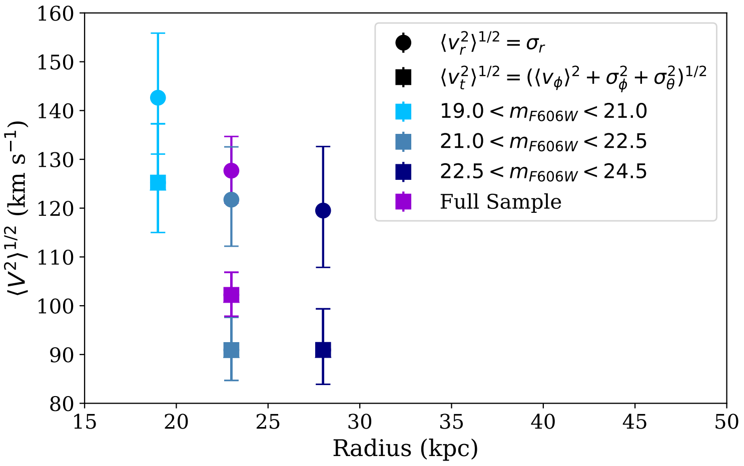

Figure 4.7: Square roots of the second moments of the radial (circular points) and tangential (square points) velocity distributions, as a function of mean Galactocentric radius. Different colors indicate the results from the analysis of the full sample (purple) as well as the three apparent magnitude bins (same colors as in Figure 4.6).

4.4 Results

In this section, we present posterior distributions for the halo velocity ellipsoid parameters. We first present the results using the full HALO7D sample, and then split our sample into three apparent magnitude bins. Finally, we consider the samples from each field separately.

4.4.1 Spherically Averaged Estimates

We first estimate the velocity ellipsoid parameters using the full HALO7D sample of 188 stars. The parameters in this model are the four ellipsoid parameters; the disk contamination in each of the four fields; as well as the absolute magnitude (and therefore distances) to each star. The resulting 1D marginalized distributions for the ellipsoid parameters are shown as the purple histograms in Figure 4.6. The left-

	$\langle v_\phi \rangle$ (km s^{-1})	σ_ϕ (km s^{-1})	σ_θ (km s^{-1})	σ_r (km s^{-1})	$\langle D \rangle$ (kpc)	$\langle r \rangle$ (kpc)	β	N_{Stars}
Full Sample	-13 ± 6	73^{+5}_{-4}	70 ± 4	128 ± 7	20	24	$0.68^{+0.04}_{-0.05}$	199
$19.0 < m_{F606W} < 21.0$	-1^{+13}_{-13}	86^{+10}_{-9}	88^{+14}_{-11}	143^{+13}_{-12}	16	19	$0.61^{+0.08}_{-0.11}$	74
$21.0 < m_{F606W} < 22.5$	-4 ± 9	66^{+7}_{-6}	61^{+7}_{-5}	122^{+11}_{-10}	18	23	$0.72^{+0.05}_{-0.07}$	73
$22.5 < m_{F606W} < 24.5$	-14 ± 9	62^{+8}_{-7}	63^{+8}_{-7}	120^{+13}_{-12}	24	28	$0.71^{+0.07}_{-0.09}$	52
COSMOS	-18 ± 9	74^{+8}_{-7}	61 ± 6	121^{+10}_{-9}	20	25	$0.67^{+0.06}_{-0.08}$	81
GOODS-N	-3^{+17}_{-16}	71^{+13}_{-11}	83^{+23}_{-18}	132^{+20}_{-16}	19	23	$0.64^{+0.12}_{-0.18}$	32
GOODS-S	-40^{+18}_{-17}	65^{+17}_{-12}	125^{+30}_{-22}	116^{+22}_{-17}	18	23	$0.14^{+0.33}_{-0.51}$	20
EGS	-1 ± 10	72^{+8}_{-7}	59^{+7}_{-6}	139^{+13}_{-11}	20	22	$0.77^{+0.05}_{-0.06}$	66

Table 4.3: Summary of the estimates of the parameters of the halo velocity ellipsoid, for the full sample, the apparent magnitude bins, and the individual HALO7D fields. Posterior medians are quoted, with errorbars giving the 16/84 percentiles.

hand panels show histograms of posterior samples for the four halo velocity ellipsoid parameters $\theta_{\rm Halo} = \{\langle v_\phi \rangle, \sigma_r, \sigma_\phi, \sigma_\theta\}$. Using the full sample of stars, we do not see a strong signature of halo rotation ($\langle v_\phi \rangle = -13 \pm 6$ km s^{-1}). We use the posterior samples of the ellipsoid parameters to derive a posterior estimate for the velocity anisotropy β; the resulting posterior distribution is shown as the purple histogram in the lower right-hand panel of Figure 4.6. We find that β is radially biased: at our mean sample distance of $\langle r \rangle = 23$ kpc, $\beta = 0.68^{+0.04}_{-0.05}$, consistent with estimates of β in the solar neighborhood (e.g., Bond et al. 2010).

In addition to modeling the full sample of stars, we also split our sample into three apparent magnitude bins. Because our distance estimates to each individual star are uncertain and probabilistic, we cannot divide our sample into different radial ranges; we therefore split the sample in apparent magnitude to study the radial variation of β. The resulting marginalized posterior distributions for the three apparent magnitude bins are shown as the blue histograms in Figure 4.6. Estimates using stars with $19.0 < m_{F606W} < 21.0$ are shown in light blue; the estimates from stars with $21.0 < m_{F606W} < 22.5$ are shown as gray blue; and $22.5 < m_{F606W} < 24.5$ are shown in dark blue. The cumulative distribution functions (CDFs) for the distances to each of the three samples are shown in the upper right-hand corner of Figure 4.6, along with the CDF for distance of the full sample.

Figure 4.7 shows the second moments of the velocity distributions as a function of the average Galactocentric distance to the sample. We see a trend of decreasing velocity dispersion with distance, in both tangential and radial motion. However, when

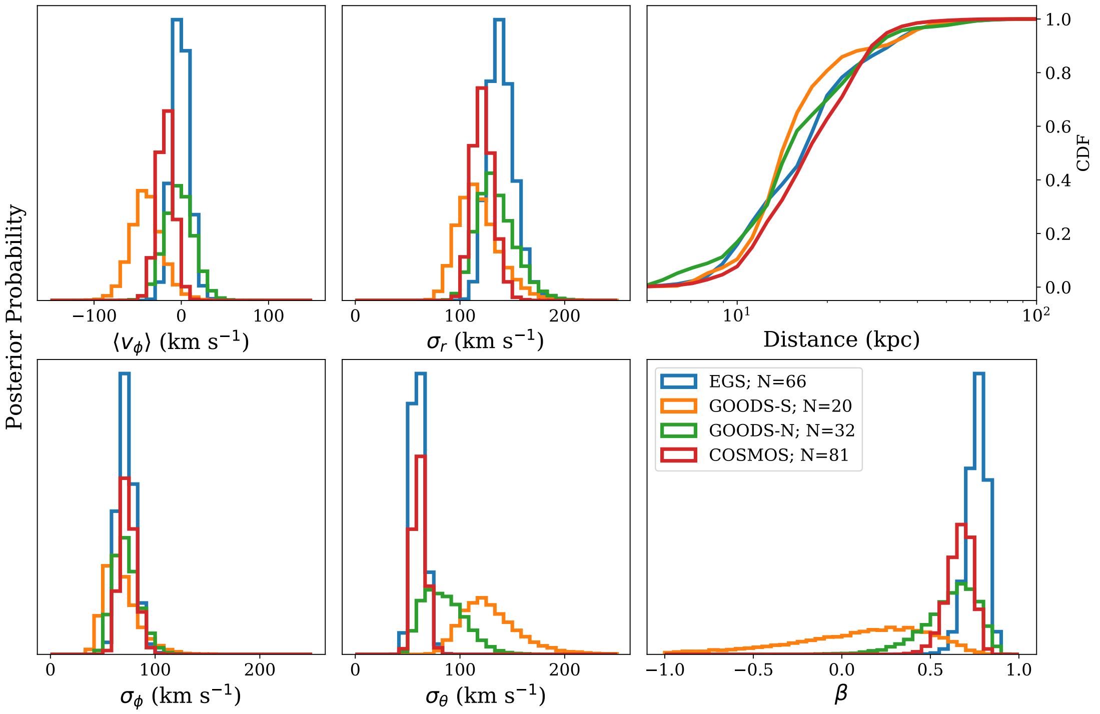

Figure 4.8: Summary of posterior results for the four fields. Left-hand panels: histograms of marginalized posterior samples for the four parameters of the halo velocity ellipsoid. Each colored histogram represents a different HALO7D field. Upper right-hand panel: CDFs for the distances of the stars in the four fields. Lower righthand panel: posterior distributions for β.

we compute the posterior distribution for β (blue histograms in lower right-hand panel of Figure 4.6), we find that all three estimates are consistent with the estimate of β from the full sample: β is radially biased for all of our spherically-averaged samples.

4.4.2 Individual Fields

In the case of modeling fields individually, the free parameters in our model include the velocity ellipsoid parameters, the distance to each star in the field, as well as the fraction of disk contamination in the field.

Posterior samples for the ellipsoid parameters in each of the four fields are shown in Figure 4.8. The left-hand panels show histograms of posterior samples for the four halo velocity ellipsoid parameters $\theta_{\text{Halo}} = \langle v_\phi \rangle, \sigma_r, \sigma_\phi, \sigma_\theta$. The upper right-hand panel shows the cumulative distribution for the distances to the four fields, and the lower right-hand panel shows the resulting posterior distribution for the velocity anisotropy.

When the four fields are treated separately, we see variation in the estimates of the velocity ellipsoid parameters. While the PDFs for GOODS-S are the broadest, because it has the smallest sample size, the GOODS-S distribution also shows a signature of rotation ($\langle v_\phi \rangle = -40^{+18}_{-17}$ km s^{-1}). The resulting estimate for β is consistent with isotropy, though also very broad, due in part to the small sample size in this field, but also due to the fact that circular orbits correspond to $\beta = -\infty$. In contrast, the estimates in the EGS field show no rotation, and the resulting estimate of β is strongly radially biased ($\beta_{\text{EGS}} = 0.77^{+0.05}_{-0.06}$).

4.4.3 Disk Contamination

The marginalized 1D posterior distributions for the disk contamination in each of the four HALO7D fields are shown in Figure 4.9. The posteriors for f_{Disk} when the fields are treated individually are the thick black histograms; the colored histograms show the estimates in a given field for the spherically averaged estimates.

Our estimates for disk contamination are low (on the order of or less than 10%); this is consistent with the predicted disk contamination levels predicted by the Besançon Galaxy Model (Robin et al. 2003; see Paper I). Because GOODS-N and GOODS-S have smaller sample sizes than EGS and COSMOS, their posterior distributions for f_{Disk} are broader, but the posterior modes are still around 10%. As is to be expected,

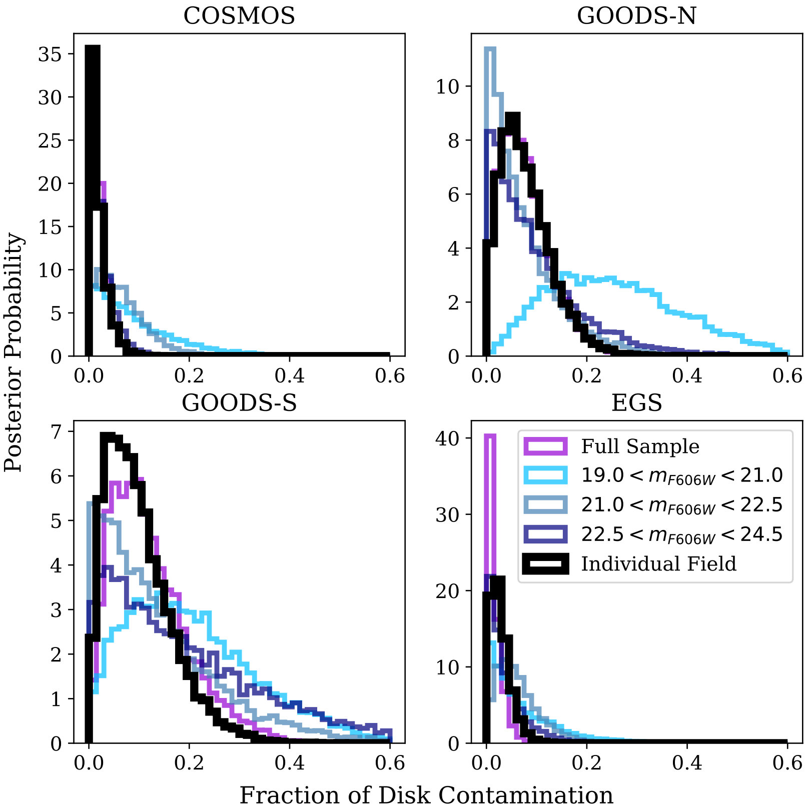

Figure 4.9: Posterior distributions for the disk contamination in the four HALO7D fields for each of the models. Black histograms indicate the posterior distributions for the fraction of disk stars when the fields are modeled individually. Colored histograms indicate the full sample (purple histograms) and the spherically averaged estimates in different apparent magnitude bins (as in Figures 4.6 and 4.7).

the disk contamination is highest for the brightest apparent magnitude bin (light blue histograms).

4.4.4 The β Radial Profile

Figure 4.10 summarizes all of our estimates of the velocity anisotropy, as a function of mean distance. Our spherically averaged estimates are plotted as circles, while the estimates of our individual fields are shown as squares. Gray points are results from other studies that used 3D kinematics to estimate β: gray triangles show the estimates of β from MW GCs, using PMs from both HST (Sohn et al. 2018) and *Gaia* (Watkins et al. 2019). The gray diamond shows the 3D estimate of β in the solar neighborhood from SDSS (Bond et al. 2010), and the gray square shows the C16 estimate of β along the line of sight towards M31. Gray crosses show the recent estimates of β from Lancaster et al. (2019) using Gaia DR2 PMs of blue horizontal branch stars.

Our spherically averaged estimates of β, which find radially biased $\beta \sim 0.6 - 0.7$, are consistent with one another and with other studies that have estimated β averaging over different parts of the sky. However, our field-to-field estimates (including the estimate from C16) show substantial variation, from strongly radially biased (EGS) to isotropic (GOODS-S) to mildly tangentially biased (M31). While the GOODS-S and M31 fields each have lower posterior estimates for β, these two fields also have the smallest sample size. Because of the way β is defined, estimates of β are sensitive to sample size and measurement uncertainties. We therefore assess how much our sample size should concern us by testing fake data. We generate 100 fake datasets (in the method described in Appendix C), from velocity distributions that have $\beta_{\rm True} = 0.75$, each containing 20 stars. Out of the 100 tests, only one of the fake datasets had a posterior distribution for β with median $\beta_{\text{Med}} \leq 0.14$ (for the full distribution of β posterior medians, see Figure C.3 in Appendix C). We do not see small sample size resulting in a systematic underestimate of β. Therefore, while the small sample size does contribute to the large uncertainty on β in this field, based on our fake data testing, we do not expect that the observed isotropic β is purely due to sample size.

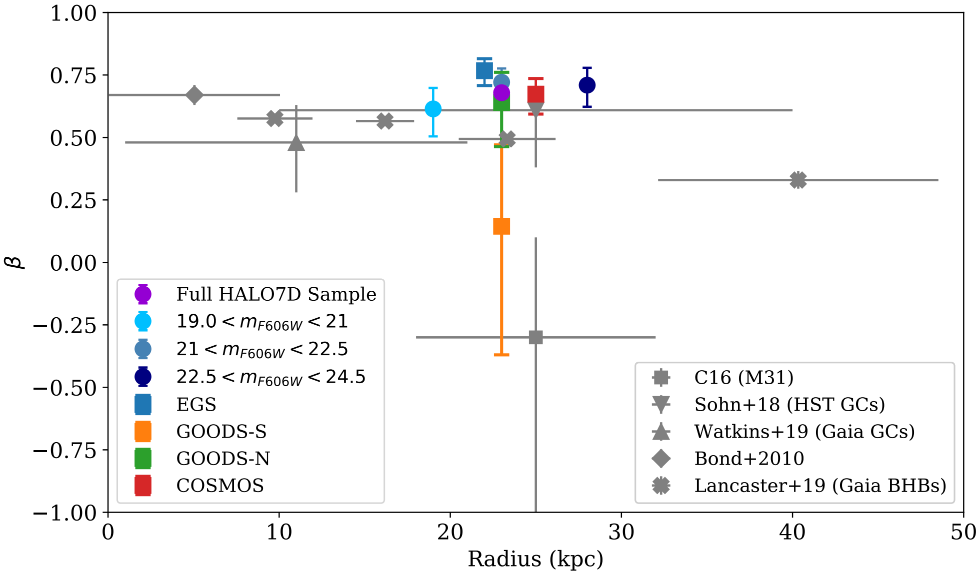

Figure 4.10: The Milky Way's radial anisotropy profile, β, as measured from 3D kinematics. Colored points indicate results from this work, while grey points indicate findings from previous work and other studies. The grey diamond shows the anisotropy estimate from Bond et al. (2010), using main sequence stars from SDSS, and the grey square shows the estimate from Cunningham et al. (2016), using 13 MW MSTO stars along the line of sight towards M31. Gray crosses show the estimates of β from Lancaster et al. 2019, using the overall estimate of β consisting of the mixture of a metal-rich, radially biased population and an isotropic, metal-poor population, from BHBs in Gaia DR2. The two recent estimates for β from MW globular clusters, using HST PMs and *Gaia* PMs, are shown with triangles (Sohn et al. 2018, Watkins et al. 2019). Square shaped points are results from individual fields, while our spherically averaged results are shown as circles. When using small fields to estimate anisotropy, β varies from mildly tangential (e.g., GOODS-S, M31) to strongly radial (e.g., EGS). However, the spherically averaged estimates are all consistently $\beta \sim 0.6$ (and consistent with solar neighborhood and GC estimates), and the posterior means increase as a function of mean sample distance.

4.5 Comparison with Other Studies

In this paper, we use the HALO7D dataset to estimate the parameters of the MW stellar halo velocity ellipsoid. We study the full HALO7D sample, the sample divided into three apparent magnitude bins, and the individual HALO7D fields. When averaging over the four HALO7D fields, we find consistent estimates for $\beta \sim 0.6 - 0.7$. Our spherically averaged results for β are consistent with results from other recent estimates of β using GCs as tracers (Sohn et al. 2018, Watkins et al. 2019). However, when we treat the four HALO7D fields separately, our estimates for the ellipsoid parameters, and thus β, show significant variation.

In their study of β profiles of simulated galaxies, Loebman et al. (2018) found that β profiles are generally increasingly radially biased as a function of radius. However, recently accreted material can cause short-lived (< 0.2 Gyr) dips in the β profiles, and longer-lived (> 0.2 Gyr) dips arise due to the disruption of the in-situ stellar halo by the close passage of a massive satellite. These "dips" in the in-situ stellar halo are more metal-rich than dips caused by the accreted stellar halo.

Several studies using LOS velocities alone have observed tangentially biased "dips" in the β profile (Sirko et al. 2004; Kafle et al. 2012, King et al. 2015); these dips occur approximately at the observed break in the MW density profile (Deason et al. 2011, Sesar et al. 2011, Watkins et al. 2009). The kinematic structure around the break radius $r_b \sim 27$ kpc is of interest in order to understand its origin. In their study of the Bullock & Johnston (2005) purely accreted stellar halos, Deason et al. (2013a) found that the buildup of stars at apocenter from a relatively early, massive accretion event, or a few synchronous events, can cause broken density profiles.

As pointed out by Hattori et al. (2017), studies of stars beyond $r \sim 15$ kpc using only LOS velocities are subject to underestimates of β. However, recent studies using *Gaia* PMs have found decreases in β around the break radius as well. Using blue horizontal branch stars in *Gaia* DR2, Lancaster et al. (2019) found that β decreases just beyond the break radius, from $\beta \sim 0.6$ at 20 kpc to $\beta \sim 0.4$ at 40 kpc (gray crosses in Figure 4.10). They argue that this is due to sharp decline in the fraction of stars belonging to a radially biased population that dominates the inner halo (i.e., the *Gaia*-Sausage) beyond its apocenter radius (which Deason et al. 2018b showed coincides with the MW break radius). Using LAMOST K-Giants with *Gaia* DR2 PMs, Bird et al. (2018) found strongly radially biased ($\beta \sim 0.8$) inside of $r \sim 25$ kpc, with β gradually decreasing beyond this radius, down to $\beta = 0.3$ at 100 kpc; however, Lancaster et al. (2019) also showed that the magnitude of the decrease observed in the Bird et al. (2018) study could be due to their treatment of measurement uncertainties.

When averaging over multiple fields, we do not see a dip in the β profile, nor a global decrease in β beyond the break radius. While our estimates are around the MW break radius, the posterior medians of our spherically averaged estimates are consistent with one another. Increasing β as a function of radius is predicted by simulations (e.g., Abadi et al. 2006, Sales et al. 2007, Rashkov et al. 2013, Loebman et al. 2018). However, we need to probe to larger distances beyond the break radius to see if this prediction is consistent with the data, or if β starts to decrease (as seen by Bird et al. 2018 and Lancaster et al. 2019).

When we treat our different lines-of-sight separately, we see potential evidence for a dip in β towards GOODS-S and M31. Based on the Loebman et al. (2018) findings, these sightlines could be dominated by material that has been recently accreted or kicked up by the passage of Sagittarius. As discussed in the Introduction, several overdensities previously believed to be accreted structures now show evidence of a potential disk origin, having been kicked out of the disk due to the passage of Sagittarius (e.g.,Price-Whelan et al. 2015; Laporte et al. 2018; Bergemann et al. 2018). One such overdensity discussed in those works is TriAnd, located along the line of sight towards M31, which is also the lowest latitude of the HALO7D fields.

Measuring abundances for stars in the HALO7D fields from their Keck spectra (McKinnon et al., in prep) will help to distinguish between the kicked-up disk scenario and the recent accretion scenario as the origin for the observed "dips" in β in GOODS-S and M31. Chemical abundances will also help to assess the origin of the strongly radially biased β estimate in EGS ($\beta_{\mathrm{EGS}} \sim 0.8$). Belokurov et al. (2018) discovered the "*Gaia*-Sausage" as a metal-rich ([Fe/H] > -1.7), radially biased ($\beta \sim 0.9$) population in *Gaia* DR1. Given that the estimate of β in EGS is more radially biased than the estimates of β in the other fields, it is possible that the sample of stars in EGS is dominated by Sausage stars. Chemical abundances will be essential in assessing to what extent the Sausage is contributing to the HALO7D sample.

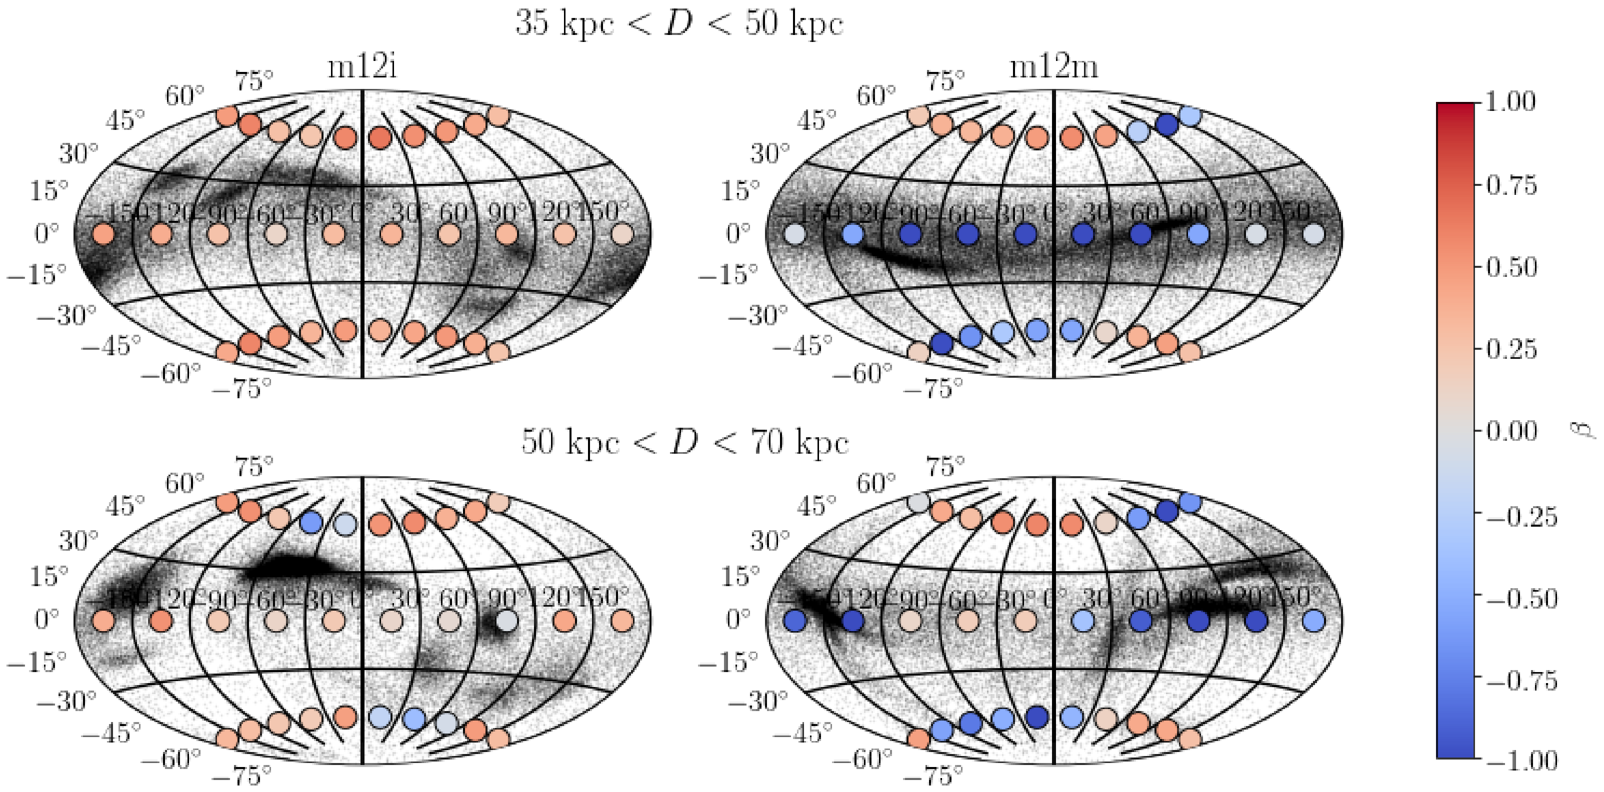

Figure 4.11: Maps of velocity anisotropy of the stellar halo in two *Latte* FIRE-2 simulations of Milky Way-mass galaxies: m12i (left) and m12m (right). The top panels show stars in the distance range 35 kpc $< D <$ 50 kpc, while the lower panels show 50 kpc $< D <$ 70 kpc. Black points show the positions of star particles in Galactic coordinates. In each halo, the sky has been subdivided into patches, and the color of the large circle in each patch indicates the velocity anisotropy for that subset of stars. Within a given distance range, each halo shows variation in its velocity anisotropy across the sky. Variation as a function of distance is also evident. In addition, the median and spread in β both vary from halo to halo: many more of the star particles in m12m are on tangentially biased orbits than in m12i.

4.6 Comparison with Simulations

When we treat the four HALO7D fields separately, we see variation in the estimates of the velocity ellipsoid parameters (and the resulting velocity anisotropy). In this section, we explore the spatial variation of velocity anisotropy in two halos from the *Latte* suite of FIRE-2 cosmological zoom-in baryonic simulations of Milky Way-mass galaxies (introduced in Wetzel et al. 2016), part of the Feedback In Realistic Environments (FIRE) simulation project.[6] These simulations were run using the Gizmo gravity plus hydrodynamics code in meshless finite-mass (MFM) mode (Hopkins 2015) and the FIRE-2 physics model (Hopkins et al. 2018). In this work, we discuss halos m12i (initially presented in Wetzel et al. 2016) and m12m (initially presented in Hopkins et al. 2018), making use of the publicly available $z = 0$ snapshots (Sanderson et al. 2018).[7] The properties of the halos of these galaxies have been shown to agree reasonably well with the properties of the MW and M31, including the stellar-to-halo mass relation (Hopkins et al. 2018); satellite dwarf galaxy stellar masses, stellar velocity dispersion, metallicities, and star-formation histories (Wetzel et al. 2016, Garrison-Kimmel et al. 2018, Escala et al. 2018, Garrison-Kimmel et al., in prep); and stellar halos (Sanderson et al. 2017, Bonaca et al. 2017). In particular, the high resolution of the *Latte* simulations (star particles have initial masses $m \approx 7000 M_\odot$ and gravitational force softening of 4 pc) means that they resolve satellite dwarf galaxies down to $M_{\rm star} \gtrsim 10^5 M_\odot$, thus resolving the galaxies that are expected to contribute the majority of mass to the formation of the stellar halos (e.g., Deason et al. 2015).

---

[6] FIRE project website: `http://fire.northwestern.edu`

[7] $z = 0$ snapshots available at `http://ananke.hub.yt`

Despite the high resolution of the *Latte* simulations, at large distances in the halos the typical spacing between star particles can be large compared to the sizes of the HALO7D fields (less than a square degree). Therefore, a detailed comparison of exactly how our selection effects, observational errors and field sizes are affecting our results is beyond the scope of this work. In addition, such a comparison may be of limited usefulness, given that variation observed across areas as small as the HALO7D fields could be due to structures below the simulation resolution limit (i.e., debris from accreted satellites with $M_{\rm star} < 10^5 M_\odot$). As a first step, we explore the spatial variation in the velocity anisotropy computed directly from the star particles in the simulation, using 30 larger fields, each spanning 36 degrees in longitude and 60 degrees in latitude.

Figure 4.11 shows the m12i simulation (left panels) and m12m (right panels). The top panels show the positions of star particles (black points), in Aitoff projection, within the distance range $35 < D < 50$ kpc; the lower panels show star particles in the distance range $50 < D < 70$ kpc. While these radial ranges are farther out than the HALO7D data, we choose these ranges to avoid the thick disks in these simulations, which are extended and kinematically hotter than the MW (Sanderson et al. 2018, Loebman et al, in prep). Star particle positions are plotted in Galactic coordinates. Galactocentric frames are defined in the method described in Section 3 of Sanderson et al. (2018), and positions are converted to Galactic coordinates using the `astropy.coordinates` package (Astropy Collaboration et al. 2013; Astropy Collaboration et al. 2018). We use the default options in `astropy.coordinates` for the position of the Sun ($R_\odot = 8.3$ kpc; Gillessen et al. 2009). The "sky" in each halo has been divided into patches, and the resulting velocity anisotropy computed from the star particles in each patch is shown by the colored points. We note that we compute the velocity anisotropy using all the star particles within a given area on the sky and radial range; we do not exclude particles in bound satellite galaxies. This choice likely affects the resulting β maps, and we plan to explore the affects of excluding and including bound satellites on β estimates in future work.

The velocity anisotropy β shows variation across the sky of a given halo, as well as with radius. In addition, these maps of velocity anisotropy are very different across the two halos: within $35 < D < 50$ kpc all patches in m12i are radially biased, but many of the patches of m12m are tangentially biased β. At $50 < D < 70$ kpc, both m12i and m12m show some tangentially biased patches and radially biased patches. The overall means and spreads of β values measured across the two halos are quite different: m12i has a mildly radially biased $\langle\beta\rangle \sim 0.3$ with a standard deviation of 0.2, whereas m12m has tangentially biased $\langle\beta\rangle \sim -0.3$ with a standard deviation of 0.7. The magnitude of the β variation observed in these two simulated galaxies is very similar to the range of β values measured across the HALO7D fields; in both Figures 4.10 and 4.11, we see variation in β over the range of $[-1, 1]$.

The differences in the β maps across these two halos are likely linked to their different accretion histories. Over these radial ranges in the simulated halos, the majority of the material mapped in Figure 4.11 is accreted, and the accreted debris in the two halos have visibly different spatial and kinematic properties. We intend to explore in future work what characteristics of a galaxy's accretion history, such as accretion times, initial orbital conditions, and masses of accreted satellites, are primarily responsible for the observed β variation patterns. Based on the Loebman et al. (2018) findings, patches with tangentially biased β could be indicating recently accreted material. Further study of the accretion histories of these simulated galaxies will help us to understand what accretion events and accretion histories give rise to different β variation patterns in galaxy halos, and what characteristics of the MW's assembly history we might be able to constrain through mapping its spatial β variation.

The β variation we observe in HALO7D and the *Latte* simulations also could have implications for the validity of MW mass estimates derived from Jeans modeling. The fundamental assumption underlying Jeans modeling is that the tracers are virialized and in dynamical equilibrium. The spatial maps and β variation observed in the Latte halos reveal that this assumption is clearly violated in the simulations. The variation in β observed with HALO7D indicates that this assumption is invalid in the MW halo as well; our results are evidence that the halo is not phase-mixed at $\langle r \rangle \sim 23$ kpc. Just how significantly the violation of the assumption of dynamical equilibrium will affect estimates of the MW mass remains to be determined. The systematic uncertainty of traditional spherical Jeans mass modeling in recovering halo masses has been observed in a number of simulations (e.g., Wang et al. 2018, Kafle et al. 2018, Eadie et al. 2018); we leave the full characterization of the effects of β variation on different approaches of MW mass estimates on the *Latte* halos to future work.

4.7 Conclusions

In this paper, we present the proper motions of distant, main sequence turnoff MW halo star candidates as measured with *HST*. These PMs are measured as a part of the HALO7D project, and have LOS velocities measured from Keck spectroscopy (Paper I). Using the 3D kinematic sample from HALO7D, we estimate the parameters of the halo velocity ellipsoid and velocity anisotropy. We estimate these parameters treating the individual survey fields separately as well as spherically averaging over all fields.

We summarize our main results as follows:

1. Using the full HALO7D sample of 188 stars, we estimate the velocity anisotropy $\beta = 0.68^{+0.04}_{-0.05}$ at $\langle r \rangle = 23$ kpc. This estimate is consistent with other recent estimates of β.
2. We estimate β from the HALO7D sample split into three apparent magnitude bins to explore the radial dependence. While estimates of velocity dispersions decrease as a function of mean sample distance, the overall estimates of β are consistent across apparent magnitude bins. Posterior medians increase as a function of mean sample distance, consistent with predictions from simulations.
3. When we treat our stars from the four HALO7D fields separately, estimates of the halo velocity ellipsoid parameters show variation from field to field. This variation could be evidence for recent accretion; it is also possible that the isotropic and tangentially biased β values from GOODS-S and M31 (respectively) are due to

the presence of kicked-up disk stars. The observed variation in β is evidence that the halo is not phase-mixed at $r \sim 23$ kpc.

4. We map the velocity anisotropy in two stellar halos from the *Latte* suite of FIRE-2 simulations and see variation in β across the skies of these two halos over a similar range to the variations observed in the HALO7D fields. In the simulated galaxies, the degree of, and patterns in, these variations are clearly tied to their different accretion histories. A more detailed study of the full accretion histories of these galaxies will shed light on the types of signatures that different accretion events can leave in β maps.

Fortunately, many of the questions raised in this work are answerable in the near future. Abundances from HALO7D spectra will provide key insights as to the origin of the observed β variation. In addition, β variation in the MW can be mapped with the *Gaia* dataset and, ultimately, LSST. Thanks to the quality of current and upcoming data, coupled with high resolution cosmological hydrodynamical simulations, we are rapidly progressing in our knowledge of our Galaxy's structure and formation.

Chapter 5

Conclusions

5.1 Summary

In this dissertation, I have presented observations, kinematic measurements, and first kinematic results from the HALO7D survey. In Chapter 2, I described how we made the first estimate of the velocity anisotropy β outside the solar neighborhood using 13 stars located along the line of sight towards M31. I argued that this observed decrease in velocity anisotropy could indicate the presence of a shell.

In Chapter 3, I described the spectroscopic observations for the HALO7D survey. I presented a new method of measuring LOS velocities from a stellar spectrum making use of Bayesian hierarchical modeling for combining multiple observations. I modeled the LOS velocity distributions as a mixture of disk and halo populations, and found consistent estimates of the LOS velocity dispersions across the four HALO7D fields. These estimates were also consistent with the estimates of σ_{LOS} from other studies. By performing our same analysis procedure on mock observations of the Bullock

& Johnston (2005) simulations, we found that a single massive accretion event, or several coincident accretion events, can give rise to consistent velocity distributions across all four sightlines.

In Chapter 4, I described the HALO7D PM measurements using deep, multi-epoch *HST*imaging. I described the new technique to measure PMs utilizing Bayesian mixture modeling to select galaxies for use in the reference frame and galaxies that should be excluded from the reference frame. I then discussed using the 3D kinematic sample from HALO7D to estimate the parameters of the velocity ellipsoid. I found very consistent estimates of $\beta = 0.6 - 0.7$ when spherically averaging over the four HALO7D fields; however, when treating the fields separately, I found variation in the estimates of β. When comparing to the *Latte* simulations, I found that these stellar halos also show variation across the sky in β, and that the maps of the two halos are very different, because of their accretion histories.

The next step for the HALO7D survey is to measure abundances (McKinnon et al., in prep). With abundances, we can start to make statements about halo formation, including in-situ/accreted populations and masses of accreted dwarfs. In addition, abundances would provide a key dimension of information for identifying and characterizing substructure in the HALO7D dataset.

However, the results from HALO7D raise many interesting questions that can be addressed with larger MW surveys and simulations. Below, I describe several ways in which we can further investigate the questions raised by the results of my dissertation work.

5.2 Future Directions

In my study of the halo kinematic structure with the HALO7D dataset (C18b), I found that velocity anisotropy not only changes as a function of radius, but also can vary across the sky. I found that the *Latte* suite of simulations also showed both spatial and radial variation in β, and that the extent to which β varies within a given halo is likely linked to its accretion history (Figure 2b). This evidence that the halo is not phased-mixed, even at $\langle r \rangle = 23$ kpc, is problematic for many MW mass estimation methods that rely on Jeans modeling. However, this observational result came from five small fields. In order to assess the extent to which our assumption of constant β (or constant $\beta(r)$) in mass modeling is violated, we must: a) study the variation of β a over larger area of the sky; b) measure β over additional deep pencil beams; and c) use simulations to understand the link between β variation and accretion history, and to quantify the impact of the observed β variation on MW mass estimates.

5.2.1 Mapping Velocity Anisotropy with Gaia

The *Gaia* dataset contains astrometric measurements for an outstanding 1 billion stars; for the first time, we now have the ability to estimate 3D kinematic properties, like velocity anisotropy β, over the full sky.

I propose to construct a map of velocity anisotropy in the MW halo using the *Gaia* dataset. This problem is far from trivial. First, because of the way β is defined ($-\infty < \beta < 1$), errorbars on estimates of β are asymmetric and sensitive to sample size and measurement uncertainties. Second, *Gaia* measurement uncertainties are large at halo distances, and most of the halo stars will not have useful distance constraints from parallaxes. Finally, selection of halo star candidates in halo kinematics studies is always challenging. Halo stars are often identified through Toomre diagrams, which is a kinematic classification (e.g., Bonaca et al. 2017, Helmi et al. 2018); members of the Sagittarius are also often removed from halo star samples based on kinematic cuts (e.g., Bird et al. 2018). However, kinematic selection in kinematic analysis leads to biased results.

Bayesian hierarchical mixture modeling provides a natural framework to address the challenges in this analysis. This modeling strategy enables the straightforward incorporation of measurement uncertainties; returns full posterior probability distributions for model parameters (namely, β); and allows for multiple populations through mixture modeling. This framework is a natural extension of the HALO7D analysis I performed in my thesis. For HALO7D, I modeled the parameters of the velocity ellipsoid in a Bayesian context, incorporating measurement uncertainties (including unknown distances) and disk contamination through a mixture model. I propose to use these tools and modeling frameworks that I developed for HALO7D to spatially map β with the *Gaia* dataset. These same techniques can be applied to kinematic measurements from other future astrometric datasets (e.g., LSST, WFIRST). Using Bayesian statistical modeling to map β variation in *Gaia* will be the most detailed study of the velocity structure in the MW halo to date.

5.2.2 The Milky Way Halo Main Sequence with *HST* Proper Motions

The full-sky *Gaia* β map proposed above will be unprecedented. However, there will be limitations to working exclusively with the *Gaia* dataset. Because of its limiting magnitude of $G \sim 20$, beyond $D \sim 15$ kpc in the halo, *Gaia* will only measure PMs for giants and evolved stars. While giants make great tracers because of their bright apparent magnitudes, giants are also rare; averaging over large areas of the sky (and thus potential inhomogeneities in the halo) is often required when estimating halo properties with giants. For dense sampling in halo populations, it is essential to also have PMs of main sequence (MS) stars; beyond $D \sim 15$ kpc, these PMs can only be measured with *HST* (Deason et al. 2013b; C18b).

To complement the above proposed *Gaia* β map, we could use *HST* PMs to map the velocity structure of the halo using MS stars in deep pencil beams. I used multi-epoch *HST* data in the CANDELS fields to measure PMs of halo stars as faint as $v \sim 24.5$ (C18b). However, there are many more *HST* pointings that have multi-epoch coverage that have yet to be analyzed for PMs.

Specifically, data used for our PM measurements must meet the following conditions: (1) observations must be obtained with either ACS/WFC or WFC3/UVIS; (2) observations must be in one of the broad-band filters F606W, F775W, F814W, or F850LP; (3) time baseline of the multiple epochs must be at least 2 yr; (4) combined exposure time in the shallower epoch must be at least one orbit long; and (5) individual exposure time must be at least 300 s long. Based on a preliminary search, we can double the sample of MW halo stars with measured *HST* PMs using exclusively archival data.

These deep pencil beams will enable the detection of fluctuations in β that will be unresolved by *Gaia*.

The *HST* PM sample also provides an exciting opportunity to measure precise PMs for the most distant stars in our Galaxy. Using color and PM information, we could identify a population of candidate MW halo giant stars with *HST* PMs. Figure 3 demonstrates how incorporating PM information into giant candidate selection vastly reduces the contamination from nearby MW dwarfs. In the magnitude range of $20 < V < 22.5$, these stars span distances 50–250 kpc, reaching well beyond the limit of 100 kpc for *Gaia* giants. In addition, these targets are well within the magnitude range for follow-up ground based spectroscopy; with LOS velocities and abundances, this project will yield *the most distant sample of stars with measured 3D kinematics in the MW.*

5.2.3 Anisotropy Variation in Cosmological Simulations

In order to understand what the spatially varying β implies about the MW's assembly, and to asses its impact on the estimates of the MW mass, one must look to theory. In order to interpret the results of C18b (as well as the above proposed projects), we could use the *Latte* simulations (introduced in Wetzel et al. 2016). The *Latte* suite of FIRE-2 cosmological zoom-in baryonic simulations of MW-mass galaxies have very high resolution, such that satellite dwarf galaxies are resolved down to $M_{\rm star} \gtrsim 10^5 M_\odot$, thus resolving the galaxies that are expected to contribute the majority of mass to the formation of the stellar halos (e.g., Deason et al. 2015).

Investigating the connection between β variation and accretion history in the *Latte* simulations is a way forward in understanding what the observed β variation tells us about our Galaxy's formation. Using the simulations, we can determine the properties of the accreted satellites that leave signatures in the β maps, such as their accretion times, initial orbital conditions, and masses. In addition, the connection between β variation and the overall accretion history properties can be quantified, such as the mass spectrum of accreted dwarfs and the distribution of accretion times. By addressing these questions in the simulations, we will learn what characteristics of the MW's assembly history can be constrained with the results from the above proposed observational projects.

The simulations also provide the opportunity to explore the implications of β variation on MW mass estimation methods. The fundamental assumption underlying Jeans modeling is that the tracers are virialized and in dynamical equilibrium. The C18b results show that this assumption is clearly violated in both the halo and the *Latte* simulations. However, what remains to be determined is how significantly the violation of the assumption of dynamical equilibrium will affect estimates of the MW mass.

We can determine the effects of the assumptions of Jeans modeling on the inferred masses and shapes of the *Latte* halos, given their observed β variation patterns. Other options for MW mass modeling can also be tested, such as the clustering of stream stars in action space (Sanderson et al. 2017) and mass-scaling methods (e.g., Li et al. 2017, Callingham et al. 2019) perform relative to Jeans modeling. The systematic uncertainty of traditional spherical Jeans mass modeling in recovering halo masses has been observed in a number of simulations (e.g., Wang et al. 2018, Kafle et al. 2018, Eadie et al. 2018); however, this has yet to be tested at the resolution of *Latte*. Testing mass estimation methods on the *Latte* halos, and observing their sensitivity to β variation, has the potential to reconcile the disparate estimates of the Galaxy's mass.

Appendix A

Testing Velociraptor

In this appendix, we discuss the testing of the Velociraptor technique described in Section 3.3 on fake data. To create fake spectra for testing, we degraded our template spectrum to a variety of signal to noise ratios. In order to realistically reproduce the noise due to the sky background for one of the HALO7D exposures, we took the noise array for an extremely faint extragalactic target that had no visible continuum or emission lines for a 20 minute exposure.

For signal to noise ratios of 3, 5, 10, and 30, we generated 180 individual spectra for each S/N. We tested Velociraptor in "single mode" (i.e., working with individual observations only) for 90 of these spectra for each S/N. The mean recovered velocities, and the standard deviations of these velocities, are shown in the top panel of Figure A.1. For the spectra with S/N=3, approximately one third of the fake sample had failed velocity measurements (i.e., chains did not successfully converge); the errorbars in Figure A.1 reflect the statistics for the successful measurements.

All projections of the posterior for one of our fake spectra can be seen in the corner plot in Figure A.2. The input values for the fake spectrum are shown as blue lines; the model successfully recovers the parameters of the fake data. The absorption line coefficients C and the continuum levels for a given spectral region are covariant; this is expected because of the way in which we parametrized the absorption line strength (see Equation 3.4).

Trace plots for 20 of the `emcee` walkers for the 11 parameters of single mode are shown in Figure A.3. The true parameter values that were used to generate the fake data are shown as thick black dashed lines. The traces are well mixed and converge successfully over the runtime of the sampler.

To test Velociraptor in hierarchical mode, we combined six fake spectra at a given signal to noise ratio, and ran Velociraptor 30 times at each S/N. The resulting mean recovered velocities, and their standard deviations, are shown in the middle panel of Figure A.1, where here the signal to noise ratio plotted on the $x-$axis refers to the signal to noise of the individual observations.

Trace plots and a corner plot for the corrected velocity, the variance of velocities σ_v^2, and the raw velocity parameters for six observations, each with S/N=5, are shown in Figures A.5 and A.4. We run `emcee` for 5,000 steps, and discard the first 3,000 as burn in. For the purposes of displaying the corner plot, we have "thinned" our chain, including every 50th sample for each walker.

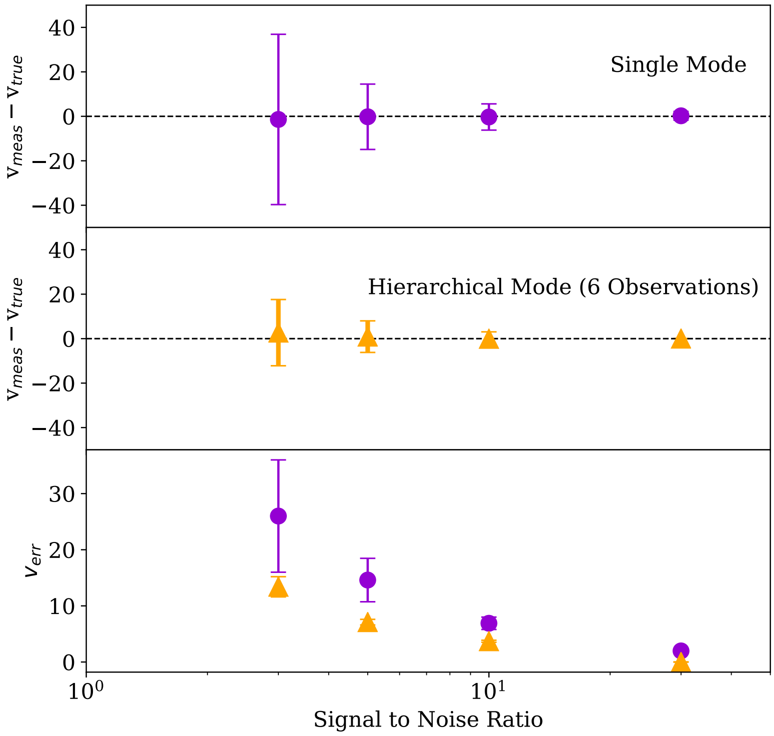

Figure A.1: Results from testing Velociraptor on fake data. Top panel: errorbars show the mean recovered velocity, and the standard deviation of the recovered velocities, for 90 runs of Velociraptor in single-mode, as a function of signal to noise. Middle panel: resulting distributions of recovered velocities when Velociraptor is run with six observations in hierarchical mode. Note that here the x-axis refers to the signal to noise of a single observation. Lower panel: velocity error (computed as half the difference between the 84th and 16th percentiles) in single-mode (purple) and hierarchical mode with six observations (orange).

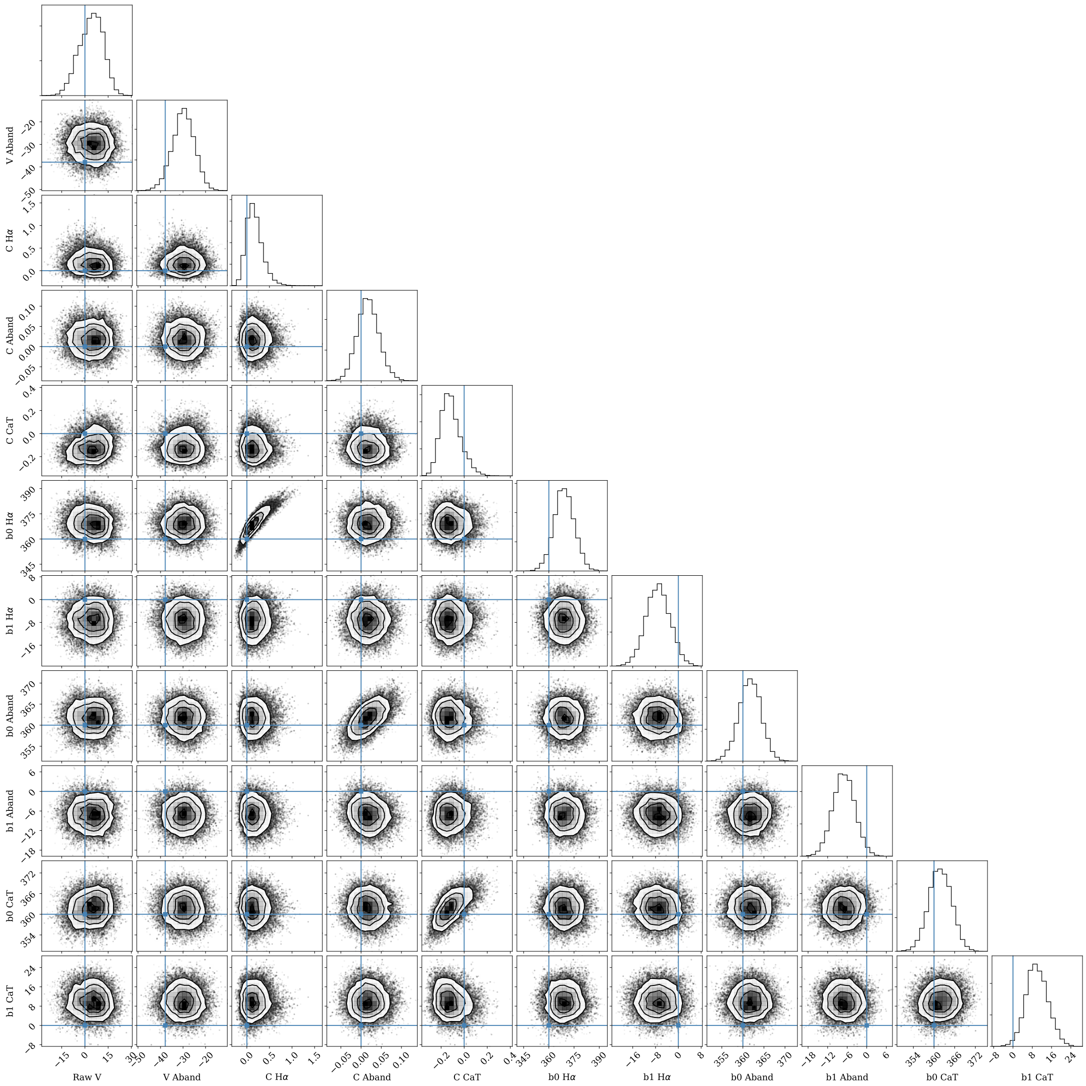

Figure A.2: Full corner plot for all 11 parameters for a fake spectrum with S/N=10. The true parameter values used to generate the fake spectrum are shown in blue. The absorption line strength parameters for a given spectral region are covariant with the continuum level; this is expected based on how we have parameterized the absorption lines.

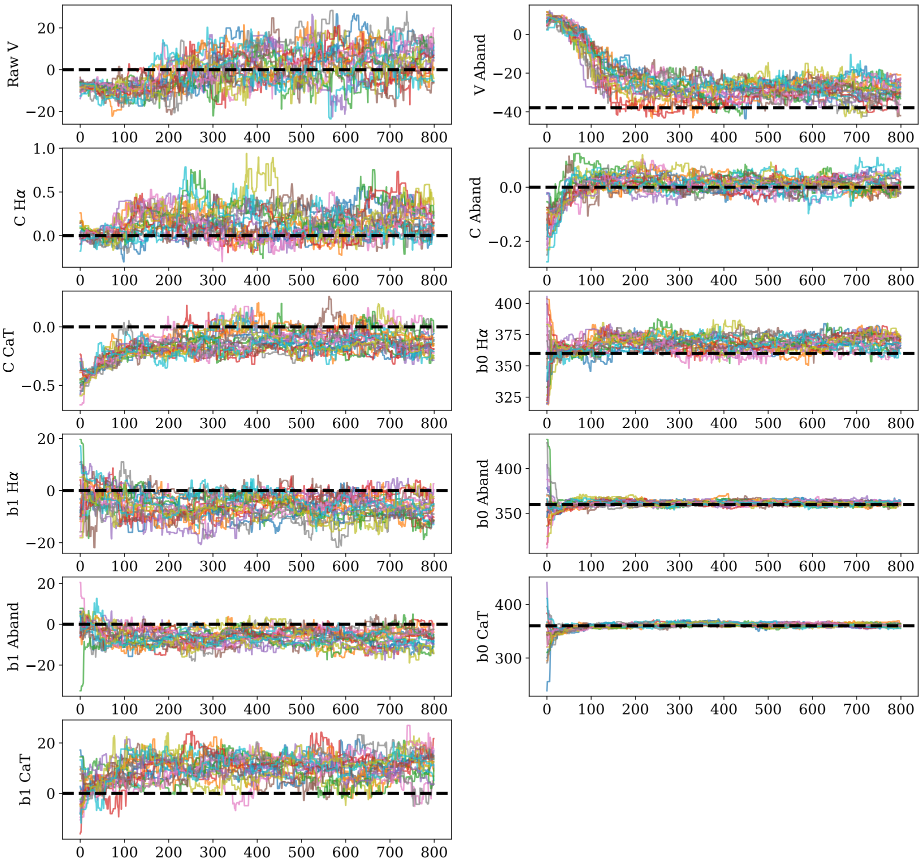

Figure A.3: Traces for all 11 single-mode parameters, for a fake spectrum generated to have S/N=10. For clarity, we show traces for only 20 randomly selected walkers. Black dashed lines indicate the true values of the model parameters used to generate this fake spectrum.

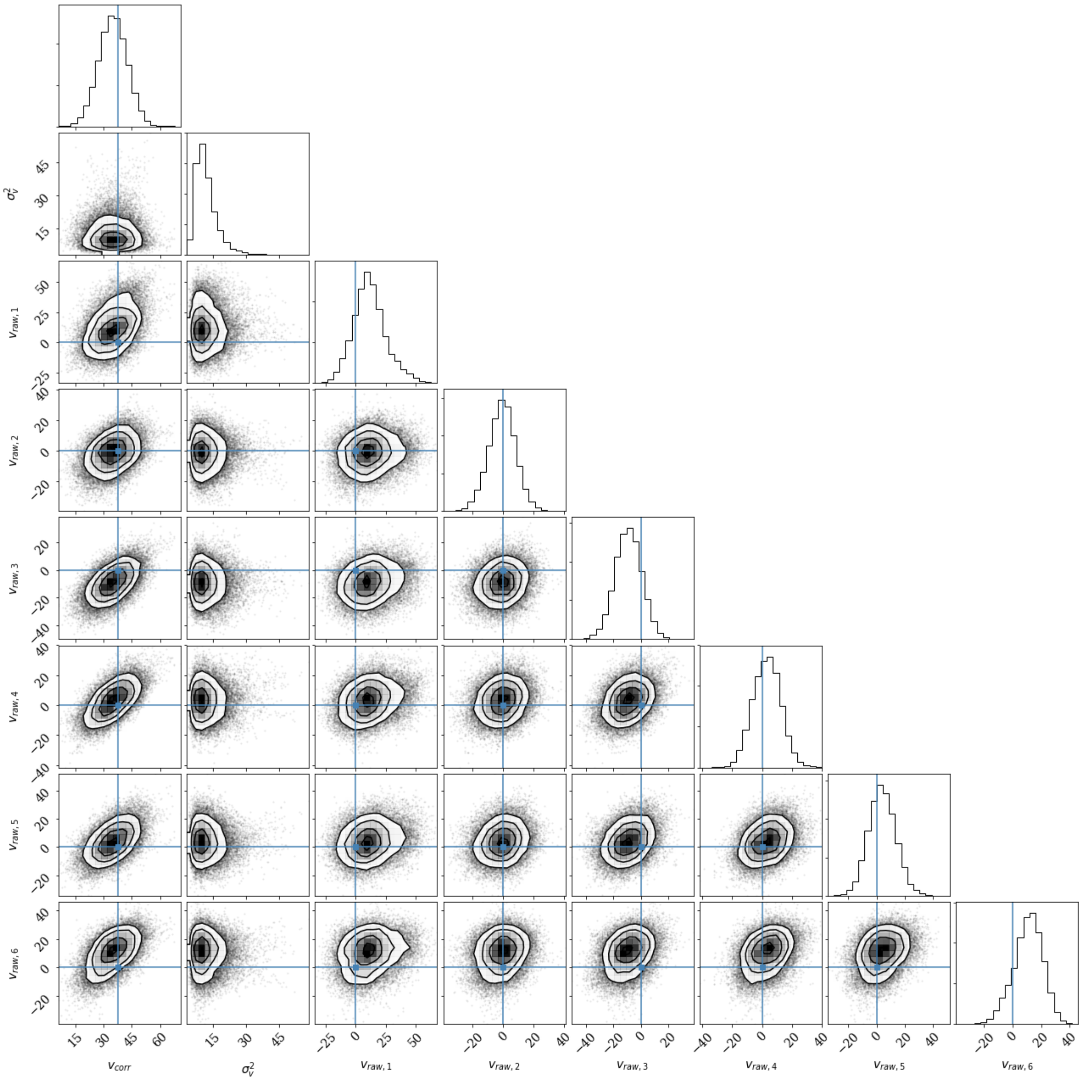

Figure A.4: Corner plot for the corrected velocity, the additional uncertainty σ_v^2, and the six raw velocities for six fake spectra that each have S/N=5. Note that we are only showing projections here for 8 out of the 68 parameters in this model. This particular run of Velociraptor ran for with 800 walkers for 500 steps. For this figure, we excluded the first 3000 steps as burn-in, and thinned the chain, showing every 50th sample.

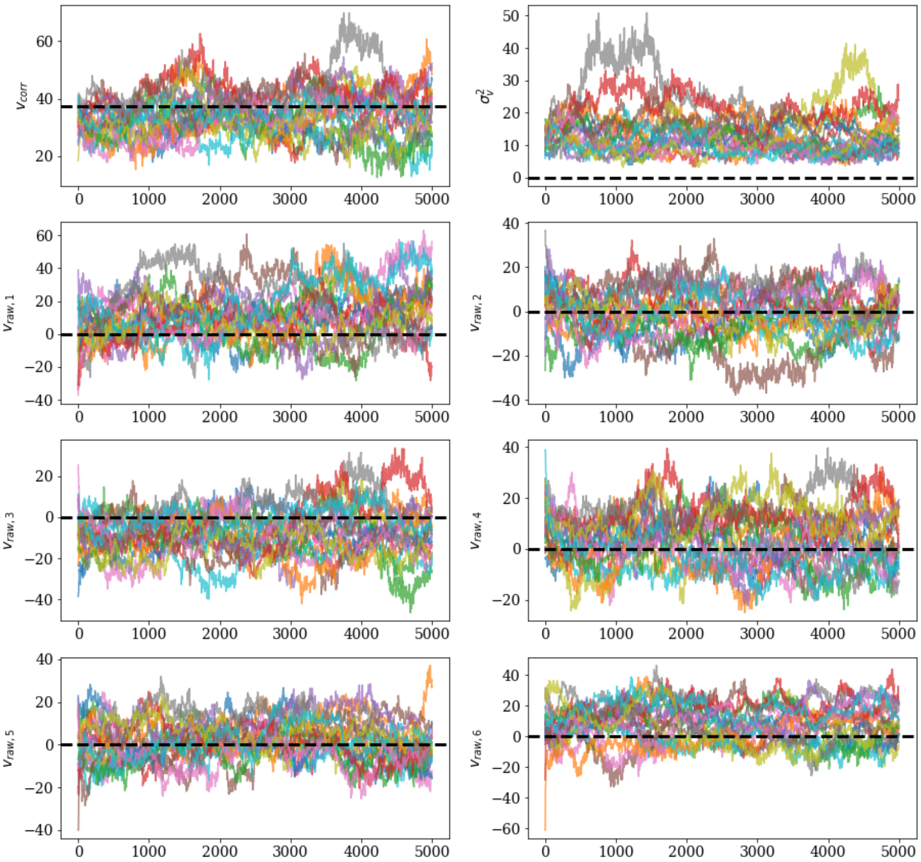

Figure A.5: Traces for the corrected velocity, the additional uncertainty σ_v^2, and the six raw velocities for six fake spectra that each have S/N=5. For clarity, we show traces for only 20 randomly selected walkers. Truths are shown as black dashed lines. Because of the complexity of the model and the large number of free parameters, the chains do not mix efficiently, and the sampler needs to be run for many iterations. Note that the true value for σ_v^2 is not recovered in this case, because all of our fake spectra were generated to have exactly the same velocity.

Appendix B

Bayesian Mixture Modeling for PM Measurements

In this appendix, we provide the details of the Bayesian mixture model used to derive estimates of PMs. We first describe our model in Section B.1, and then the Gibbs sampling algorithm used to sample from the posterior distribution for model parameters in Section B.1.2.

B.1 Proper Motion Model

To map one image onto another, we perform a six-parameter linear transformation:

$$\begin{pmatrix} A & B & x_t - \delta u \\ C & D & y_t - \delta v \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} u \\ v \\ 1 \end{pmatrix} = \begin{pmatrix} u_{ref} \\ v_{ref} \\ 1 \end{pmatrix}, \tag{B.1}$$

where $\delta u, \delta v$ represent the change (in pixels) of a star from one image to another due to proper motion (so, for galaxies, $\delta u = \delta v = 0$).

In our model, we treat stars and galaxies separately. For stars:

$$\begin{aligned} Au + Bv + x_t - \delta u &\sim \mathrm{N}(u_{ref}, \sigma_*^2) \\ Cu + Dv + y_t - \delta v &\sim \mathrm{N}(v_{ref}, \sigma_*^2), \\ u_{imref}, v_{imref} &\sim \mathrm{N}(u_{ref}, v_{ref}, \sigma_*^2) \end{aligned} \tag{B.2}$$

where u_{imref}, v_{imref} are the measured positions in the defined reference image, whereas u_{ref}, v_{ref} are the positions of the object in the reference epoch, which are free parameters. Because stars' central positions are well measured, we define $\sigma_* = 0.02$ pixels. $\delta u, \delta v$ represent the shift in pixels from image 1 to image 2, which can be converted to proper motions North and West, respectively:

$$PMN = \frac{\delta u \times 50 \text{ mas/pix}}{\Delta t}; \tag{B.3}$$

$$PMW = \frac{\delta v \times 50 \text{ mas/pix}}{\Delta t}, \tag{B.4}$$

where Δt is measured in years. We describe the galaxy positions as a two-component

Gaussian mixture model. Defining a fixed location in an image as the galaxy's precise position is not trivial and sometimes fails, given that galaxies are resolved sources with complex morphologies. We therefore consider "good" galaxies (i.e., galaxies with well measured positions) and "bad" galaxies (galaxies with poorly measured positions).

$$\begin{aligned} Au + Bv + x_t &\sim \mathrm{N}(u_{ref}, \sigma^2) \\ Cu + Dv + y_t &\sim \mathrm{N}(v_{ref}, \sigma^2)\cdot \\ u_{imref}, v_{imref} &\sim \mathrm{N}(u_{ref}, v_{ref}, \sigma^2) \end{aligned} \tag{B.5}$$

For "good" galaxies, $\sigma = 0.1$ pixels, whereas for "bad" galaxies, $\sigma = 3$ pixels.

An example of the initial classification of "good" and "bad" galaxies is shown in Figure B.1. Figure B.1 shows the change in positions in pixels, in the distortion-corrected frame, for objects in two *HST* images, taken seven years apart. Black points show the positions of galaxies initially classified as "good"; these are clustered at $(0, 0)$, because they were used in the reference frame for the linear transformation. Positions of galaxies initially classified as "bad" are shown as grey crosses. In our Bayesian mixture model, we allow galaxies to move in and out of the reference frame probabilistically. Pink points show the change in the positions of the stars in the images. These stars have a mean motion and scatter relative to the stationary reference frame of distant galaxies; these relate to the dynamical quantities of interest estimated in this study.

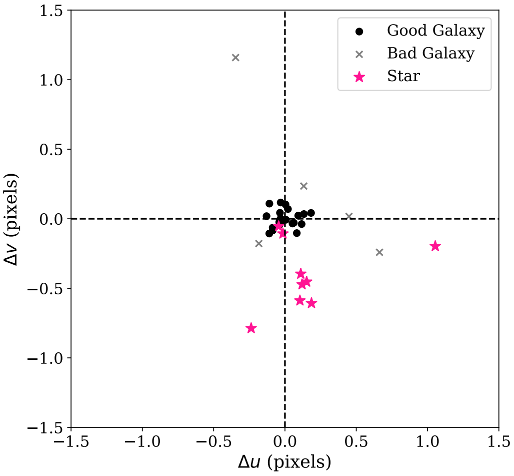

Figure B.1: An example of the linear transformation method on two images, j8pu44cvq (taken in 2004) and jboa38c2q (from 2011). Axes represent the change in pixels, in the distortion-corrected frame (u, v), for objects in the two images, after applying the six parameter linear transformation. Grey crosses indicate the change in positions for the galaxies initially classified as "bad"; black points are the positions of "good" galaxies used in the reference frame. The change in positions of the "good" galaxies are clustered at $(0, 0)$. Pink stars show change in the positions of the stars in these two images.

B.1.1 Gaussian Mixture Models written with Indicator Variables

Mixture models can be expressed in different ways. For a two component mixture model, the likelihood of a given data point can be written as

$$p(x|\theta) = \lambda N(x|\theta_1, \sigma_1^2) + (1-\lambda)N(x|\theta_2, \sigma_2^2) \tag{B.6}$$

Where λ is the fraction of objects in the underlying population that belong to distribution 1. However, sums in probability calculations make posterior sampling more difficult. To improve our sampling efficiency, we can re-write the above equation using indicators z_j:

$$p(x, z|\theta) = (\lambda N(x|\theta_1, \sigma_1^2))^{z_1}((1-\lambda)N(x|\theta_2, \sigma_2^2))^{z_2}. \tag{B.7}$$

In this construction, for a given step in the MCMC chain, the indicator $z_{ih} = 1$ if datapoint x_i is associated with component h and $z_{ih} = 0$ otherwise.

Our full posterior thus takes the form:

$$
\begin{aligned}
&p(\theta|u_1...u_k, v_1...v_k, u_{imref}, v_{imref}) \propto \\
&\prod_{k=1}^{N_{im}} \prod_{j=1}^{N_{stars}} \exp\left\{ \frac{1}{2\sigma_*^2}(Au_{jk} + Bv_{jk} + x_{t,k} - \delta u_{jk} - u_{ref})^2 \right. \\
&\qquad \left. + \frac{1}{2\sigma_*^2}(Cu_{jk} + Dv_{jk} + y_{t,k} - \delta v_{jk} - v_{ref})^2 \right\} \\
&\times \prod_{j=1}^{N_{gals}} \left(f_{g,k} \exp\left\{ \frac{1}{2\sigma_g^2}(Au_{jk} + Bv_{jk} + x_{t,k} - u_{ref})^2 \right. \right. \\
&\qquad \left. \left. + \frac{1}{2\sigma_*^2}(Cu_{jk} + Dv_{jk} + y_{t,k} - v_{ref})^2 \right\} \right)^{z_{g,jk}} \\
&\times \left((1 - f_{g,k}) \exp\left\{ \frac{1}{2\sigma_b^2}(Au_{jk} + Bv_{jk} + x_{t,k} - u_{ref})^2 \right. \right. \\
&\qquad \left. \left. + \frac{1}{2\sigma_b^2}(Cu_{jk} + Dv_{jk} + y_{t,k} - v_{ref})^2 \right\} \right)^{z_{b,jk}}
\end{aligned}
\tag{B.8}
$$

where $f_{g,k}$ is the fraction of good galaxies in image k, and z_{jg} is the indicator for galaxy j in image k. By construction, if a galaxy has a "good" position in image k, $z_{g,jk} = 1$ and $z_{b,jk} = 0$ (i.e. a galaxy can only belong to one mixture component at a time).

B.1.2 Gibbs Sampling Algorithm

To sample from the posterior distribution for our parameters, we use Gibbs sampling. In a Gibbs sampler, we sample directly from the conditional posterior distributions for each parameter. Gibbs samplers can only be used if the full conditional distributions of the parameters can be written in closed form, which is usually only the case when conjugate priors (or, in special cases, reference priors) have been used.

Our Gibbs sampling algorithm consists of the following steps:

1. Initialize the transformation parameters for each image using standard linear-least squares. If the image is in the same epoch as the reference image, use the star positions to match frames. Otherwise, use the positions of the "good" galaxies. Initial values for PMs are averaged over the images, and initial values for the reference positions are those in the reference image.

2. For each star, we draw from the conditional posterior distributions for PMN and PMW, as well as the conditional posterior distributions for the reference positions. The conditional distributions for proper motions are:

$$PMW \sim N(\frac{50 \text{ mas/pix}}{k\sum \Delta t_k^2} \sum_{k=1}^{N_{im,k}} \delta u_k \times \Delta t_k, (50 \text{ mas/pix})^2 \times \frac{\sigma_*^2}{\sum \Delta t^2}) \tag{B.9}$$

$$PMN \sim N(\frac{50 \text{ mas/pix}}{k\sum \Delta t_k^2} \sum_{k=1}^{N_{im,k}} \delta v_k \times \Delta t_k, (50 \text{ mas/pix})^2 \times \frac{\sigma_*^2}{\sum \Delta t^2}) \tag{B.10}$$

3. For each galaxy, we first loop over each image, including the reference image, and draw an indicator. We allow a galaxy to be "good" in some subset of images and "bad" in another. We draw the indicator for a given galaxy as a Bernoulli variable with probability:

$$\frac{f_{gal} \times N(u|u_{ref}, \sigma_g^2) \times N(v|v_{ref}, \sigma_g^2)}{f_{gal} \times N(u|u_{ref}, \sigma_g^2) \times N(v|v_{ref}, \sigma_g^2) + (1 - f_{gal}) \times N(u|u_{ref}, \sigma_b^2) \times N(v|v_{ref}, \sigma_b^2)} \tag{B.11}$$

We then draw from the conditional posterior distributions for the reference frame

Parameter	μ/V	V^{-1}
A	$\sum \frac{u_j \times u_{ref,j}}{\sigma_j^2} - B \sum \frac{u_j \times v_j}{\sigma_j^2} - x_t \sum \frac{u_j}{\sigma_j^2} + \sum \frac{\delta u_j u_j}{\sigma_j^2}$	$\sum \frac{u_j^2}{\sigma_j^2}$
B	$\sum \frac{v_j \times u_{ref,j}}{\sigma_j^2} - A \sum \frac{u_j \times v_j}{\sigma_j^2} - x_t \sum \frac{v_j}{\sigma_j^2} + \sum \frac{\delta u_j v_j}{\sigma_j^2}$	$\sum \frac{v_j^2}{\sigma_j^2}$
C	$\sum \frac{u_j \times v_{ref,j}}{\sigma_j^2} - D \sum \frac{u_j \times v_j}{\sigma_j^2} - y_t \sum \frac{u_j}{\sigma_j^2} + \sum \frac{\delta v_j u_j}{\sigma_j^2}$	$\sum \frac{u_j^2}{\sigma_j^2}$
D	$\sum \frac{v_j \times v_{ref,j}}{\sigma_j^2} - C \sum \frac{u_j \times v_j}{\sigma_j^2} - y_t \sum \frac{u_j}{\sigma_j^2} + \sum \frac{\delta v_j v_j}{\sigma_j^2}$	$\sum \frac{v_j^2}{\sigma_j^2}$
x_t	$A \sum \frac{u_j}{\sigma_j^2} + B \sum \frac{v_j}{\sigma_j^2}$	n_{obj}/σ_j^2
y_t	$C \sum \frac{u_j}{\sigma_j^2} + D \sum \frac{v_j}{\sigma_j^2}$	n_{obj}/σ_j^2

Table B.1: Parameters of the conditional posterior distributions for the image transformation parameters. Conditional posteriors for all 6 image transformation parameters are normal distributions with mean μ and variance V. Sums are over all objects in an image. σ_j=0.02 for stars,σ_j=0.1 for "good" galaxies, and σ_j=3. for "bad" galaxies.

positions of each object.

4. Finally, we loop over each image, drawing from the conditional posterior distributions for the image parameters $(A, B, C, D, u_t, y_t, f_{gal})$ (see Table B.1). The f_{gal} parameters are drawn from beta distributions:

$$p(f_{gal,k}|...) \sim \text{Beta}(n_{goodgals,k} + 1, n_{badgals,k} + 1) \tag{B.12}$$

where $n_{goodgals,k}, n_{badgals,k}$ are computed by summing the indicators for each population at that step in the chain.

Appendix C

Fake Data Testing for Ellipsoid Model

In this Appendix, we discuss how we tested our method for estimating the parameters of the velocity ellipsoid using fake data.

To create fake data for a given line of sight for this model, we:

1. Generate samples from our kernel density estimate for M_{F81W} vs $M_{F606W} - M_{F814W}$ based on the weighted VandenBerg et al. (2006) isochrones.
2. Assign each draw an apparent magnitude, drawn from a uniform distribution in m_{F814W} over the range $[19, 24.5]$.
3. Given the resulting distances from the draws in apparent and absolute magnitudes, we use Monte Carlo acception/rejection to keep draws consistent with the MW density profile of Deason et al. (2011).

4. Assign stars velocities in spherical coordinates, based on random draws from normal distributions.

5. Convert D, V_r, V_ϕ, V_θ to μ_l, μ_b, v_{LOS} using the `astropy.coordinates` package. Given that we do not use `astropy.coordinates` to perform the velocity transformations in our ellipsoid modeling code, this step provides an additional check on our coordinate transformations.

6. Draw fake measured values from normal distributions centered on μ_l, μ_b, v_{LOS}, with dispersions corresponding to measurement uncertainties. For the purposes of this testing, we assign PM uncertainties of 0.2 mas yr^{-1} and LOS velocity uncertainties based on a fit of the relation between apparent magnitude and LOS velocity error shown in Figure 7 of Paper I.

We generate fake disk stars using a similar method, except drawing stars from the density profile and velocity distributions for our disk model. Figure C.1 shows the posterior distribution for the halo ellipsoid parameters when our analysis is performed on a fake dataset. This particular fake dataset contains 100 halo star and 50 disk stars in the GOODS-N field. Values for the parameters used to generate the data are shown in blue.

Results from testing 30 fake halo datasets, each with 100 stars, are shown in Figure C.2. Top panels show histograms of posterior medians for each simulated dataset; bottom panels are histograms of the errors measured in each dataset (computed as half the difference of the 84 and 16 percentiles). The errors in the posterior distributions are reasonable given the observed spread in posterior medians. The resulting distribution of posterior medians for β are shown in the lefthand panel of Figure C.3.

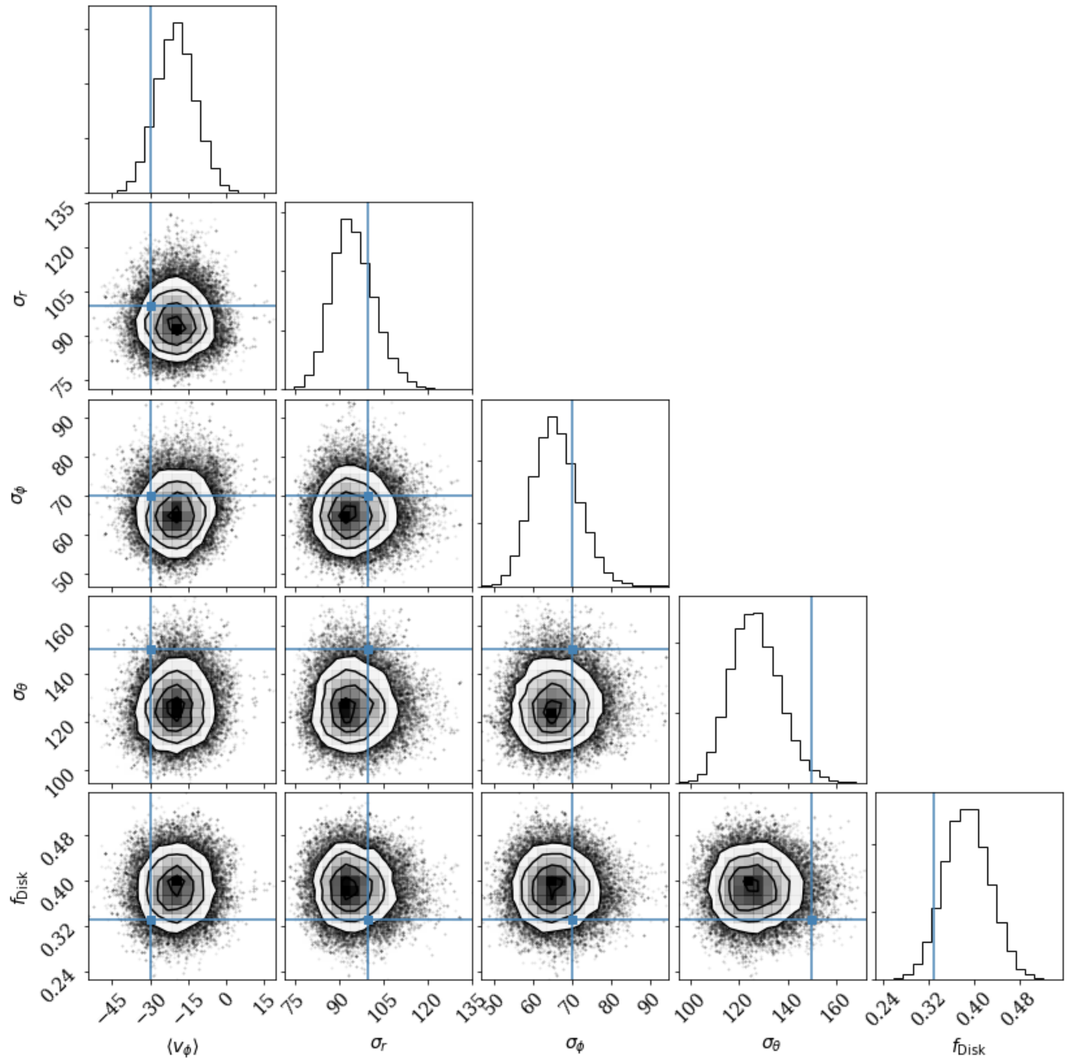

Figure C.1: Resulting projections of posterior samples for fake GOODS-N data. This fake sample contained 100 halo stars and 50 disk stars. The true values of the distributions used to generate the data are shown in blue.

C.1 Sensitivity to Sample Size

In order to assess how the sample size of the GOODS-S field is affecting the estimate of β in that field, we generated 100 fake datasets, each containing 16 stars. These datasets were generated from velocity distributions that have $\beta_{\text{True}} = 0.75$. Figure C.3 shows the distribution of the resulting posterior medians for β when we model this fake dataset. Out of the 100 fake datasets, only one had a posterior median with $\beta < 0.14$.

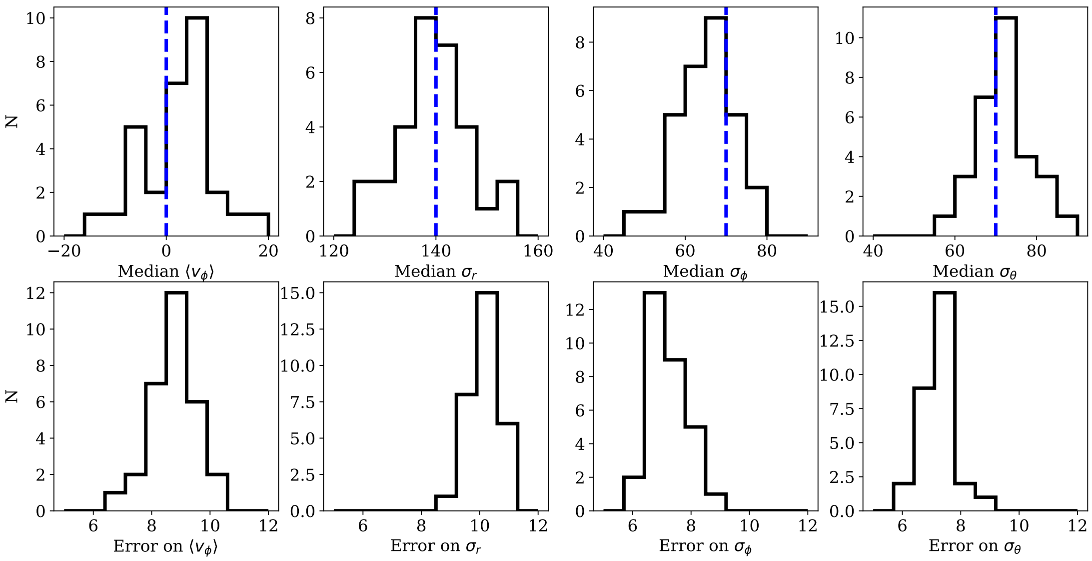

Figure C.2: Top panel: Distributions of posterior medians for the halo velocity ellipsoid parameters recovered from 30 fake datasets, each with 100 stars. Parameter values used to generate the fake data are shown as blue vertical dashed lines. Bottom panel: histograms of the the error estimates for each parameter.

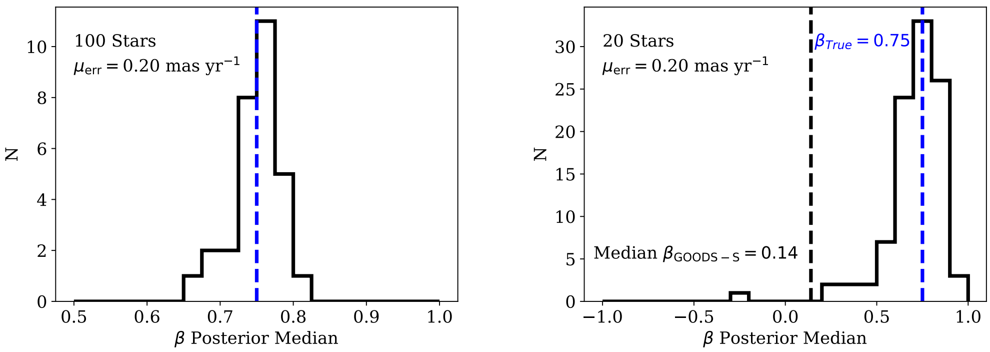

Figure C.3: Histograms of posterior medians for the estimates of β from fake data testing. Lefthand panel: distribution of β estimates from 30 fake datasets, each containing 100 stars, with PM uncertainties of 0.2 mas yr^{-1}. Righthand panel: the estimates of β from 100 fake datasets, each containing 20 stars, with PM errors of 0.20 mas yr^{-1}. For both sets of fake datasets, radial velocity uncertainties were assigned as a function of apparent magnitude (see Figure 7 of Paper I). None out of the 100 fake datasets yielded a posterior median estimate of $\beta < 0.14$; the minimum posterior median estimate was $\beta_{\rm min} = 0.20$.

Bibliography

Abadi, M. G., Navarro, J. F., & Steinmetz, M. 2006, MNRAS, 365, 747

Anderson, J., & Bedin, L. R. 2010, Publications of the Astronomical Society of the Pacific, 122, 1035

Anderson, J., & King, I. R. 2006, PSFs, Photometry, and Astronomy for the ACS/WFC, Tech. Rep. Instrument Science Report ACS 2006-01,

Anderson, J., & van der Marel, R. P. 2010, ApJ, 710, 1032

Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, A&A, 558, A33

Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, AJ, 156, 123

Barro, G., Pérez-González, P. G., Gallego, J., et al. 2011, ApJS, 193, 13

Beers, T. C., Carollo, D., Ivezić, Ž., et al. 2012, ApJ, 746, 34

Belokurov, V., Erkal, D., Evans, N. W., Koposov, S. E., & Deason, A. J. 2018, MNRAS, 478, 611

Bergemann, M., Sesar, B., Cohen, J. G., et al. 2018, Nature, 555, 334

Bertin, E., & Arnouts, S. 1996, A&AS, 117, 393

Binney, J. 2010, MNRAS, 401, 2318

Binney, J., & McMillan, P. 2011, MNRAS, 413, 1889

Binney, J., & Tremaine, S. 2008, Galactic Dynamics: Second Edition

Bird, S. A., Xue, X.-X., Liu, C., et al. 2018, ArXiv e-prints, arXiv:1805.04503

Bonaca, A., Conroy, C., Wetzel, A., Hopkins, P. F., & Kereš, D. 2017, ApJ, 845, 101

Bond, N. A., Ivezić, Ž., Sesar, B., et al. 2010, ApJ, 716, 1

Bovy, J. 2015, The Astrophysical Journal Supplement Series, 216, 29

Brown, T. M., Ferguson, H. C., Smith, E., et al. 2004, AJ, 127, 2738

Brown, T. M., Smith, E., Ferguson, H. C., et al. 2009, ApJS, 184, 152

Brown, W. R., Geller, M. J., Kenyon, S. J., & Diaferio, A. 2010, AJ, 139, 59

Bullock, J. S., & Johnston, K. V. 2005, ApJ, 635, 931

Callingham, T. M., Cautun, M., Deason, A. J., et al. 2019, MNRAS, 484, 5453

Capak, P., Aussel, H., Ajiki, M., et al. 2007, ApJS, 172, 99

Cappellari, M., & Emsellem, E. 2004, PASP, 116, 138

Carollo, D., Beers, T. C., Lee, Y. S., et al. 2007, Nature, 450, 1020

Carollo, D., Beers, T. C., Chiba, M., et al. 2010, ApJ, 712, 692

Cohen, J. G., Sesar, B., Bahnolzer, S., et al. 2017, ApJ, 849, 150

Conroy, C., Bonaca, A., Naidu, R. P., et al. 2018, ArXiv e-prints, arXiv:1805.05954

Cooper, A. P., Parry, O. H., Lowing, B., Cole, S., & Frenk, C. 2015, MNRAS, 454, 3185

Cooper, A. P., Cole, S., Frenk, C. S., et al. 2010, MNRAS, 406, 744

Cooper, M. C., Newman, J. A., Davis, M., Finkbeiner, D. P., & Gerke, B. F. 2012, spec2d: DEEP2 DEIMOS Spectral Pipeline, Astrophysics Source Code Library, ascl:1203.003

Cunningham, E. C., Deason, A. J., Rockosi, C. M., et al. 2019, ApJ, 876, 124

Cunningham, E. C., Deason, A. J., Guhathakurta, P., et al. 2016, ApJ, 820, 18

Cunningham, E. C., Deason, A. J., Sanderson, R. E., et al. 2018, arXiv e-prints, arXiv:1810.12201

Deason, A. J., Belokurov, V., & Evans, N. W. 2011, MNRAS, 416, 2903

Deason, A. J., Belokurov, V., Evans, N. W., & An, J. 2012, MNRAS: Letters, 424, L44

Deason, A. J., Belokurov, V., Evans, N. W., & Johnston, K. V. 2013a, ApJ, 763, 113

Deason, A. J., Belokurov, V., & Koposov, S. E. 2018a, ApJ, 852, 118

Deason, A. J., Belokurov, V., Koposov, S. E., & Lancaster, L. 2018b, ApJ, 862, L1

Deason, A. J., Belokurov, V., & Weisz, D. R. 2015, MNRAS, 448, L77

Deason, A. J., Van der Marel, R. P., Guhathakurta, P., Sohn, S. T., & Brown, T. M. 2013b, ApJ, 766, 24

Deason, A. J., Belokurov, V., Hamren, K. M., et al. 2014, MNRAS, 444, 3975

Dehnen, W., McLaughlin, D. E., & Sachania, J. 2006, MNRAS, 369, 1688

Eadie, G., Keller, B., & Harris, W. E. 2018, ApJ, 865, 72

Eadie, G. M., Springford, A., & Harris, W. E. 2017, ApJ, 835, 167

Escala, I., Wetzel, A., Kirby, E. N., et al. 2018, MNRAS, 474, 2194

Faber, S. M., Phillips, A. C., Kibrick, R. I., et al. 2003, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 4841, Instrument Design and Performance for Optical/Infrared Ground-based Telescopes, ed. M. Iye & A. F. M. Moorwood, 1657–1669

Font, A. S., McCarthy, I. G., Crain, R. A., et al. 2011, MNRAS, 416, 2802

Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, Publications of the

Astronomical Society of the Pacific, 125, 306

Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, PASP, 125, 306

Foreman-Mackey, D., Price-Whelan, A., Ryan, G., et al. 2014, triangle.py v0.1.1

Gaia Collaboration. 2018, VizieR Online Data Catalog, I/345

Gaia Collaboration, Prusti, T., de Bruijne, J. H. J., et al. 2016, A&A, 595, A1

Gaia Collaboration, Helmi, A., van Leeuwen, F., et al. 2018a, A&A, 616, A12

Gaia Collaboration, Brown, A. G. A., Vallenari, A., et al. 2018b, A&A, 616, A1

Garrison-Kimmel, S., Hopkins, P. F., Wetzel, A., et al. 2018, ArXiv e-prints, arXiv:1806.04143

Gilbert, K. M., Guhathakurta, P., Beaton, R. L., et al. 2012, ApJ, 760, 76

Gillessen, S., Eisenhauer, F., Trippe, S., et al. 2009, ApJ, 692, 1075

Gnedin, O. Y., Brown, W. R., Geller, M. J., & Kenyon, S. J. 2010, ApJ, 720, L108

Goodman, J., & Weare, J. 2010, Communications in Applied Mathematics and Computational Science, 5, 65

Grogin, N. A., Kocevski, D. D., Faber, S. M., et al. 2011, ApJS, 197, 35

Guo, Y., Ferguson, H. C., Giavalisco, M., et al. 2013, ApJS, 207, 24

Hattori, K., Valluri, M., Loebman, S. R., & Bell, E. F. 2017, ApJ, 841, 91

Helmi, A., Babusiaux, C., Koppelman, H. H., et al. 2018, ArXiv e-prints, arXiv:1806.06038

Hernitschek, N., Cohen, J. G., Rix, H.-W., et al. 2018, ApJ, 859, 31

Hopkins, P. F. 2015, MNRAS, 450, 53

Hopkins, P. F., Wetzel, A., Kereš, D., et al. 2018, MNRAS, 480, 800

Jeans, J. H. 1915, MNRAS, 76, 70

Jeffery, E. J., Smith, E., Brown, T. M., et al. 2011, AJ, 141, 171

Johnston, K. V., Bullock, J. S., Sharma, S., et al. 2008, ApJ, 689, 936

Johnston, K. V., Sheffield, A. A., Majewski, S. R., Sharma, S., & Rocha-Pinto, H. J. 2012, ApJ, 760, 95

Kafle, P. R., Sharma, S., Lewis, G. F., & Bland-Hawthorn, J. 2012, ApJ, 761, 98

Kafle, P. R., Sharma, S., Robotham, A. S. G., Elahi, P. J., & Driver, S. P. 2018, MNRAS, 475, 4434

Kalirai, J. S. 2012, Nature, 486, 90

Kallivayalil, N., van der Marel, R. P., & Alcock, C. 2006a, ApJ, 652, 1213

Kallivayalil, N., van der Marel, R. P., Alcock, C., et al. 2006b, ApJ, 638, 772

Kallivayalil, N., van der Marel, R. P., Besla, G., Anderson, J., & Alcock, C. 2013, ApJ, 764, 161

King, III, C., Brown, W. R., Geller, M. J., & Kenyon, S. J. 2015, ArXiv e-prints, arXiv:1506.05369

Kirby, E. N., Cohen, J. G., Guhathakurta, P., et al. 2013, ApJ, 779, 102

Koekemoer, A. M., Faber, S. M., Ferguson, H. C., et al. 2011, ApJS, 197, 36

Lancaster, L., Koposov, S. E., Belokurov, V., Evans, N. W., & Deason, A. J. 2019, MNRAS, 486, 378

Laporte, C. F. P., Johnston, K. V., Gómez, F. A., Garavito-Camargo, N., & Besla, G. 2018, MNRAS, 481, 286

Leauthaud, A., Massey, R., Kneib, J.-P., et al. 2007, ApJS, 172, 219

Li, T. S., Balbinot, E., Mondrik, N., et al. 2015, ArXiv e-prints, arXiv:1509.04296

Li, Z.-Z., Jing, Y. P., Qian, Y.-Z., Yuan, Z., & Zhao, D.-H. 2017, ApJ, 850, 116

Loebman, S. R., Valluri, M., Hattori, K., et al. 2018, ApJ, 853, 196

Mahmud, N., & Anderson, J. 2008, Publications of the Astronomical Society of the Pacific, 120, 907

Martin, N. F., Ibata, R. A., Rich, R. M., et al. 2014, ApJ, 787, 19

McCarthy, I. G., Font, A. S., Crain, R. A., et al. 2012, MNRAS, 420, 2245

McMillan, P. J. 2011, MNRAS, 414, 2446

Muzzin, A., Marchesini, D., Stefanon, M., et al. 2013, ApJS, 206, 8

Naab, T., Khochfar, S., & Burkert, A. 2006, ApJ, 636, L81

Nayyeri, H., Hemmati, S., Mobasher, B., et al. 2017, ApJS, 228, 7

Newman, J. A., Cooper, M. C., Davis, M., et al. 2013, The Astrophysical Journal Supplement Series, 208, 5

Nissen, P. E., & Schuster, W. J. 2010, A&A, 511, L10

Perryman, M. A. C., de Boer, K. S., Gilmore, G., et al. 2001, A&A, 369, 339

Price-Whelan, A. M., Johnston, K. V., Sheffield, A. A., Laporte, C. F. P., & Sesar, B. 2015, MNRAS, 452, 676

Rashkov, V., Pillepich, A., Deason, A. J., et al. 2013, ApJ, 773, L32

Reid, M. J., Menten, K. M., Zheng, X. W., et al. 2009, ApJ, 700, 137

Robin, A. C., Reylé, C., Derrière, S., & Picaud, S. 2003, A&A, 409, 523

Sales, L. V., Navarro, J. F., Abadi, M. G., & Steinmetz, M. 2007, MNRAS, 379, 1464

Sanderson, R. E., Garrison-Kimmel, S., Wetzel, A., et al. 2017, ArXiv e-prints, arXiv:1712.05808

Sanderson, R. E., Wetzel, A., Loebman, S., et al. 2018, ArXiv e-prints, arXiv:1806.10564

Schönrich, R. 2012, MNRAS, 427, 274

Schönrich, R., Binney, J., & Dehnen, W. 2010, MNRAS, 403, 1829

Sesar, B., Hernitschek, N., Dierickx, M. I. P., Fardal, M. A., & Rix, H.-W. 2017, ApJ, 844, L4

Sesar, B., Jurić, M., & Ivezić, Ž. 2011, ApJ, 731, 4

Sharma, S., Bland-Hawthorn, J., Johnston, K. V., & Binney, J. 2011, ApJ, 730, 3

Sheffield, A. A., Johnston, K. V., Majewski, S. R., et al. 2014, ApJ, 793, 62

Sheffield, A. A., Majewski, S. R., Johnston, K. V., et al. 2012, ApJ, 761, 161

Simon, J. D., & Geha, M. 2007, ApJ, 670, 313

Sirianni, M., Jee, M. J., Benítez, N., et al. 2005, PASP, 117, 1049

Sirko, E., Goodman, J., Knapp, G. R., et al. 2004, AJ, 127, 914

Slater, C. T., Nidever, D. L., Munn, J. A., Bell, E. F., & Majewski, S. R. 2016, ApJ, 832, 206

Smith, M. C., Evans, N. W., Belokurov, V., et al. 2009, MNRAS, 399, 1223

Sohn, S. T., Anderson, J., & van der Marel, R. P. 2012, ApJ, 753, 7

Sohn, S. T., Besla, G., van der Marel, R. P., et al. 2013, ApJ, 768, 139

Sohn, S. T., van der Marel, R. P., Carlin, J. L., et al. 2015, ApJ, 803, 56

Sohn, S. T., Watkins, L. L., Fardal, M. A., et al. 2018, ArXiv e-prints, arXiv:1804.01994

Sohn, S. T., Majewski, S. R., Muñoz, R. R., et al. 2007, ApJ, 663, 960

Sohn, S. T., van der Marel, R. P., Kallivayalil, N., et al. 2016, ApJ, 833, 235

Sohn, S. T., Patel, E., Besla, G., et al. 2017, ApJ, 849, 93

Stefanon, M., Yan, H., Mobasher, B., et al. 2017, ApJS, 229, 32

Tissera, P. B., White, S. D. M., & Scannapieco, C. 2012, MNRAS, 420, 255

Toloba, E., Li, B., Guhathakurta, P., et al. 2016, ApJ, 822, 51

VandenBerg, D. A., Bergbusch, P. A., & Dowler, P. D. 2006, ApJS, 162, 375

Wang, W., Han, J., Cole, S., et al. 2018, MNRAS, 476, 5669

Watkins, L. L., van der Marel, R. P., Sohn, S. T., & Evans, N. W. 2019, ApJ, 873, 118

Watkins, L. L., Evans, N. W., Belokurov, V., et al. 2009, MNRAS, 398, 1757

Wetzel, A. R., Hopkins, P. F., Kim, J.-h., et al. 2016, ApJ, 827, L23

Xue, X. X., Rix, H. W., Zhao, G., et al. 2008, ApJ, 684, 1143

Xue, X. X., Ma, Z., Rix, H. W., et al. 2016, VizieR Online Data Catalog, J/ApJ/784/170

Yesuf, H. M., Koo, D. C., Faber, S. M., et al. 2017, ApJ, 841, 83

Zolotov, A., Willman, B., Brooks, A. M., et al. 2009, ApJ, 702, 1058

Zolotov, A., Willman, B., Brooks, A. M., et al. 2010, ApJ, 721, 738