Title
The Elusive Three: On Black Holes, Dark Matter, and Neutrinos

ISBN
9798189291647

Author
Smith, Tyler Benjamin

Publication Date
2026-09-02

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA,
IRVINE

The Elusive Three: On Black Holes, Dark Matter, and Neutrinos

DISSERTATION

submitted in partial satisfaction of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

in Physics

by

Tyler Benjamin Smith

Dissertation Committee:
Professor Manoj Kaplinghat, Co-Chair
Professor Timothy Tait, Co-Chair
Professor Kevork Abazajian

2026

DEDICATION

To my daughters Autumn and Violet for always being a light to guide me.
To my wife Patrisha for her love and support.
To my mother for her strength.
In memoriam: Brian (Kaotik) Andresen (1982-2022)

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

First and foremost, I am eternally thankful to my family. I never imagined that I would pursue a PhD. My daughters Autumn and Violet gave me the motivation to re-enroll in school, to pursue something I could be happy doing while also providing for my family, and the motivation to endure and complete this PhD. To my wife Patrisha, you were the perfect partner through this and I couldn't have done it alone. Your support and love helped me get through this program and so much more. I'm so amazed that you were able to complete your Bachelor's while I worked on this PhD. I think we are what would qualify as a power couple. I think we absolutely crushed it and provided a great example for our girls who will go on to do great things.

Thank you to my advisors Manoj Kaplinghat and Tim Tait. Your guidance over the years have kept me on track and have helped shape me into the scientist I am today. I appreciate the time you spent working through complex topics with me, for hearing out my thoughts, and helping me develop an intuition for the problems we tackled over the years. I appreciate the support to go to conferences and summer schools where I could network and share my research with others. I am thankful for the many, *many*, letters of recommendation you submitted on my behalf over the years. These letters helped frame my work and accolades and assisted me in obtaining several fellowships and job offers.

Thank you to Kevork Abazajian, my third committee member. While we didn't get to work with each other, I am thankful for the discussions we had and your support. Thank you to Shirley Li, I've thoroughly enjoyed chatting with you over the years. Your insight into academia and research has been invaluable. Thank you to Franklin Dollar and Steph Sallum who were Physics and Astronomy Community Excellence (PACE) advisors. Additionally, thank you to all of the PACE team, Ryan Rice-Smith, Stephanie Urbano-Stawinski, Francisco Mercado, Tae Baxter, Jack Lubin, Helena Escudero, Jake Rudolph, Vidya Venkatesan, Mica Bacani, and Jay Krishnan. PACE was a big part of my PhD experience, from sitting in on Fire-up Fridays myself to leading splinter sessions at AAS on Mentorship and Leadership.

Thank you to the Gravitational Wave Paleontology Group at UCSD. In particular, thank you so much to Floor Broekgaarden. I appreciate you welcoming me into your group, for collaborating on the work that led to Chapter 3 with me, and your guidance and mentorship since I met you. Your genuine enthusiasm and kind heartedness are appreciated more than you know. Thanks Shasha Arani, Sasha Levina, Manasvini Komandur, Steffani Grondin, Kyle Rocha, and everyone else in the group for the welcoming environment, the friendliness, the comments, and feedback.

Many thanks to the Family Resource Center (FRC) here at UCI. Betsy Latimore, your support and love for my family is deeply appreciated. Camila Haro, Sam Lopez, Jenny Iaconelli, Miki Hughes, and everyone part of the FRC you all made it so fun to attend FRC events and to parent together. Further thanks to the many organizations I was a part of outside of PACE and the FRC, including SACNAS, PGC, and UNITY. I've had a blast throwing community events, conducting outreach, and partying with you all.

Thank you to my group members, Maya Silverman, Michael Ryan, Sophia Gad-Nasr, Kevin Andrade, Igor Palubski, Josh Hill, and all the undergrads! I appreciate all the feedback, discussions, hacking, practice talks, and friendship. Thank you to my office mates Pranav Premnath and Patricia Fofie, for keeping it real and letting me keep my wall a mess of posters. Thanks to Patrick Staudt for all the help with my bicycle, for riding and lifting with me, and for letting me cat-sit Hugo. Thank you Matt Dowicz for waking up at the crack of dawn to come lift with me at ARC, and for letting me lift at your place when the ARC was closed. Thanks to DaVante Cain for being a great neighbor and helping me truck stuff around. Thank you to Wataru Hayashi, Abby Bault, and Alec Hewitt for the lifting seshs. Thanks to the other particle theory and astrophysics students who supported me over the years: Jason Arakawa, Jessica Howard, Max Fieg, Loraine Ascencio, Johanna Paine, Cameron Moffett-Smith, Izzy Ginnett, India Bhalla-Ladd, Andrew Evans, Josh Kaler, Olti Myrtaj, Jason Baretz, Claire Rogers, Courtney Klein, Cannon Vogel, Rebecca Riley, Manuel Paul, and many more!

The work in Chapter 2 is a reprint of the work in *Tyler B. Smith and Manoj Kaplinghat 2026 ApJ 997 231.*

The work in Chapter 3 is as it was submitted to the Astrophysical Journal and is currently under review with expected publication later this year. The coauthors listed in this publication are Floor Broekgaarden, Sasha Levina, Amedeo Romagnolo, Manasvini Komandur, Melanie Santiago, and Kyle A. Rocha.

The work in Chapter 4 is an adaptation of the work that will be submitted to Astrophysical Journal later this year. The coauthors listed in this publication are Manoj Kaplinghat, Michael Ryan, and Maya Silverman.

The work in Chapter 5 is an adaptation of the work that will be submitted to Physical Review D later this year. The coauthors on this work are Shirley W. Li, John Beacom, Luke F. Roberts, and Francesco Capozzi.

The work in this thesis was in part supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. 1839285, the LANL–SoCalHub–UCI fellowship, and the ARCS fellowship.

VITA

Tyler Benjamin Smith

EDUCATION

Doctor of Philosophy in Physics **2026**
University of California, Irvine *Irvine, Ca*

Bachelor of Science in Physics **2020**
University of California, Riverside *Riverside, Ca*

RESEARCH EXPERIENCE

Graduate Research Assistant **2026**
University of California, Irvine *Irvine, California*

LANL-UCI-SoCalHub Fellow **2024–2025**
University of California, Irvine *Irvine, California*

NSF Graduate Research Fellow **2021–2024**
University of California, Irvine *Irvine, California*

UCI Departmental Fellow **2020–2021**
University of California, Irvine *Irvine, California*

TEACHING EXPERIENCE

Teaching Assistant **2025–2026**
University of California, Irvine *Irvine, Ca*

Adjunct Faculty **2024–2026**
Irvine Valley College *Irvine, Ca*

Adjunct Faculty **2024-2026**
Chaffey College *Rancho Cucamonga, Ca*

Adjunct Faculty **2024**
Orange Coast College *Costa Mesa, Ca*

FIRST AUTHOR REFEREED JOURNAL PUBLICATIONS

Massquerade: Impacts of Mass Ratio Reversals on Binary Black Hole Merger Rates and Mass Distributions **2026**
ApJ (In Review)

Isolated Binary Black Hole Formation and Merger Rates from Galaxy Evolution **2026**
ApJ

IN PREPARATION

Self-Interacting Dark Matter in the First Billion Years **2026**
Submission planned for September to ApJ

Measuring Energy and Deleptonization of the Next Galactic Supernova Without Spectral Assumptions **2026**
Submission planned for September to PRD

Other Works

How Common Are Common Envelopes? Quantifying Their Role in Forming Gravitational-Wave Sources **2026**
arXiv:2606.0532

Physics and Astronomy Community Excellence: Structured Near-Peer Mentorship in Doctoral Degree Programs **2026**
AEJ

Conserved Currents are not Anomaly Safe **2024**
arXiv:2401.02483

Reimagining the Astronomy PhD for the 21st Century **2023**
AAS

Dark Sector Studies with Neutrino Beams **2022**
arXiv:2207.06898

The Forward Physics Facility: Sites, experiments, and physics potential **2021**
arXiv:2109.10905

ABSTRACT OF THE DISSERTATION

The Elusive Three: On Black Holes, Dark Matter, and Neutrinos

By

Tyler Benjamin Smith

Doctor of Philosophy in Physics

University of California, Irvine, 2026

Professor Manoj Kaplinghat, Co-Chair

Professor Timothy Tait, Co-Chair

This dissertation investigates three of the most elusive phenomena in modern astrophysics: black holes, dark matter, and neutrinos, studied through theoretical modeling and large-scale numerical simulation to confront a new generation of gravitational-wave, electromagnetic, and neutrino observations.

The first part addresses the astrophysical population of merging binary black holes (BBHs) detected by LIGO-Virgo-KAGRA. Using the population-synthesis code SEVN combined with a galaxy-evolution framework spanning cosmic star formation, metallicity, and galaxy stellar mass, I construct a model for the volumetric BBH merger rate density and its dependence on primary mass, secondary mass, and mass ratio, finding that the observed primary-mass distribution requires either a top-heavy initial mass function in low-metallicity dwarf galaxies or a substantial dynamical-capture contribution above $\sim 30\,\mathrm{M}_{\odot}$. I then examine mass ratio reversal (MRR), binaries in which the initially less massive star forms the more massive compact object, across the SEVN and COMPAS codes. Both codes show that MRR leaves a distinct imprint on the merger-rate landscape, though the size and character of that imprint depends on the treatment of mass transfer and supernova physics. This makes clear that the primary-mass distribution inferred by LVK is not a direct tracer of the initially more

massive stars, but a superposition of physically distinct evolutionary populations, an effect that must be accounted for to robustly connect gravitational-wave observations to massive binary evolution.

The second part turns to the nature of dark matter. Motivated by JWST's discovery of unexpectedly massive, rapidly assembled galaxies in the early Universe, I use cosmological N-body simulations to test whether self-interacting dark matter (SIDM) can enhance early star formation relative to cold dark matter (CDM). I find that SIDM leaves the halo mass function unchanged but produces systematically rounder halos and suppressed subhalo survival by $z \approx 6$, with the total accretion rate onto hosts remaining comparable to CDM — indicating that self-interactions primarily enhance post-infall subhalo disruption rather than large-scale structure growth. Using TNG50-1, I estimate that this disruption liberates gas that could plausibly boost star-formation efficiency. I further show that gravothermal collapse, though analytically expected for a sizable portion of the simulated population, is not realized in these simulations, an absence attributable to the merger-dominated environment of early halo assembly.

The third part develops a model-independent framework for extracting the total emitted energy and net deleptonization of the next Galactic core-collapse supernova from its neutrino signal, using B-spline spectral unfolding applied jointly across Super-Kamiokande, DUNE, and JUNO. I show that this approach recovers the total emitted energy with percent-level precision without assuming a parametric spectral shape, that a nonzero net deleptonization can be robustly established despite its larger reconstruction uncertainty, that DUNE's dedicated ν_e channel is indispensable to the reconstruction, and that the resulting energy measurement is precise enough to serve as a direct, model-independent probe of the proto-neutron star's mass.

Together, these studies demonstrate how upcoming and current observational facilities, spanning gravitational waves, deep imaging, and neutrino detection, can be used to test and constrain fundamental physics across some of the most extreme environments in the Universe.

Chapter 1

Introduction

For millennia humanity has studied the cosmos, beginning with the charting of the paths of the Sun/Moon and cataloging of stars. As time progressed, so did our understanding of the Universe and the mathematical formalisms we use to describe it. The instruments with which we observe and study the Universe have also drastically improved over the centuries. While our ancestors only had light at their disposal to study the Universe, we exist at a unique time in Human history with the advent of gravitational wave (GW) and neutrino observatories which allow us to 'see' beyond the electromagnetic spectrum. We sit at the forefront of the era of *Multi-Messenger Astrophysics*.

In addition to these new tools to study the cosmos, traditional light-based telescopes have improved dramatically. Decades of work have led to a vast array of observatories and experiments that have come online recently and will continue to do so over the coming decades. Never before in Human history have we had such a wealth of data to compare our theoretical models against. I am grateful for the countless hours of work that have brought the field into such a data-rich epoch; each new dataset makes it an exciting time to be a theorist.

This thesis is largely centered on studying the Universe through these various vantage points, in particular, I study what I call the *elusive three*: black holes, dark matter, and neutrinos. Each is a phenomenon probed via different physics and cutting-edge experiments — some already online and taking data, others yet to come — covering a vast range of wavelengths, frequencies, and energies. This includes electromagnetic facilities spanning wavelength and survey scale, e.g. JWST, Roman, Euclid, DESI, the Vera C. Rubin Observatory, and the upcoming Square Kilometre Array, alongside the gravitational-wave and neutrino observatories discussed below.

1.1 Black Holes

Black holes have been hypothesized for over a century, it took Karl Schwarzschild less than two months after Einstein published his theory of General Relativity to obtain a solution which form the basis for black holes (Schwarzschild, 1916). It took only slightly longer for the discovery of gravitational waves as a solution to the field equations (Einstein, 1918). Around a century later, the until then hypothetical gravitational waves, were observed for the first time in the LIGO detectors (Abbott et al., 2016). The ensuing decade has seen gravitational wave observations grow from this first detection to over 250 confirmed detections encompassed in the fifth gravitational-wave transient catalog (GWTC-5) (The LIGO Scientific Collaboration et al., 2026a,b,c).

In the frequency band in which the LIGO-Virgo-KAGRA (LVK) collaboration observes, gravitational waves are sourced via the merger of two compact objects, neutron star – neutron star (NSNS) mergers, neutron star – black hole (NSBH), or binary black hole (BBH) mergers. The overwhelming majority of confirmed GW events originate from the final source, BBH mergers. These BBH mergers and their GW signatures allow us to infer the astrophysical population of merging BBHs. From this, we can then attempt to understand the nature of binary stars and stellar evolution. On the flip-side, and the approach taken in Chapter 2 and Chapter 3, one can start from models of stellar physics, binary evolution, and galactic/cosmological evolution and construct the underlying astrophysical population of BBHs. In the rest of this introduction I will detail the ingredients needed to theoretically model the astrophysical merger rate density of binary black holes.

1.1.1 Stellar Evolution

Stars start their lives on the main sequence when they begin burning hydrogen, this is referred to as the zero age main sequence (ZAMS). For massive stars ($\gtrsim 8\,\mathrm{M}_{\odot}$), the progenitors of compact objects, they progress through their lives through successive burning stages beginning with hydrogen burning into helium, followed by helium burning, and so on up through silicon burning (Hoyle, 1946; Burbidge et al., 1957; Cameron, 1957). This burning allows them to produce elements beyond those created in the early Universe, i.e. hydrogen, helium, and lithium (deuterium — an isotope of hydrogen — is also produced copiously during Big Bang nucleosynthesis) (Fields et al., 2014).

Upon exhaustion of their nuclear fuel, massive stars end their lives via supernova explosions collapsing into either neutron stars ($8 \lesssim \mathrm{M}_{\mathrm{ZAMS}}/\mathrm{M}_{\odot} \lesssim 20$) or black holes ($\mathrm{M}_{\mathrm{ZAMS}}/\mathrm{M}_{\odot} \gtrsim 20$) (Heger et al., 2003). Heavier elements, beyond those created during stellar evolution or supernovae explosions, include the r-process elements (see Cowan et al. (2021) for a review). Gravitational wave mergers, in particular NSNS mergers, are now understood to be a dominant site of r-process enrichment in the Universe (Lattimer & Schramm, 1974; Symbalisty & Schramm, 1982; Chen et al., 2024). Confirmation of r-process nucleosynthesis, in the context of compact object mergers, was observed in the multi-messenger event corresponding to the NSNS merger GW170817 (Abbott et al., 2017a,b) and its electromagnetic follow ups (see Metzger (2017) for an overview). This highlights the amazing physics we can deduce from the incoming wealth of data!

The thresholds stated above for NS or BH formation, are not hard and fast, rather they are estimates which vary based on properties such as metallicity, supernova engine, and winds prescription. These properties are also responsible for the mapping between the ZAMS mass and the corresponding remnant mass. I will discuss them in turn.

Metallicity

Metallicity, Z, quantifies the mass fraction of a star's material in elements heavier than helium. Metallicity is intricately related to the environment in which a stellar population forms — the gas content of its host galaxy. We will see that physical processes, such as stellar winds, are highly dependent on metallicity.

Supernova Engine

The supernova engine details how a supernova process unfolds, I consider two types in this thesis: delayed or rapid explosions (Fryer et al., 2012). The difference between the two is the time at which the SN explosion occurs post bounce, for the rapid engine this is < 250ms and for the delayed it can long exceed this value. In the population synthesis codes used throughout, the rapid supernova engine results in a mass gap between neutron stars and black holes between 3–5 $M_{\odot}$. This gap was physically motivated based on the lack of observations in this mass regime, however GW observations and EM observations have since put the idea of a *lower* mass gap to rest (Abbott et al., 2020, 2021a,b; LIGO Scientific Collaboration et al., 2025; The LIGO Scientific Collaboration et al., 2026a) (a small discussion on this is included in Chapter 2).

Stellar Winds

Stellar winds have large effects on stellar evolution (Meynet et al., 1994), in particular they greatly affect how much mass a star retains to form a compact object, but they also can alter stellar evolutionary tracks which are the backbone of modern population synthesis codes (codes used to simulate stellar and binary evolution). Here I focus on the effects of mass loss from massive O/B-type stars. In particular, stellar winds are line driven (Lucy & Solomon, 1970): metal ions in the stellar wind, via their resonance spectral lines, absorb and re-emit UV photons, the momentum transfer from the absorption is radially outward, while the re-emission is isotropic and averages out; this creates an outward propagating stellar wind. This wind continues to accelerate outward according to CAK theory (Castor et al., 1975; Vink, 2025): as the material moves outward, it sees Doppler-shifted UV radiation, allowing it to resonantly absorb previously inaccessible photons from the photosphere and continue accelerating. If the resulting radiative acceleration exceeds the star's inward gravitational acceleration, the wind can be driven to escape velocity, producing a steady outflow of mass. The strength of this line-driven mass loss depends sensitively on stellar luminosity and metallicity (Abbott, 1982; Pauldrach et al., 1986; Kudritzki et al., 1987; Romagnolo et al., 2026a), since the wind-driving opacity comes from resonance transitions of metal ions, then fewer metals means fewer driving lines and weaker winds.

1.1.2 Binary Physics

For single stellar evolution, there is a clear, well-defined mapping between initial stellar parameters (e.g. mass, metallicity, SN engine) and final remnant mass. Stars, however, are typically found in binaries and triples (Sana et al., 2012; Moe & Di Stefano, 2017) and the mapping deviates significantly from the single stellar evolution (SSE) prediction (see Figure 2.2 in Chapter 2). These changes arise from the initial parameters such as masses, orbital separation, and eccentricity, but also from the binary interactions, e.g. mass transfer, common envelope evolution, and supernova natal kicks. I briefly discuss these interactions below.

Mass Transfer

Mass transfer between two stars proceeds via one (or both) stars filling their Roche lobes, a tear-shaped volume element contained with an equipotential surface of the star's gravity (Hurley et al., 2002). Once this volume is filled the outer gas from the more massive donor star can transfer to the companion via the Lagrangian point joining their Roche lobes, termed Roche-lobe overflow (RLO). This mass can be accreted by the accretor (smaller) star, or if the quantity of gas is too large to be accreted the gas can engulf the accretor's Roche lobe leading to a common envelope (discussed below). In addition to RLO, stars can accrete stellar material which is shed by their companion's stellar winds via Bondi-Hoyle-Lyttleton accretion (Hoyle & Lyttleton, 1939; Bondi, 1952). This form of mass transfer is non-conservative and can strip the system of angular momentum. Mass transfer can greatly alter the mass of both the donor and accretor star, with the donor losing mass and the accretor gaining mass. This alters the ZAMS-to-remnant mass relationship and can in some cases lead to the accretor star becoming the more massive star or compact object via mass-ratio reversal which is the main topic of Chapter 3.

Common Envelope

As mentioned in Section 1.1.2, a binary system in which both Roche lobes are filled with gas is called a common envelope. In order for stars to survive their evolution to end in a binary compact object system, they need to start sufficiently far apart such that both stars have room to expand into their giant phase without causing a stellar merger. However, the large separation required is not able to be closed, via gravitational wave emission alone, in a Hubble time. Common envelope is a solution to this problem, see Mandel & Farmer (2022) and references therein. Energy is stored within the envelope and upon ejection angular momentum is stripped from the system, this leads to a tighter orbit in which GW emission can then lead to a merger.

Supernova natal kicks

Supernova explosions are not perfectly symmetrical, this asymmetry can lead to compact objects receiving a high velocity natal kick at birth. The distribution of kicks is typically assumed to follow a Maxwellian distribution with 1D root-mean-square dispersion σ. In the literature, the typical dispersion assumed is 265 km/s derived in Hobbs et al. (2005). However, a missing Jacobian factor accounting for their logarithmic bin sizes was discovered by Disberg & Mandel (2025). Furthermore, Disberg & Mandel (2025) found that a log-normal distribution — peaking at 150–200 km/s, notably lower than the Hobbs et al. (2005) value — is also consistent with the data alongside a single/double Maxwellian. It is important to model these kicks properly when computing binary compact object merger populations, as they directly set the fate of the binary. A sufficiently large kick can unbind the system entirely, while a surviving binary's post-kick orbital separation and eccentricity which can be driven either tighter or wider than before the explosion. In the latter case, this sets the GW inspiral timescale (Peters, 1964), and hence whether and how quickly the system merges (Kalogera, 2000; Hurley et al., 2002).

1.1.3 Cosmological Integration

The stellar and binary physics is important for system to system merger effects, but the population of stars which form and contribute to the local merger rate density also depend on the galactic environment the stars are formed in and their cosmological evolution. Taken together, the galactic and cosmological ingredients are referred to as cosmic integration. The important ingredients in cosmic integration are discussed below.

Initial Mass Function

The initial mass function (IMF), $\xi(M)$, details the number of stars formed per unit mass interval as a function of stellar mass. Typically, the IMF follows a power-law distribution of the form $M^{-\alpha}$, where α dictates how steep the power-law is. Throughout this thesis I follow the broken power-law IMF from Kroupa (2001), which in the high-mass regime ($M > 0.5\,\mathrm{M}_{\odot}$) finds $\alpha = 2.3$ in agreement with the results from Salpeter (1955). In Chapter 2 we develop a functional form for the IMF which allows for a top-heavy IMF for more massive stars, based on the observation of top-heavy IMFs in dwarf galaxies and globular clusters (Marks et al., 2012; Geha et al., 2013; Gennaro et al., 2018; Weatherford et al., 2021), which brings our calculated BBH mass distribution into better agreement with the inferred distribution based on GW observations. While often assumed to be universal, this assumption has been challenged by observational evidence for environment-dependent variation (e.g., Bastian et al., 2010; Zhang et al., 2018), further motivating the top-heavy extension developed in Chapter 2.

Galaxy Stellar Mass Function

Similar to the IMF, the galaxy stellar mass function (GSMF), $\phi(M_\star, z)$, describes the number density of galaxies per unit comoving volume, per dex in stellar mass, as a function of stellar mass and redshift. The typical functional form for the GSMF follows a single- or double-Schechter fit (e.g., Baldry et al., 2012). A lot of previous work was done at only local redshifts, however recent measurements have allowed us to constrain the GSMF down to redshifts of $z \sim 7.5$ (e.g., Weaver et al., 2023). Galaxies can be difficult to observe especially when they are small, as such, most GSMFs only provide observational data down to stellar masses of $10^8\,\mathrm{M}_\odot$ in the local Universe, the limit worsens at higher redshift ($z \approx 3$) raising to $10^{9-10}\,\mathrm{M}_\odot$. The GSMF is important because the metallicity of the gas is related to the stellar mass of the galaxy, as will be discussed in the next section. Knowing how many small galaxies, typically with lower metallicity, there is at a given redshift is important for the quantification of the merging BBHs from that galaxy.

Galaxy Mass – Metallicity Relation

The galaxy mass–metallicity relation (MZR) describes the relationship between galaxy stellar mass and gas-phase oxygen abundance for star-forming galaxies (e.g. Tremonti et al., 2004). Recently, the MZR has been constrained up to redshifts as high as $z \approx 10$ based on JWST measurements (Nakajima et al., 2023; Curti et al., 2024), confirming that galaxies are systematically more metal-poor, for fixed stellar mass, at high redshift. In the local

Universe, the MZR has been constrained in galaxies with stellar mass as low as $\sim 10^6\,\mathrm{M}_\odot$ (Lee et al., 2006). The MZR is essential for cosmic integration because gas-phase metallicity sets the metallicity of the stars forming from that gas, a key input for population synthesis calculations of compact-object formation and merger. Combined with the GSMF, which gives the number density of galaxies at each stellar mass and redshift, the MZR allows the metallicity-dependent star-formation history to be computed self-consistently across cosmic time.

Since observational MZR constraints are limited to discrete mass and redshift bins, cosmological simulations offer a complementary approach: constructing a continuous, redshift-evolving, mass-dependent functional form for the MZR, calibrated against the observations described above. In Chapter 2, I adopt the relationship derived from the FIRE simulations by Ma et al. (2016).

Star-Forming Main Sequence and Star Formation Rate Density

The star-forming main-sequence (SFMS) is a correlation between the star formation rate in a galaxy and its stellar mass. Characterizing the evolution of the SFMS's normalization, slope, and scatter with redshift has provided insight into the physical processes governing star formation across cosmic time (e.g., Brinchmann et al., 2004; Speagle et al., 2014; Popesso et al., 2023).The SFMS is often parameterized as a single power law, though several studies find a flattening at high stellar mass, attributed to the onset of quenching (e.g., Whitaker et al., 2014; Lee et al., 2015; Popesso et al., 2019). The convolution of the SFMS with the GSMF over stellar mass at a given redshift yields the star-formation rate density (SFRD) — the total mass of stars formed per unit comoving volume per unit time. This quantifies how much star formation is occurring at each epoch and, combined with the metallicity-dependent star formation history described above, sets which stars are available to form the compact-object binaries that merge at later times, as described in Chapter 2.

1.2 Dark Matter

Evidence for dark matter was first detailed in the works of Jan Oort (Oort, 1932) and Fritz Zwicky (Zwicky, 1933). Oort found that the observed velocities of stars near the Sun required mass beyond what was directly visible, and Zwicky reached the same conclusion from the velocity dispersion of galaxies in the Coma cluster. While Oort's observed discrepancy was later revised with improved measurements of the galactic disk (Kuijken & Gilmore, 1989), Zwicky's cluster-scale result has been repeatedly confirmed and remains foundational for the modern case for dark matter.

Following these earlier observations, evidence for dark matter was reinforced in the rotation curves of spiral galaxies by Rubin & Ford (1970), in which the rotational velocity of the galaxies was found to be constant at large radii rather than falling off as predicted by Kepler's laws (King, 1963). Similar observations were later carried out for dwarf galaxies where the same features were observed (e.g., Carignan & Freeman, 1988). On the galactic scale, further evidence for DM has been observed in galaxy-galaxy lensing (Brainerd et al., 1996; Hoekstra et al., 2004; Minor et al., 2017).

Further evidence is apparent in the growth of large-scale structure (LSS) observed in galaxy surveys today. LSS requires density perturbations to begin growing well before recombination; because baryons remain tightly coupled to photons until this epoch, and radiation pressure suppresses the growth of these coupled perturbations, a decoupled, non-baryonic matter component is required to seed structure early enough to match the clustering of galaxies observed today (Peebles, 1982; Blumenthal et al., 1984), consistent with the matter power spectrum measured by modern galaxy surveys (Tegmark et al., 2004; Cole et al., 2005; Adame et al., 2025). Independently, the cosmic microwave background's power spectrum requires a DM composition of $\sim 25\%$ in order to get the relative heights of the acoustic peaks correct (Hinshaw et al., 2013; Hou et al., 2014; Aghanim et al., 2020).

The growth of LSS and the CMB, discussed above, already provide strong evidence that the missing mass behaves as a distinct, non-baryonic particle component rather than simply a modification of gravity. In addition to these, one of the most direct and visually compelling pieces of evidence for this particle interpretation comes from the Bullet Cluster, a pair of colliding galaxy clusters whose gravitational lensing map is spatially offset from their baryonic matter (Clowe et al., 2006). Dark matter, assumed collisionless, does not interact with itself and thus passes through the incoming cluster largely undisturbed. Baryonic matter, on the other hand, interacts with itself, producing ram pressure and shock heating that slows the gas during the collision. This produces a clear spatial separation between the lensing-inferred mass and the baryonic gas, which modified theories of gravity cannot readily reproduce without invoking a DM-like component of their own (Cirelli et al., 2024).

Despite this evidence, the nature of DM remains a mystery: we currently only have evidence for its gravitational interaction, and we do not know if it couples to the Standard Model. Depending on the particle nature of dark matter we expect different observational properties, e.g. whether dark matter halos are densely packed in their inner regions or whether their inner density is relatively constant. Galaxies, and the dark matter halos (gravitationally bound clumps of dark matter) they inhabit, remain ideal testing grounds to understand the nature of DM: their dynamics are gravitationally dominated by dark matter, so tracers such as stellar and gas kinematics offer a direct probe of the underlying DM distribution, particularly in the dense central regions most sensitive to its particle nature.

1.2.1 Collisionless Cold Dark Matter

The largely accepted model throughout the 20th century was that of collisionless cold dark matter, in other words, non-relativistic, non-luminous (or non-baryonic) matter which couples only to gravity (Peebles, 1982, 1984; Blumenthal et al., 1984; Bahcall et al., 1999). It is this model of dark matter which is foundational to the Standard Model of Cosmology, $\Lambda_{\rm CDM}$, and it had much success predicting the large-scale structure of the Universe (Springel et al., 2006). $\Lambda_{\rm CDM}$ works exceptionally well more broadly, e.g. explaining both the cosmic microwave background (Aghanim et al., 2020) and baryon acoustic oscillations (Adame et al., 2025) (see Frenk & White (2012) for a review on the success of $\Lambda_{\rm CDM}$).

Small-Scale Challenges to CDM

Despite these successes at large scales, CDM has faced longstanding challenges on the scale of individual galaxies and their substructure. In the CDM paradigm, DM halos (clumps of DM) grow via hierarchical assembly of smaller halos (Press & Schechter, 1974; Lacey & Cole, 1993). It is these DM halos in which galaxies form, via the accretion of gas, its subsequent cooling, and fragmentation into stars (e.g., White & Rees, 1978; Wechsler & Tinker, 2018). This hierarchically assembled structure implies that subhalos can exist within the gravitational potential of a given host halo. It is at this level of small-scale structure in which CDM has historically had trouble accounting for various observations.

The first of these problems is the **core-cusp problem**, which states that halos comprised of CDM are expected to have a density profile which goes like $\rho \propto r^{-1}$ (Dubinski & Carlberg, 1991; Navarro et al., 1996b; Cole & Lacey, 1996) in the central region, i.e. are *cuspy*, while observations of rotation curves have been shown to be better fit by an isothermal profile (Kuzio de Naray et al., 2008; Oh et al., 2015; Li et al., 2020; Dehghani et al., 2020), i.e. constant in the inner region, or *cored*.

Another problem linked to the core-cusp problem in CDM is the **diversity problem**. CDM profiles have an expected universal density profile (Navarro et al., 1997) based on DMO simulations, however there exists a diversity in observed galactic rotation curves (Flores & Primack, 1994; de Blok et al., 2001; Kuzio de Naray et al., 2010), meaning that it is not simply core vs. cusp, but rather a variety of inner density slopes which exist in nature.

The third problem which arises is the **too-big-to-fail problem**, which states that CDM simulations produce subhalos that are overly dense compared to the Milky Way's satellites, and unless the Milky Way is a statistical fluke, we should observe their luminous counterparts, i.e. they cannot be completely dark given their high density (Boylan-Kolchin et al., 2011, 2012).

There is historically a fourth problem which CDM failed to explain, the so-called **missing-satellites problem**. The issue was that CDM predicts halos down to subsolar mass scales, and that, at the time, there were just too few observed MW satellites to match the CDM prediction. The issue was later alleviated as more satellites were discovered (e.g., Willman et al., 2005; Drlica-Wagner et al., 2015; Homma et al., 2016).

Baryonic Solutions to the Small-Scale Challenges

The tensions described above were originally identified by comparing dark-matter-only (DMO) simulations directly to observed galaxies, simulations that, by construction, omit baryonic processes. A natural question, then, is how much of each discrepancy reflects a genuine shortcoming of CDM itself, rather than the absence of baryonic physics in the theoretical prediction being tested against it. Simulations that incorporate baryonic physics have since alleviated several of these tensions, which I detail below (see also the review by Sales et al. (2022)).

In DMO simulations, satellites orbit within a purely collisionless host halo; once a central galaxy is included, its deeper gravitational potential strips mass from infalling satellites more aggressively than in the DMO case, lowering their present-day central masses below DMO expectations (Zolotov et al., 2012; Brooks & Zolotov, 2014). This effect, enhanced tidal stripping, offers a resolution to the too-big-to-fail problem.

Turning to the core-cusp problem, supernova feedback offers a resolution: repeated bursts of star formation followed by supernova-driven gas outflows produce fluctuations in the gravitational potential which transfer energy to the surrounding dark matter, transforming an initially cuspy profile into a shallower core (Navarro et al., 1996a; Governato et al., 2010; Pontzen & Governato, 2012). The efficiency of this process has been shown to depend on the stellar-to-halo mass ratio (Di Cintio et al., 2014) and also on the burstiness of star formation (Pontzen & Governato, 2012; Teyssier et al., 2013).

Baryonic physics also complicates the observational side of this picture. Non-circular gas motions, driven by bars, outflows, or general non-equilibrium kinematics in actively star-forming dwarfs, can bias the rotation velocity inferred from observed gas kinematics low relative to the true circular velocity, mimicking a cored profile even when the underlying halo remains cuspy (Rhee et al., 2004; Oman et al., 2019).

Despite three decades of progress along both of these lines, the diversity problem in particular remains unresolved: no baryonic feedback prescription has yet reproduced, self-consistently, both the full observed range of inner density slopes and the large-scale successes of ΛCDM (Sales et al., 2022). This motivates continued interest in solutions that modify the dark sector directly. Self-interacting dark matter (SIDM), discussed next, is one such possibility: rather than relying on baryonic feedback, self-interactions between dark matter particles redistribute energy within the halo directly, offering a physically distinct mechanism for core formation that, unlike baryonic feedback alone, aims to reproduce the diversity of observed rotation curves while preserving CDM's successes on larger scales.

1.2.2 Self-Interacting Dark Matter

Interestingly, the problems noted above can be alleviated if we relax the collisionless assumption on CDM and allow it to have self-interactions. Allowing for a richer dark sector, in which dark matter particles interact via a dark force-carrying mediator, is a natural extension given the complex nature of the Standard Model (e.g., Holdom, 1986a,b; Feng et al., 2009; Boddy et al., 2014). Simple elastic scattering between dark matter particles, via a dark force carrying mediator, termed self-interacting dark matter (SIDM) in the literature

Spergel & Steinhardt (2000), has been shown to alleviate these tensions, (see Tulin & Yu (2018); Adhikari et al. (2025) for a review).

Does the James Webb Space Telescope observe a tension with Λ_{CDM}?

Recent observations by the James Webb Space Telescope (JWST) have further brought the CDM paradigm under scrutiny. Examples include observations of galaxies with stellar masses exceeding $10^{10}\,\mathrm{M}_{\odot}$ within the first $\sim$ 500–700 Myr (Labbé et al., 2023; Xiao et al., 2024), unusually high star-formation efficiencies ($>$ 5–10%; Boylan-Kolchin 2023; Harikane et al. 2024; Pérez-González et al. 2025; Xiao et al. 2024), and sizes or luminosities unexpected at such early times (Wang et al., 2023; Carniani et al., 2024). These observations point toward rapid, efficient baryon assembly at high redshifts which puts a strain on the standard galaxy formation models in the Λ_{CDM} paradigm, and thus finding explanations consistent with this model have been sought. Some of the proposed remedies to reduce this tension include bursty star formation (Sun et al., 2023) or a top-heavy IMF (Trinca et al., 2024; Ventura et al., 2024; Hutter et al., 2025; Lu et al., 2025; Durrant et al., 2026). However, it is plausible that the issue lies with Λ_{CDM} itself. Similar to SIDM solutions to the aforementioned problems, perhaps it can also remedy this recent tension with CDM.

Beyond this specific question, however, the high-redshift Universe now opened up by JWST constitutes a qualitatively new regime for testing SIDM models altogether. Dark matter halos at these early epochs are younger and occupy earlier stages of gravothermal evolution than their well-studied low-redshift counterparts, so observables tied to this evolutionary state offer constraining power on SIDM parameter space that is largely inaccessible to established low-redshift probes such as dwarf galaxy rotation curves or cluster mergers. In this sense, JWST's high-redshift galaxy population represents not just a possible application of SIDM, but a new avenue for constraining it.

In Chapter 4 I pursue both threads: I explore whether SIDM can account for the rapid assembly of early galaxies, and I use the gravothermal evolution of SIDM halos at these epochs as a probe of SIDM models in its own right, applying it to *Little Red Dots* (LRDs), a puzzling population of red and compact objects at $z \gtrsim 4$ (Labbé et al., 2023; Kocevski et al., 2023; Harikane et al., 2023a; Kokorev et al., 2024; Greene et al., 2024; Williams et al., 2024; Matthee et al., 2024; Labbe et al., 2025; Akins et al., 2025; Kocevski et al., 2025).

1.2.3 Gravothermal Collapse

Central to the LRD explanation above is a generic — and, at sufficiently late times, unavoidable — consequence of self-interactions: gravothermal collapse. SIDM halos initially resemble their CDM counterparts, since scattering rates are low at the densities present at early times, preserving the same predictions for early-Universe physics such as the CMB (Dave et al., 2001). However, halos are not isolated, static systems: velocity dispersion rises outward from the halo center, so the dense central regions are initially cooler than their surroundings. Self-interactions conduct heat inward, heating and expanding the core and thermalizing it into an isothermal profile. Eventually, continued expansion leaves the core hotter than the surrounding envelope, reversing the temperature gradient and driving heat flow outward instead. The core then contracts, densities rise, and scatterings become more frequent, triggering a runaway loss of energy and the onset of gravothermal collapse (Lynden-Bell & Wood, 1968; Balberg et al., 2002).

Because halos form and begin their gravothermal evolution at different times, and because the timescale itself depends sensitively on a halo's density and velocity dispersion, halos observed at a common epoch can occupy very different points along the coring-to-collapse sequence — some still core-forming, others near maximum core, others already contracting. This spread naturally produces a range of central densities among halos of similar mass, offering an explanation for the diversity problem. It also addresses the core-cusp problem more directly than a single fixed-core prescription would: rather than predicting uniform cores across the halo population, gravothermal evolution predicts cores of varying size and depth set by each halo's individual formation history.

An Analytic Prescription for Collapse

These qualitative trends can be made quantitative: the timescale of gravothermal collapse for an isolated halo can be calculated following the analytic framework of Balberg et al. (2002) and Outmezguine et al. (2023) in which the time to collapse is given by:

$$t_{\rm coll} = \frac{1}{c_1}\,\frac{6-\alpha}{(3\alpha-2)-n\,(\alpha-2)}\,t_{c,0}\,, \tag{1.1}$$

where $\alpha = 4/\sqrt{\pi}$, $n = 4.5 \times 10^{-5}$, and $c_1 = 1.9 \times 10^{-3}$ are constants calibrated from gravothermal evolution models, and the numerical prefactor evaluates to approximately 400. The collisional timescale $t_{c,0}$ is:

$$t_{c,0} = \frac{2}{3\,\alpha\,C}\,(\sigma_0\,\rho_{c,0}\,v_{c,0})^{-1}\,, \tag{1.2}$$

where σ_0 is the cross section per unit DM particle mass. C is an order-one numerical constant fit to simulations and varies in the literature from $\sim$ 0.45–0.84 (Essig et al., 2019; Outmezguine et al., 2023; Nishikawa et al., 2020; Palubski et al., 2024; Koda & Shapiro, 2011; Yang et al., 2023). The core density $\rho_{c,0}$ and core velocity dispersion $v_{c,0}$ are evaluated at the onset of gravothermal evolution, and Outmezguine et al. (2023) show these can be derived directly from the NFW profile of the halo:

$$\rho_{c,0} = 2.4\,\rho_s\,, \qquad v_{c,0} = 0.64\,V_{\rm max}\,, \tag{1.3}$$

where ρ_s is the NFW scale density and $V_{\rm max}$ is the maximum circular velocity.

This analytic model is based on idealized, isolated halo simulations; the time to collapse can be affected by environmental factors (Dave et al., 2001), e.g. mergers (Silverman et al., 2026).

I will discuss how the analytic approximation holds up in a fully cosmological simulation in Chapter 4.

1.3 Supernova Neutrinos

Neutrinos were postulated by Wolfgang Pauli in 1930 (Brown, 1978) to preserve energy and momentum conservation in beta decay, and their existence was experimentally confirmed by Clyde Cowan and Frederick Reines (Cowan et al., 1956), work for which Reines was later awarded the Nobel Prize in Physics while at UC Irvine. Neutrinos remain a source of mystery: predicted massless in the Standard Model, they have since been measured to have mass via the discovery of neutrino oscillations (Fukuda et al., 1998; Ahmad et al., 2002). Understanding the origin of the neutrino mass and mixing properties remains a central focus of current particle physics, motivating a new generation of large-scale neutrino observatories. But neutrinos are not only a probe of fundamental physics — they are also a powerful astrophysical messenger. Chapter 5 focuses on one such application: supernova neutrinos, which shed light on the explosive death of massive stars.

1.3.1 Core-Collapse Supernova

Core-collapse supernova (CCSN), the death of a massive star, is one of the most energetic events in the Universe (Ott et al., 2013). Around 99% of this energy is released not via luminous photons but instead by neutrinos, roughly 10^{53} erg. Neutrinos are emitted across their entire flavor spectrum, ν_e, $\bar{\nu}_e$, and the heavier species ν_μ,$\bar{\nu}_\mu$, ν_τ, and $\bar{\nu}_\tau$ (collectively referred to as ν_x from here on). The importance of neutrinos in SN explosions was first pointed out by Colgate & White (1966) and Arnett (1966), as pointed out in a review by Janka (2017). To date, we have only confidently detected neutrinos from one supernovae explosion: SN1987A, located 50kpc away in the LMC. Unfortunately, there was no observed SN1987B, but the optimism was there. SN1987A resulted in ~ 24 detected neutrinos in the

Irvine Michigan Brookhaven (IMB), Kamiokande II, and Baksan experiments (Bionta et al., 1987; Hirata et al., 1987; Alekseev et al., 1987). Considering the estimated Galactic CCSN rate of a few per century (Li et al., 2011; Adams et al., 2013; Rozwadowska et al., 2021), and that the last confirmed Galactic event, Kepler's Supernova, occurred over four centuries ago, remaining prepared to extract the full multi-messenger and flavor information from the next such event is a pressing priority.

CCSNe inevitably lead to the formation of either a neutron star or a black hole. In the neutron star case, forming the neutron-rich core requires copious electron capture, $e^- + p \to n + \nu_e$, which preferentially produces ν_e. Detecting this flavor asymmetry, a probe of the deleptonization of the core, will be feasible with current and upcoming Earth-based detectors, as discussed in Chapter 5

Once a star begins silicon burning, it starts losing radiation and thermal pressure support against gravity, relying instead on electron degeneracy pressure to support the core. When the core exceeds the Chandrasekhar mass ($\sim 1.4\,\mathrm{M}_\odot$), it begins to collapse. Electron neutrinos are produced via electron capture on nuclei, but their free streaming is halted once the core reaches densities of order $10^{11}\,\mathrm{g/cm^3}$, at which point they become trapped. The collapse halts abruptly once nuclear densities are reached, causing the core to bounce and drive a shock wave outward. As this shock dissociates infalling nuclei into free nucleons, rapid electron capture on free protons produces a brief, intense burst of electron neutrinos known as the neutronization burst. The shock subsequently stalls as it loses energy to further dissociation and neutrino emission, while the proto-neutron star settles toward weak equilibrium, sustaining a high neutrino luminosity across all flavors. The stalled shock can be revived by neutrino heating — the partial reabsorption of neutrinos emitted from the proto-neutron star's neutrinosphere just behind the shock — resulting in a successful explosion (Bethe & Wilson, 1985).

If the shock cannot be revived through neutrino heating, e.g. if the mass accretion rate remains too high for the neutrino luminosity to overcome the ram pressure of infalling material (Burrows & Goshy, 1993), then the shock fails to reach a successful explosion, and the collapse continues. Without a mechanism to halt further accretion, the proto-neutron star continues to grow until it exceeds its maximum stable mass and collapses further to a black hole.

The specific flavor and timing signatures imprinted on the neutrino burst from infall, through the neutronization burst, to the post-bounce accretion and cooling phases, encode the explosion mechanism itself (Li et al., 2021), motivating a detailed look at what a full flavor, Earth-based detection of the next Galactic CCSN would actually look like.

1.3.2 Super-Kamiokande: Inverse Beta Decay

The best detection channel for a single flavor is inverse beta decay (IBD), $\bar{\nu}_e + p \to e^+ + n$, which proceeds on free protons. Super-Kamiokande, a water Cherenkov detector operating for over two decades with a fiducial mass of 22.5 kt (Suzuki, 2019), has the longest-running IBD sensitivity among current-generation detectors; its recently loaded gadolinium enables a delayed-coincidence neutron tag that cleanly separates IBD from other interaction channels. For a fiducial supernova at 10 kpc, Super-K alone is expected to detect $\sim 7 \times 10^3$ IBD events, the dominant channel for reconstructing the $\bar{\nu}_e$ spectrum. IBD is also available at JUNO and the future Hyper-Kamiokande detector, both of which provide complementary statistics and energy resolution, as discussed in Chapter 5.

1.3.3 DUNE: ν_e Charged-Current Absorption on Argon

The ν_e flavor is best probed via charged-current absorption on argon, $\nu_e + {}^{40}\mathrm{Ar} \to e^- + {}^{40}\mathrm{K}^*$, a channel uniquely available at the Deep Underground Neutrino Experiment (DUNE) (Abi et al., 2021, 2020). DUNE is a liquid argon time-projection chamber under construction at the Sanford Underground Research Facility, with a planned 40 kt fiducial mass across four detector modules. For a fiducial supernova at 10 kpc, DUNE is expected to detect $\sim 4 \times 10^3$ events in this channel. Because no other detector considered here isolates ν_e directly, DUNE plays an essential role in disentangling the ν_e contribution from the mixed-flavor signals seen elsewhere.

1.3.4 JUNO: Neutrino–Electron Elastic Scattering

The heavy-lepton flavors, ν_x, are the hardest to detect directly, since no charged-current channel exists for them at these energies. Neutrino–electron elastic scattering, $\nu + e^- \to \nu + e^-$, proceeds for all flavors via a combination of charged- and neutral-current amplitudes, making it the primary channel used to constrain ν_x. JUNO, a 20 kt liquid scintillator detector in Jiangmen, China (An et al., 2016; Abusleme et al., 2025a), offers excellent energy resolution and a low detection threshold for this channel. For a fiducial supernova at 10 kpc, JUNO is expected to detect $\sim 5 \times 10^2$ scattering events, roughly half attributable to ν_e, with the remainder split between $\bar{\nu}_e$ and ν_x, reflecting the larger charged-current cross section for ν_e. Because this channel mixes contributions from all flavors, extracting a clean ν_x measurement requires combining it with the flavor-specific IBD and DUNE channels described above, and is challenge I address in Chapter 5.

1.4 Outline of this Thesis

The three preceding sections have introduced the physical motivation for each of the topics addressed in this dissertation. The chapters that follow are organized around these three themes and are largely self-contained, each grounded in a distinct set of observational drivers and methodologies; while black holes, dark matter, and neutrinos do not share a single unifying physical mechanism, each chapter reflects a common approach: using theoretical and computational modeling to extract what upcoming or current observational facilities can reveal about fundamental physics.

Chapter 2 develops a model for the astrophysical population of merging binary black holes, combining stellar and binary evolution from the SEVN code with a galaxy-evolution framework spanning cosmic star formation history, metallicity, and galaxy stellar mass. I show how the resulting primary-mass, secondary-mass, and mass-ratio distributions compare to LIGO-Virgo-KAGRA observations, and what these comparisons imply for the initial mass function and formation channels of massive black holes.

Chapter 3 extends this population synthesis approach to examine mass ratio reversal, in which binary interactions cause the initially less massive star to form the more massive compact object. Comparing predictions from SEVN and COMPAS, I quantify how this process reshapes the observed BBH mass distribution and identify the evolutionary pathways responsible.

Chapter 4 shifts focus to the particle nature of dark matter, using cosmological N-body simulations to test whether self-interacting dark matter can help explain the unexpectedly massive, rapidly assembling galaxies observed by JWST at high redshift. I examine the imprint of self-interactions on halo shapes, subhalo survival, and gravothermal evolution, and assess the implications for early gas accretion and star formation.

Chapter 5 addresses the next Galactic core-collapse supernova, developing a model-independent framework for reconstructing the total emitted energy and net deleptonization of the explosion from its neutrino signal across Super-Kamiokande, DUNE, and JUNO. I show how this reconstruction connects directly to the mass and neutronization of the newly formed proto-neutron star.

Each chapter concludes with its own discussion of results and implications for future work.

Chapter 2

Isolated Binary Black Hole Formation and Merger Rates from Galaxy Evolution

2.1 Abstract

The LIGO-Virgo-KAGRA (LVK) collaboration has detected over 150 confirmed gravitational wave events through O4a. Binary black hole (BBH) systems represent the overwhelming majority of these observations. We construct a model for the population of the BBHs based on the distribution of metallicities in galaxies and state-of-the-art stellar evolution models implemented through the Stellar EVolution N-body (`SEVN`) code. We calculate the redshift evolution of the total merger rate of BBHs and the differential rates with respect to primary mass, secondary mass, and the mass ratio. We explore variations in the delay-

time distribution's (DTD) power-law index and show that it affects the total merger rate's spectral shape, but primarily acts as an amplitude shift on the differential rates. When comparing to the primary mass distribution, our results indicate that either the average IMF in dwarf galaxies must be top heavy, or most of the 30–40 $M_{\odot}$ black holes must be formed through a dynamical capture mechanism. For masses greater than about $50\,M_{\odot}$, the predicted number of BBH systems plummet to zero, revealing the well-known mass gap due to the pair instability mechanism and mass loss in binary systems.

2.2 Introduction

In 2015, the LIGO collaboration discovered the first direct evidence for gravitational waves (GWs) with the detection of event GW150914 Abbott et al. (2016). Currently observing run four (O4) is nearing completion, with over 150 confirmed GW events detected by the LIGO–Virgo–KAGRA (LVK) collaboration Abbott et al. (2019a); Abbott et al. (2021); Abbott et al. (2021a); Abac et al. (2025) up to O4a. Among the GW events are binary neutron star (BNS), neutron star–black hole (NS–BH), and binary black hole (BBH) mergers encompassing a wide mass spectrum.

These GW observations have revolutionized our understanding of compact objects. Prior to the first GW detections, observations of black holes in X-ray binaries were largely limited to masses of $\lesssim 20\,M_{\odot}$ Farr et al. (2011); Ozel et al. (2010) which was consistent with the maximum black hole mass calculated in Fryer & Kalogera (2001). While observational evidence pointed to an upper mass $\lesssim 30\,M_{\odot}$, theoretical modeling of low-metallicity systems asserted a higher threshold around $\sim 80\,M_{\odot}$ Belczynski et al. (2010a); Mapelli et al. (2009). Thus, the detection of GW150914 Abbott et al. (2016) which featured component masses of $36\,M_{\odot}$ and $29\,M_{\odot}$, was not a complete surprise Mandel & Farmer (2022), but still challenged existing assumptions primarily related to the lack of observation of such massive black holes and prompted a revision of black hole formation theories. Subsequent merger events have pushed beyond observational gaps, placing increasing pressure on theoretical models.

Two notable examples of changing perspectives are related to the so-called mass gaps in the compact object mass spectrum. The first mass gap, the *lower-mass gap*, occurs between the highest-mass neutron stars and the lowest-mass black holes. Neutron stars have long had a theoretical upper limit of $\sim 2.2\,\mathrm{M}_{\odot}$, termed the Tolman–Oppenheimer–Volkoff (TOV) limit Bombaci (1996); Kalogera & Baym (1996); Rezzolla et al. (2018). Likewise, black holes were previously thought to only form down to a lower limit of $\sim 5\,\mathrm{M}_{\odot}$ Farr et al. (2011), possibly resulting from core-collapse supernova physics Abbott et al. (2023). This gap was postulated due to the dearth of observations of black holes below this threshold. Gravitational wave events GW 190814 Abbott et al. (2020), GW 200115 Abbott et al. (2021a,b), and a first result of the O4 run Abac et al. (2024), directly challenge the lower-mass gap expectation. While traditional astrophysical processes can explain the presence of black holes in this range, studies have also shown that lower-mass gap black holes could originate from primordial black holes Afroz & Mukherjee (2024).

The second mass gap is predicted by the pair instability mechanism Fowler & Hoyle (1964); Barkat et al. (1967); Fryer et al. (2001); Belczynski et al. (2016); Woosley et al. (2007); Woosley (2017, 2019), termed the *pair-instability mass gap* or the *upper-mass gap*. When a star's core temperature reaches $\sim 10^9\,\mathrm{K}$ the high-energy photons gain sufficient kinetic energy to annihilate into electron-positron pairs. This conversion reduces the radiation pressure of the star inducing gravitational collapse, leading to ignition of explosive oxygen/silicon burning. For helium core masses $30 \lesssim M_{He}/\mathrm{M}_{\odot} \lesssim 60$ the core contracts and ignites burning, but it is not sufficient to completely disrupt the star. The star expands and cools to equilibrium, and this process can repeat leading to a series of pulsations lasting from a few hours to $\sim 10,000$ years Woosley (2017). For higher mass cores $60 \lesssim M_{He}/\mathrm{M}_{\odot} \lesssim 130$ the explosion disrupts the entire star, resulting in a pair-instability supernova with no remnant. For core masses exceeding this limit, $M_{He} \gtrsim 130\,\mathrm{M}_{\odot}$, the star collapses directly to a black hole as a result of the pair instability Heger et al. (2003).

That these mass gaps are populated according to LVK observations Abbott et al. (2020a,b); Abbott et al. (2021a, 2023); Abac et al. (2025) highlight the increasing discrepancy between current theoretical expectations and observations. There are currently > 150 total confirmed events in O4a, with more to come by the end of the fourth observing run. The new data will shed light on compact objects and may also present new challenges for understanding their nature and origin. In light of this, it is imperative to develop sophisticated models that detail the formation and evolutionary pathways of these compact objects, particularly in the context of mergers.

In this study, we specifically focus on binary black hole systems that form and evolve through the isolated channel. The isolated pathway consists of binary systems where the component stars interact with each other before merger but have negligible interaction with external pairs or the environment Tutukov et al. (1973); Bethe & Brown (1998); O'Shaughnessy et al. (2010); Dominik et al. (2013, 2015); Belczynski et al. (2016); Stevenson et al. (2017); Elbert et al. (2018); Baibhav et al. (2019); Spera et al. (2019); Santoliquido et al. (2020); Bouffanais et al. (2021); Broekgaarden & Berger (2021); García et al. (2021); Stevenson & Clarke (2022); Broekgaarden et al. (2022); Iorio et al. (2023a); Boesky et al. (2024a). In contrast, the dynamical channel involves interactions among single or multi-star systems within densely populated stellar environments, such as globular clusters, nuclear star clusters, or young star clusters Portegies Zwart & McMillan (2000); O'Leary et al. (2006); Banerjee et al. (2010); Farr et al. (2017); Mapelli et al. (2021); Santoliquido et al. (2020); Gerosa & Fishbach (2021); Sedda et al. (2023). In the dynamical channel, stars frequently engage in close encounters or exchanges, leading to a series of complex dynamical interactions.

It is common in the literature to adopt a framework based on the use of a population synthesis code coupled with galactic and cosmological relations such as the star formation rate (SFR), galaxy stellar mass function (GSMF), and mass-metallicity relation (MZR) O'Shaughnessy et al. (2010); Neijssel et al. (2019a); Santoliquido et al. (2020); Bouffanais et al. (2021); Broekgaarden et al. (2022); Broekgaarden & Berger (2021); Stevenson & Clarke (2022); Boesky et al. (2024a). We take a similar approach, utilizing the Stellar EVolution N-body code `SEVN` Spera et al. (2015); Spera & Mapelli (2017); Spera et al. (2019); Mapelli et al. (2020); Iorio et al. (2023a), which incorporates the PARSEC tracks Bressan et al. (2012); Tang et al. (2014); Chen et al. (2015); Marigo et al. (2017); Nguyen et al. (2022) as our choice of population synthesis model. We differ from previous studies by adopting an up-to-date, redshift-evolving GSMF based on observational constraints from the COSMOS survey Weaver et al. (2023), which provides star-forming measurements out to $z \sim 5.5$. Beyond this range, we extrapolate the observed trend. We note that Neijssel et al. (2019a) also consider an evolving GSMF, though theirs is derived from the EAGLE simulations Furlong et al. (2015). This choice showcases how the number density of binary black holes that merge within the age of the Universe evolves over redshift.

In conjunction with the GSMF, we account for the redshift evolution of the mass-metallicity relation (MZR) by adopting the results of cosmological zoom-in simulations from Ma et al. (2016). Their results remain consistent with observed MZR trends up to $z = 3$ from Tremonti et al. (2004); Gallazzi et al. (2005); Erb et al. (2006); Mannucci et al. (2009); Zahid et al. (2011); Kirby et al. (2013); Steidel et al. (2014); Sanders et al. (2015). This redshift range covers the bulk of current LVK detections, given the short stellar lifetimes and delay-time distributions of massive binary systems as discussed in detail below. Furthermore, we adopt the star-forming main sequence in Speagle et al. (2014) and constrain the normalization by the core-collapse supernova rate (CCSN) from the ZTF survey Perley et al. (2020).

Our goal is to show how our particular choices of galactic and cosmological parameterizations, grounded in empirical observations, can reproduce the total merger rate and component mass distributions inferred from the LVK collaboration. We highlight how the isolated channel alone cannot reconcile the LVK results, and what it would take from star formation and stellar physics to do so. These results further strengthen the argument for inclusion of the dynamical population, in particular at higher masses.

While our general approach is common, we further highlight how our study goes beyond previous works. Most of the works on isolated binary evolution cited above have focused primarily on the total merger rate density and its evolution with redshift. Though there have also been efforts to examine specific aspects of the black hole mass distributions. For example, Broekgaarden et al. (2022) investigated the chirp mass distribution inferred from gravitational wave observations, while Neijssel et al. (2019a) analyzed the total mass distribution, though without incorporating the effects of pair-instability supernovae. Their justification was that similar high-mass cutoffs can emerge from particular combinations of galaxy and stellar evolution relations. Additionally, Tanikawa et al. (2022) performed model comparisons against the cosmic merger rate and its mass derivatives across Pop I-III systems. However, unlike our work, Tanikawa et al. (2022) does not incorporate empirically constrained galaxy-based ingredients such as the evolving GSMF and MZR.

We extend this body of work by modeling not only the total merger rate, but also the differential merger rate density with respect to primary mass, secondary mass, and mass ratio. This enables a more detailed and physically grounded comparison with both current LVK data and future observations. Notably, the secondary mass spectrum remains relatively underexplored, though see recent work from Sadiq et al. (2024); Farah et al. (2024). We also briefly examine the impact of the DTD on the redshift evolution of the merger rate, finding that even modest variations can lead to noticeable differences. While this manuscript was under review, LVK released O4a, providing an updated and expanded catalog of GW events. We include comparisons with these most recent results in this work.

Our paper is structured as follows. In Section 2.3.1, we begin by discussing the simulations used to inform and constrain our analysis. Following this, we construct the population of black holes in Section 2.3.2. The calculation of the merger rate density is detailed in Section

2.3.4. In Section 2.4, we compare the results of our model to those of the LVK collaboration. We conclude in Section 2.5.

2.3 Methodology

2.3.1 Zero-Age Main Sequence to Compact Object Evolution with `SEVN` Simulations

The first ingredient to properly model the binary black hole population is stellar evolution and dynamics. To encapsulate the complex and broad physical processes involved, we use the stellar evolution code `SEVN v2.7.3` (Spera et al., 2015; Spera & Mapelli, 2017; Spera et al., 2019; Mapelli et al., 2020; Iorio et al., 2023a) to evolve stars from their zero-age main sequence (ZAMS) masses to their final remnant masses, with a focus on binary mergers. We adopt the default parameters unless noted otherwise, e.g. PARSEC tracks with overshooting parameter $\lambda = 0.5$, and refer the reader to Spera et al. (2019) and Iorio et al. (2023a) for further details on these prescriptions and the effects of modifying them.

We evolve $\sim 8 \times 10^7$ binary configurations per 20 distinct metallicities spanning $-4 \leq \log_{10}(Z/Z_\odot) \leq 0.2$. Progenitor masses for the primary star are uniformly sampled in the range $10 \leq M/\mathrm{M}_\odot \leq 200$ with a spacing of $1\,\mathrm{M}_\odot$. To obtain a representative sample of binary systems, following the sampling of $M_{1,\mathrm{ZAMS}}$ and $\log_{10} Z$, we sample the mass ratio, eccentricity, and period from distributions following the results in (Sana et al., 2012):

$$P(q_{ZAMS}) \propto q_{ZAMS}^{-0.1}, \qquad q_{ZAMS} = \frac{M_2}{M_1}, \qquad q_{ZAMS} \in [0, 1], \tag{2.1a}$$

$$P(\mathcal{P}) \propto \mathcal{P}^{-0.55}, \qquad \mathcal{P} = \log_{10}(\mathrm{P/day}), \qquad \mathcal{P} \in [0.15, 5.5], \tag{2.1b}$$

$$P(e) \propto e^{-0.42}, \qquad e \in [0, 0.9]. \tag{2.1c}$$

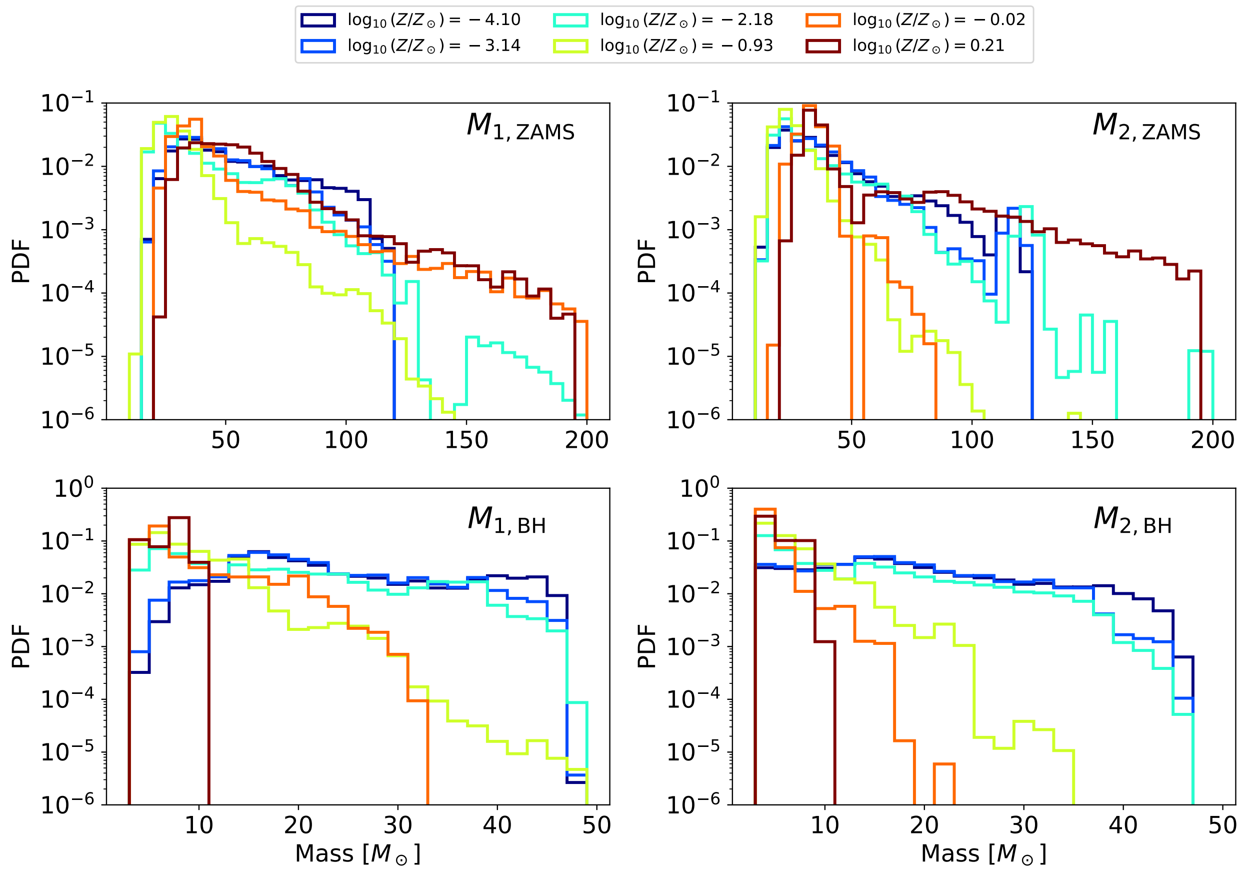

Figure 2.1: Probability density functions of stellar and remnant masses across metallicity. The top row shows the zero-age main sequence (ZAMS) masses of the primary (left; $M_{1,\mathrm{ZAMS}}$) and secondary (right; $M_{2,\mathrm{ZAMS}}$) stars in merging binary systems, while the bottom row shows the corresponding black hole masses, where $M_{1,\mathrm{BH}}$ denotes the more massive black hole in the binary. Each curve represents a different metallicity, as indicated by the color scheme, and all distributions are weighted by the initial mass function and normalized to highlight relative contributions.

In our model, we assume zero initial spins for the progenitor stars. This assumption is well justified, as the majority of the star's primordial angular momentum is lost during its evolution due to Roche-lobe overflow mass transfer and wind mass loss (Qin et al., 2018; Bavera et al., 2020). These studies show that as the star expands, its angular momentum is transferred to the outer layers, which are subsequently removed through the aforementioned processes. Specifically, Bavera et al. (2020) highlight that the initial angular momentum is mostly lost by the time a star becomes a helium star, and Qin et al. (2018) emphasize that most of the angular momentum a star might have had is lost during this expansion phase, making the initial spin largely irrelevant for the final spin of the black hole. By assuming zero initial spin we have simplified the model without sacrificing accuracy, as the black hole's spin is dominated by processes such as mass transfer, tidal synchronization, and accretion in the binary system's later evolutionary stages.

We restrict our analysis to binary systems containing two black holes that merge within a Hubble time ($\sim$ 14 Gyr). Figure 2.1 shows the distributions of both initial stellar masses (ZAMS) and final black hole (BH) masses for our population of binaries. Because we uniformly sample the primary stellar mass across the full range, the systems are weighted by the IMF to produce a physically meaningful distribution. We define the primary black hole mass, $M_{1,\mathrm{BH}}$, as the more massive BH in the binary. For consistency, we define its progenitor, $M_{1,\mathrm{ZAMS}}$, as the primary star that formed this black hole. Each panel in Figure 2.1 includes contributions from six different metallicity bins spanning our full range. The distributions are normalized per metallicity to highlight the contributions across the mass spectrum.

In the ZAMS distributions (top panels), we observe a decline in the contribution from very massive stars ($\gtrsim 100\,\mathrm{M}_\odot$) as metallicity decreases. This is because low-metallicity stars retain more mass and are more likely to enter the pair-instability regime, preventing black hole formation. In contrast, high-metallicity stars experience strong mass loss through stellar winds, which prevents them from reaching the pair-instability threshold, allowing them to contribute even at the highest initial masses. However, these high-metallicity stars only produce black holes with masses below $\sim 30\,\mathrm{M}_\odot$ due to extensive wind mass loss (Romagnolo et al., 2024). Low-metallicity stars, on the other hand, contribute more uniformly across the BH mass range and can form BHs up to $\sim 45\,\mathrm{M}_\odot$ in our simulations.

To further illustrate the initial-to-final mass relationship, we compare the initial stellar mass to the final black hole mass in Figure 2.2. This figure is inspired by Figure 5 in Spera et al. (2019), but differs in key aspects. They examine systems which end up in compact object binaries and find that the majority of binary systems follow the results from single stellar evolution (SSE), while we include the cut $t_{\rm merge} < t_{\rm Hubble}$ and find most of these systems do not track SSE. The red line indicates the SSE case in which each progenitor and metallicity combination produces a unique black hole mass in `SEVN`. In contrast, the binary stellar evolution (BSE) model exhibits a wide range of black hole masses for each progenitor mass-metallicity combination. Each subplot represents the BSE case via hexagonal binning, width of $\sim 2\,\mathrm{M}_{\odot}$, where the colorbar logarithmically represents the density. It is well known that merging BBHs have a greater efficiency at low metallicities as is apparent in our figure. Additionally, at intermediate metallicities ($-2 \lesssim \log_{10}(Z/Z_{\odot}) \lesssim -1$), mass transfer in binaries can produce BHs from progenitors in the mass range that would not do so under SSE. For low-metallicity systems there exists an island above the SSE line at $M_{\rm ZAMS} \sim 20\mathrm{M}_{\odot}$. These systems are predominantly populated by systems where mass ratio reversals takes place, i.e. the secondary progenitor becomes the primary black hole. Mass ratio reversals occur in $\sim 20\%$ of our merged systems. For the most metal-rich systems, the evolutionary path follows the SSE case, which is due to intense stellar winds disrupting binary systems and leaving only those that evolve in relative isolation.

We note that prescription for the pair-instability mechanism we use is termed M20 in `SEVN` following the fit from Spera & Mapelli (2017) to 1D hydrodynamical simulations from Woosley (2017). Under this prescription a star will undergo pulsational pair instability if the pre-supernova helium core mass falls between $32 \leq M_{\rm He}/\mathrm{M}_{\odot} \leq 64$ and a pair-instability supernova for $64 \leq M_{\rm He}/\mathrm{M}_{\odot} \leq 135$ (Iorio et al., 2023a). In the simulations from Woosley (2017) black holes were not found to form between $\sim 50 - 130\,\mathrm{M}_{\odot}$, however as is clear from Figure 2.2 we obtain a small number of black holes within this regime. These massive black holes form in the `SEVN` code as the remnants of low-metallicity progenitor stars of $\sim 100\,\mathrm{M}_{\odot}$, and avoid the pair instability mechanism during their late dredge-up phase as the mass of the helium core is reduced to just below the threshold required for PISN, $\sim 32\,\mathrm{M}_{\odot}$ (Iorio et al., 2023a). The star can then end its life with roughly the mass at the end of helium depletion phase, producing a massive BH that lies in the pair-instability mass gap. These black holes evolve in relative isolation, mimicking the SSE track, and are exceptionally rare. Thus, predictions for black holes with masses above $\sim 50\,\mathrm{M}_{\odot}$ should be interpreted with care.

It is clear that the mass distribution of surviving binary black hole masses is intricately shaped by the progenitor characteristics and the orbital configurations of each system, in agreement with Spera et al. (2019); Broekgaarden et al. (2022); Iorio et al. (2023a). To accurately capture this complex relationship, we introduce an efficiency parameter, $\varepsilon(M_{\mathrm{ZAMS}}, Z)$. This parameter represents the likelihood of a star, for a given progenitor mass and metallicity, to form a black hole of any mass, marginalized over the orbital parameters. This efficiency factor is defined as:

$$\varepsilon(M_{\mathrm{ZAMS}}, Z) = \frac{\mathcal{N}_{\mathrm{sur}}(M_{\mathrm{ZAMS}}, Z)}{\mathcal{N}_{\mathrm{sim}}(M_{\mathrm{ZAMS}}, Z)}\,, \tag{2.2}$$

where $\mathcal{N}_{\mathrm{sur}}(M_{\mathrm{ZAMS}}, Z)$ represents the number of primary binary black hole progenitors that successfully merge within a Hubble time as determined from our sampling, and $\mathcal{N}_{\mathrm{sim}}(M_{\mathrm{ZAMS}}, Z)$ denotes the total number of simulated progenitor configurations. The results from Equation 2.2 are shown in Figure 2.3. This figure demonstrates that stars with lower metallicity have a higher likelihood to form black hole remnants, consistent with findings in various works including, but not limited to Belczynski et al. (2010b); Mapelli et al. (2010); Stevenson et al. (2017); Chruslinska et al. (2019); Giacobbo & Mapelli (2018); Neijssel et al. (2019a); Spera et al. (2019); Broekgaarden et al. (2022); Schiebelbein-Zwack & Fishbach (2024). While the efficiency factor for creating black holes in general is higher at lower metallicities, the progenitors are also more likely to undergo pair-instability supernova as indicated by the dip

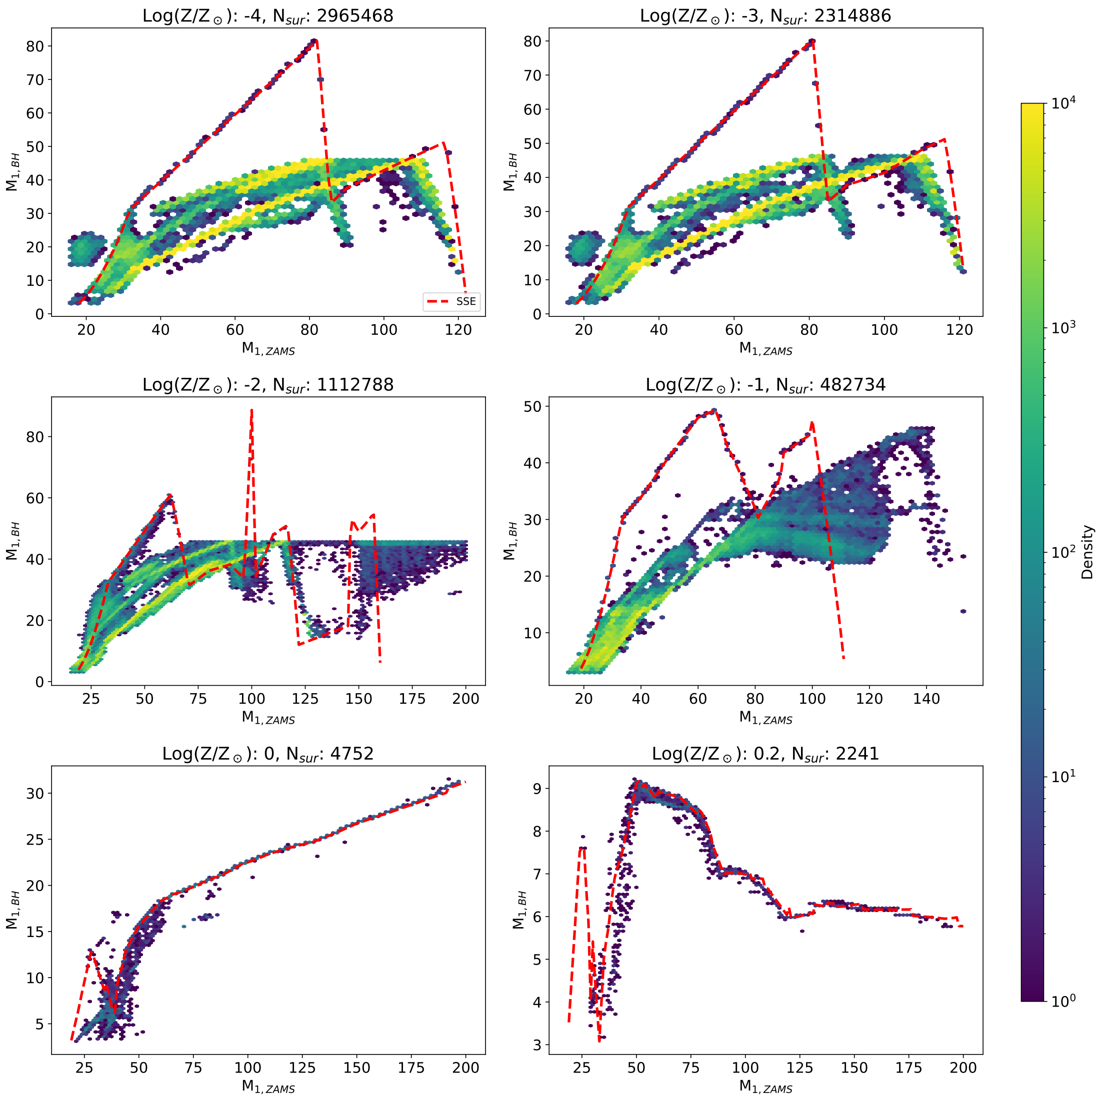

Figure 2.2: Density plots of primary black hole mass ($M_{1,\mathrm{BH}}$) versus the primary ZAMS mass ($M_{1,\mathrm{ZAMS}}$) for six metallicities. Each panel lists the metallicity and the number of surviving systems (N_{surv}), illustrating the steep decline in BBH formation efficiency toward higher metallicities. The color scale shows the logarithmic density of systems in hexagonal bins of $2\,M_\odot$, while the red dashed lines indicate single-star evolution (SSE) predictions. The initial-final mass relation deviates substantially from the SSE case, reflecting the effects of binarity.

in the maximum M_{ZAMS} for black holes at $\log_{10}(Z/Z_\odot) \lesssim -2.5$. This limits the growth and quantity of otherwise fairly massive black holes.

We construct the joint distribution of the primary black hole mass, M_1, and mass ratio, $p(M_1, q \mid M_{\mathrm{ZAMS}}, Z)$, derived from the population of surviving binary black hole systems.

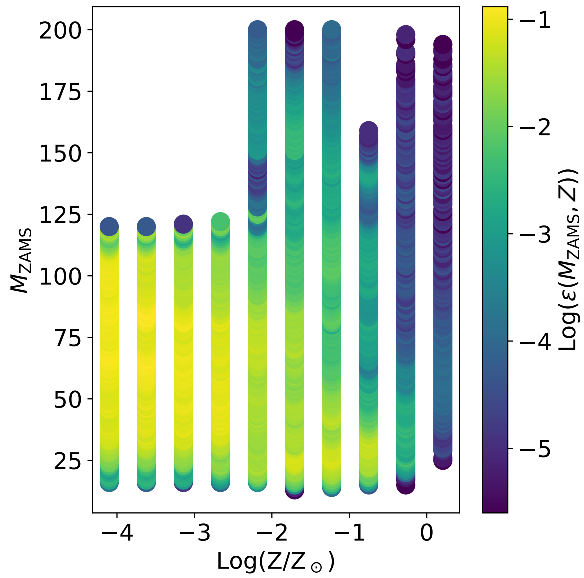

Figure 2.3: The efficiency for black holes to form given a particular progenitor mass, $M_{\rm ZAMS}$, and metallicity, Z. Low-metallicity progenitors experience a higher efficiency and overall survivability probability as expected. In contrast, higher-metallicity systems, although they may avoid pair instability, have a much lower probability of survival as indicated by the reduced efficiency factor.

The black hole mass ratio is defined as $q \equiv M_2/M_1$. For each combination of progenitor mass $M_{\rm ZAMS}$ and metallicity Z, we use a two-dimensional kernel density estimate over the M_1–q space to determine the density. The resulting joint distribution $p(M_1, q \mid M_{\rm ZAMS}, Z)$ is then normalized by our efficiency function from Equation 2.2. Finally, we interpolate the normalized distribution bilinearly over a grid of Z and $M_{\rm ZAMS}$ values, ensuring full representation across the parameter space.

2.3.2 Galactic Binary Black Hole Population

Building on the results from Section 2.3.1, we develop a model for the population of binary black holes that merge within a Hubble time. Our goal is to quantify the growth and abundance of these systems in galaxies of varying mass across cosmic time. The formation rate of such objects is captured by the following integral:

$$\frac{d^2N}{dM_1\,dt}(M_\star, z) = \frac{S(z, M_\star)}{\langle M_p\rangle(M_\star, z)} \int_0^1 dq \int_{Z_l}^{Z_u} dZ\, \rho(Z|M_\star, z)$$
$$\times \int_{M_{ZAMS,\mathrm{low}}}^{M_{ZAMS,\mathrm{high}}} dM_{\mathrm{ZAMS}}\, p(M_{\mathrm{ZAMS}}|z, M_\star)\, p(M_1, q|Z, M_{\mathrm{ZAMS}})\,, \tag{2.3}$$

The limits Z_u and Z_l are taken from the bounds of our simulations, $Z_l = 1.5 \times 10^{-6}$ and $Z_u = 3.2 \times 10^{-2}$. The limits of the mass integral are from $M_{\mathrm{ZAMS,low}} = 8\,\mathrm{M}_\odot$, the mass needed to produce a CCSN, up to $M_{\mathrm{ZAMS,high}} = 200\,\mathrm{M}_\odot$, rather than the Eddington limit, in light of the various stars observed with estimated masses greater than $150\,\mathrm{M}_\odot$ (Crowther et al., 2010; Bestenlehner et al., 2014, 2020). We note that, due to the suppression by the IMF, this change in limits has a negligible impact on the overall result. Although the lower bound does not coincide with the $M_{\mathrm{ZAMS}} = 10\,\mathrm{M}_\odot$ limit of our simulations, no merging BBHs form with primary masses between 8—10 $M_\odot$, so the distribution $p(M_1, q \mid M_{\mathrm{ZAMS}}, Z)$ naturally vanishes in this range. The distribution function of primary black holes, $p(M_1, q \mid M_{\mathrm{ZAMS}}, Z)$, is calculated as given in Section 2.3.1, normalized by $\varepsilon(M_{\mathrm{ZAMS}}, Z)$ defined in Equation 2.2.

The initial mass function (IMF), i.e., the distribution of stars per unit M_{ZAMS} at a given redshift and galaxy mass is given by $p(M_{\mathrm{ZAMS}}|z, M_\star)$. It is defined such when integrated over M_{ZAMS}, we get back unity. Here we assume a power-law of the form M_{ZAMS}^{ξ} for the IMF *over the range of stellar masses that gives rise to black holes* and we will constrain this using the core collapse supernova rate. To account for the possibility of a top-heavy IMF, we allow the exponent ξ to change as a function of stellar mass using the relation $M_p^{-\alpha+\beta\log_{10}(M_{\mathrm{ZAMS}}/8)}$. Here, β sets the stellar mass-dependent flattening or steepening of the IMF at higher masses. We choose β such that it enhances the high-mass tail while remaining consistent with observational constraints, such as those found in Marks et al. (2012) for globular clusters.

In this work, we present results based on two IMFs, the constant slope of -2.3 (i.e., $\beta = 0$) consistent with the Kroupa IMF (Kroupa, 2001) and the running-slope model with $\beta = 0.3$. Note that we choose a slope of -2.3 at $8\mathrm{M}_\odot$ for both models to facilitate comparison between the two models. We have not performed a scan of parameters to determine what would best fit the LIGO data, as this exercise would be premature given the uncertainties that still need to be explored. This comparison also highlights the potential impact of metallicity-driven variations of the IMF, the black hole mass distribution, and the number density. There is evidence suggesting potential variations at the high-mass end due to environmental factors such as metallicity in dwarf galaxies and globular clusters (Marks et al., 2012; Geha et al., 2013; Gennaro et al., 2018; Weatherford et al., 2021), as well as work that has explored variations in the IMF and highlighted its impact (Chruślińska et al., 2020; Tanikawa et al., 2022). In particular, the work by Tanikawa et al. (2022) also examine a top-heavy IMF, $\propto M_{\mathrm{ZAMS}}^{-1}$, however only for very metal poor stars, those outside the metallicities we consider here. In contrast, we gradually flatten ours across the mass space for stars large enough to end their lives in a core-collapse supernova.

The metallicity distribution function, $\rho(Z \mid M_\star, z)$, is assumed to follow a log-normal distribution, though this is shown to be a simplifying assumption (van Son et al., 2023a), within the limits defined by the range of our simulations $-4.1 < \log_{10}(Z/Z_\odot) < 0.2$, and evolves over cosmic time (Zahid et al., 2013). The mean, $\langle Z(M_\star, z)\rangle$, is obtained from Ma et al. (2016), who fit to high-resolution cosmological zoom-in simulations of the evolution of galaxy mass-metallicity relations, in agreement with results up to z=3 including Tremonti et al. (2004); Gallazzi et al. (2005); Erb et al. (2006); Mannucci et al. (2009); Zahid et al. (2011); Kirby et al. (2013); Steidel et al. (2014); Sanders et al. (2015). We use the gas-phase metallicity, with the functional form restated here for convenience:

$$\log\left(\frac{Z_{\rm gas}}{Z_\odot}\right) = 0.35\left[\log\left(\frac{M_\star}{M_\odot}\right) - 10\right] + 0.93\, e^{-0.43z} - 1.05\,. \tag{2.4}$$

For the variance used in the $\rho(Z \mid M_\star, z)$, we use $\sigma(M_\star, z) = 0.4$, constant across redshift, which is the median value found in the compilation of low redshift stellar metallicity observations found in Simon (2019) and the values found in Gallazzi et al. (2005) which have a mean redshift value of $z = 0.13$. If instead we chose to use a value akin to that used in Santoliquido et al. (2020), i.e. $\sigma(M_\star, z) = 0.2$, the total binary black hole merger rate varies by less than a factor of 2 overall.

The function outside the integrand in Equation 2.3, $S(z, M_\star)$, represents the total stellar mass formed per year in a galaxy of stellar mass $M_\star$ at redshift z, often referred to as the *star-forming main sequence*. We adopt the empirical relation from Speagle et al. (2014), constructed from a meta-analysis of 25 previous studies. In their formulation, the main sequence is expressed as a function of the cosmic time t (in Gyr), which we convert to redshift using $t = t(z)$. This choice, when convolved with our GSMF, yields a star-formation rate density which we compare to observational works in Appendix 2A. For convenience, the functional form is given here:

$$\mathrm{SFR}(M_\star, t(z)) = 10^{-(6.51 - 0.11\, t(z))}\left(\frac{M_\star}{M_\odot}\right)^{0.84 - 0.026\, t(z)} M_\odot\, \mathrm{yr}^{-1}. \tag{2.5}$$

In Equation 2.3, we note that the average progenitor mass is defined with respect to the IMF defined over the entire star-forming range, i.e., $\langle M_{\rm ZAMS}\rangle(M_\star, z) = \int dM_{\rm ZAMS}\, M_{\rm ZAMS}\, p(M_{\rm ZAMS}|M_\star, z)$.

If we write $p(M_{\rm ZAMS}|z, M_\star) = C M_{\rm ZAMS}^{\xi}$ then we find that

$$\frac{C}{\langle M_{\rm ZAMS}\rangle(M_\star, z)} = \frac{f_{bin} K^{CC}}{2\int_{m_l^{CC}}^{m_u^{CC}} dM_{\rm ZAMS}\, M_{\rm ZAMS}^{\xi}}\,, \tag{2.6}$$

where K^{CC} is the number of supernova formed per unit stellar mass as given in (Botticella et al., 2008) and we assume that this is constant. The limits in the integral are from $m_l^{CC} = 8\,\mathrm{M}_\odot$ to $m_u^{CC} = 200\,\mathrm{M}_\odot$. The volumetric supernova rate is then given by $R^{CC}(z) = K^{CC}\psi(z)$. We use our SFRD today, $\psi(z=0)$, and $R^{CC}(z \approx 0)$ from the ZTF survey (Perley et al., 2020) to obtain the value of $K^{CC} = f_\psi R^{CC}_{ZTF}(z \approx 0)/\psi(0)$. We introduce f_ψ as a correction factor to account for the model dependency of the SFR, we find a value of $f_\psi = 0.4$ to best match our model's rate to LVKs. The overall amplitude is sensitive to shifts in the prescriptions we use, such as the SFR or binary fraction and can vary by about an order of magnitude (Neijssel et al., 2019a). While a more rigorous treatment would begin with the specific star formation rate, measuring it to high redshift is challenging. Pursuing that path, though compelling, lies beyond our current scope and is left for future work.

The factor of 1/2 in Equation 2.6 is included to avoid double counting our systems, as we are only considering the primary mass of the binary. The overall binary fraction of stars in the galaxy is denoted by f_{bin}. The study by Sana et al. (2008) finds the intrinsic binary fraction to be $f_{\rm bin} \sim 0.7$, while Moe & Di Stefano (2017) report that O-type stars are more commonly found in triples ($n = 3$) and quadruples ($n = 4$), with $f_{n\geq 3} = 0.73$ and $f_{\rm bin} = 0.21$. Triples and quadruples contributing to the merger rate would consist of a closer-orbit binary pair with additional stellar companions. While we do not explicitly factor in $n \geq 3$ systems, our model accounts for a subset of their evolution. Thus, $f_{\rm bin} \in [0.21, 0.94]$. For consistency with both studies, we adopt $f_{\rm bin} = 0.7$ throughout.

2.3.3 Volumetric Binary Black Hole Formation Rate

In the last section we derived the per-galaxy instantaneous production rate of merging binary black holes with primary mass M_1 at a given redshift, $d^2N(M_\star, z)/dM_1dt$ (Equation 2.3). In this section we will extend this analysis to a volumetric population by convolving the galaxy stellar mass function (GSMF), $\phi(M_\star, z)$, which quantifies the number of galaxies per unit volume per dex in galaxy stellar mass with our binary black hole formation rate, defined as:

$$\dot{n}_{\rm BBH}(M\star, M_1, z) = \phi(M_\star, z)\, \frac{d^2N_{\rm BBH}}{dM_1\, dt}(M_\star, z)\,, \tag{2.7}$$

The quantity $\dot{n}_{\rm BBH}(M\star, M_1, z)$ represents the contribution to the volumetric BBH formation rate density per primary mass interval and per dex in galaxy stellar mass. To account for the distribution of galaxies across stellar mass and redshift, we adopt an observationally constrained star-forming GSMF. For $z < 0.2$, we use the double-Schechter fit from Baldry et al. (2012). For $0.2 \leq z \leq 5.5$, we adopt the GSMF from Weaver et al. (2023), who fit a double-Schechter function for $z \lesssim 2.5$ and a single-Schechter function at higher redshift. To ensure smooth evolution between bins, we linearly interpolate the Schechter parameters between the bin centers. We extrapolate the redshift evolution of the Schechter parameters beyond $z = 5.5$ using a power-law dependence on $(1 + z)$, calibrated to the observed trends at $z \geq 2$:

$$M^*(z) \approx 1.42 \times 10^{12}\, (1 + z)^{-2.376}, \tag{2.8}$$

$$\phi_1(z) \approx 7.52 \times 10^{-4}\, (1 + z)^{-0.882}. \tag{2.9}$$

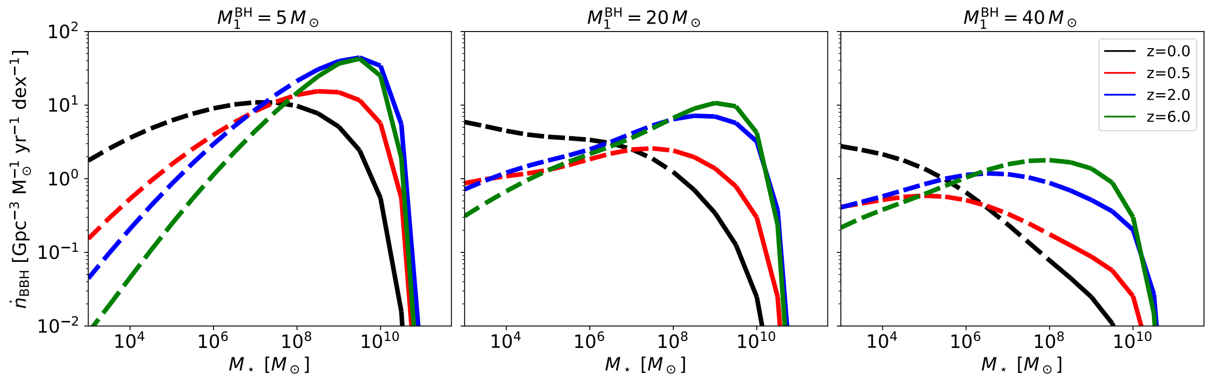

Figure 2.4: Differential binary black hole (BBH) formation rate density, $\dot{n}_{\rm BBH}(M_\star, M_1, z)$, as a function of galaxy stellar mass for three representative primary black hole masses (left to right: $M_1 = 5$, 20, and $40\,{\rm M}_\odot$). Each panel shows four redshifts ($z = 0$, 0.5, 2, and 6), illustrating how BBH formation shifts across the galaxy population over cosmic time. Binary black hole formation migrates to lower galactic stellar masses as the Universe evolves, this migration occurs quicker for more massive BBHs. Dashed segments indicate the regime below $M_\star = 10^8\,{\rm M}_\odot$, where the galaxy stellar mass function is extrapolated to lower stellar masses where observational data is not yet constraining.

In this extrapolation, we keep α_1 fixed at its constant value, consistent with the single-Schechter regime in Weaver et al. (2023). These high-redshift systems contribute minimally to the low-redshift merger rate density due to the short lifetimes of progenitor stars and the preference for shorter delay-times in our systems.

The results of Equation 2.7 are shown in Figure 2.4. Each panel corresponds to a different primary black hole mass, illustrating how the differential BBH formation rate density, $\dot{n}_{\rm BBH}(M_\star, M_1, z)$, depends on the stellar mass of the host galaxy and on cosmic time. At early epochs near the peak of cosmic star formation ($z \simeq 2$), the formation of both $5\,{\rm M}_\odot$ and $20\,{\rm M}_\odot$ black holes occurs primarily in relatively massive galaxies, whereas the highest-mass ($40\,{\rm M}_\odot$) black holes are already associated with lower-mass systems. By $z \simeq 0.5$ (~ 4.5 Gyr ago), lower-mass galaxies begin dominating the production of binary black holes, with the most massive black holes forming preferentially in smaller hosts. At the present epoch, the formation of $40\,{\rm M}_\odot$ black holes peaks in the lowest mass galaxies, while even $5\,{\rm M}_\odot$ systems predominantly form in smaller dwarf galaxies. The dashed portions of each curve mark the region below $M_\star = 10^8\,\mathrm{M}_\odot$, where the galaxy stellar mass function is extrapolated beyond the range constrained by current observations; results in this regime should therefore be interpreted with caution. While Chruslinska & Nelemans (2019) show that extrapolating the GSMF to low stellar masses can lead to an overestimation of the star formation rate at high redshift, our adopted GSMF does not exhibit such an excess at early times, although it tends to overestimate the SFR at our current epoch. Overall, the figure highlights a clear shift of black hole formation toward progressively lower-mass galaxies as the Universe evolves.

2.3.4 The Merger Rate Density

To obtain the merger rate density, we require the delay-time distribution (DTD), or the time between binary formation, taken to be the time the stars enter the main sequence, and black hole merger. The delay time, τ, depends on orbital parameters, stellar evolution prescriptions, and the mass-loss history. We calculate the DTD, $P(\tau)$, directly from our simulations, as shown in Figure 2.5. Due to the high mass of the stars that we consider, their average lifetimes are negligible in comparison to the merger time. We fit a power-law function of the form $P(\tau) \propto \tau^{-0.85}$ to the DTD, similar to that found in Neijssel et al. (2019a); Mukherjee & Moradinezhad Dizgah (2022); Fishbach & Kalogera (2021); Karathanasis et al. (2023). We will show in the following section how even modest changes to the DTD spectral index shapes the merger rate. Convolving our Equation 2.3 with the GSMF and the DTD we obtain the differential merger rate density, at a given time t_0:

$$\frac{d\mathcal{R}}{dM_1} = \int_0^{t_0} \int_{M_{min}}^{M_{max}} \phi(M_\star, z(t_0-\tau')) \frac{d^2 N_{BBH}}{dM_1\, dt}(M_\star, t_0-\tau') P(\tau') d\tau' dM_\star\,. \tag{2.10}$$

The limits of integration over galaxy stellar mass are set to $M_{\mathrm{min}} = 10^8\,\mathrm{M}_\odot$ and $M_{\mathrm{max}} = 10^{12}\,\mathrm{M}_\odot$, consistent with the observational range of the galaxy stellar mass function from

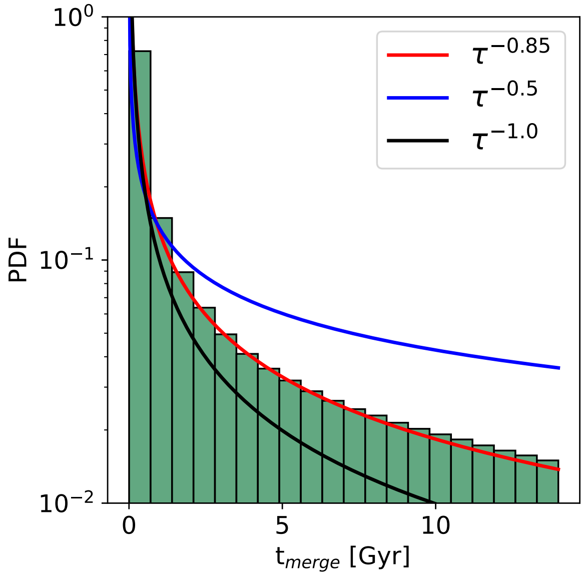

Figure 2.5: Power-law fit to the merger time distribution from the `SEVN` simulations. The distribution is steeply peaked at short delays, implying that the merger rate density today will be dominated by binaries formed more recently in cosmic history. We find the best fit to our data is a power-law $p(\tau) \propto \tau^{-\gamma}$ with spectral index $\gamma = 0.85$. We include a steeper distribution of $\gamma = 1.0$ (Fishbach & Kalogera, 2021) and a flatter more extreme value of $\gamma = 0.5$ for comparison.

Weaver et al. (2023). Although galaxies with smaller stellar masses exist, they are not well constrained observationally, and extrapolating the GSMF below this limit can lead to an overestimation of the star formation rate at low redshift. We find that extending the lower limit down to $10^3\,\mathrm{M}_\odot$ affects the inferred merger rate density in the next section by less than a factor of two.

2.4 Results

2.4.1 Primary Mass Spectrum

In the previous section, we detailed the calculation of the merger rate density given in Equation 2.10. Throughout this section, we compare our results to inferred distributions based on the gravitational wave observations from the LVK collaboration and find significant agreement with their findings.

In the most recent gravitational-wave transient catalog (Abac et al., 2025), the LVK Collaboration models the merger rate density and its derivatives using both parametric and non-parametric approaches. In Figure 2.6, we compare our results to two of their reference models: the *Broken Power Law + Two Peaks* (BPL+2P) model, which describes a broken power-law continuum between a minimum and maximum mass with two additional Gaussian peaks at $\sim 10\,\mathrm{M}_\odot$ and $\sim 35\,\mathrm{M}_\odot$, and the non-parametric *B-Spline* model, which represents the distribution using one-dimensional basis splines. To compare with these, we highlight two different forms for the power-law index, $\xi = -\alpha + \beta \log_{10}(M_{\mathrm{ZAMS}}/8\mathrm{M}_\odot)$, our constant-slope ($\alpha = 2.3, \beta = 0.0$) which matches the fiducial Kroupa IMF (Kroupa, 2001) and our running-slope ($\alpha = 2.3, \beta = 0.3$) where we introduce a non-zero β to gradually flatten the IMF at higher masses. We find that the running-slope model exhibits greater agreement with LVK observations, suggesting that mass-dependent variations in the IMF, required mainly in dwarf galaxies, play a significant role in shaping the BBH mass distribution.

The inferred peak at $\sim 35\,\mathrm{M}_\odot$ is apparent across our models, associated with the pair-instability mechanism. The smoothing toward lower masses, present in the LVK *BPL+2P* and *B-Spline* models, is attributed to factors such as metallicity and binary evolution, which naturally blur the edge of the proposed lower mass gap (Abbott et al., 2019b). Our model exhibits a similar tapering at low masses but with the primary peak shifted to slightly lower values ($\sim 6\,\mathrm{M}_\odot$ compared to $\sim 10\,\mathrm{M}_\odot$ in the LVK results). This offset could potentially be reduced by varying the binary evolution parameters in `SEVN`; however, doing so would require running new simulations and is left to future work.

Our model demonstrates good agreement with the observed differential merger rate across much of the primary mass spectrum, including the tapering at low masses and the peaks. However, discrepancies arise at the high-mass end of the spectrum, where the model underestimates the number of binary black holes required to populate the observed tail of the

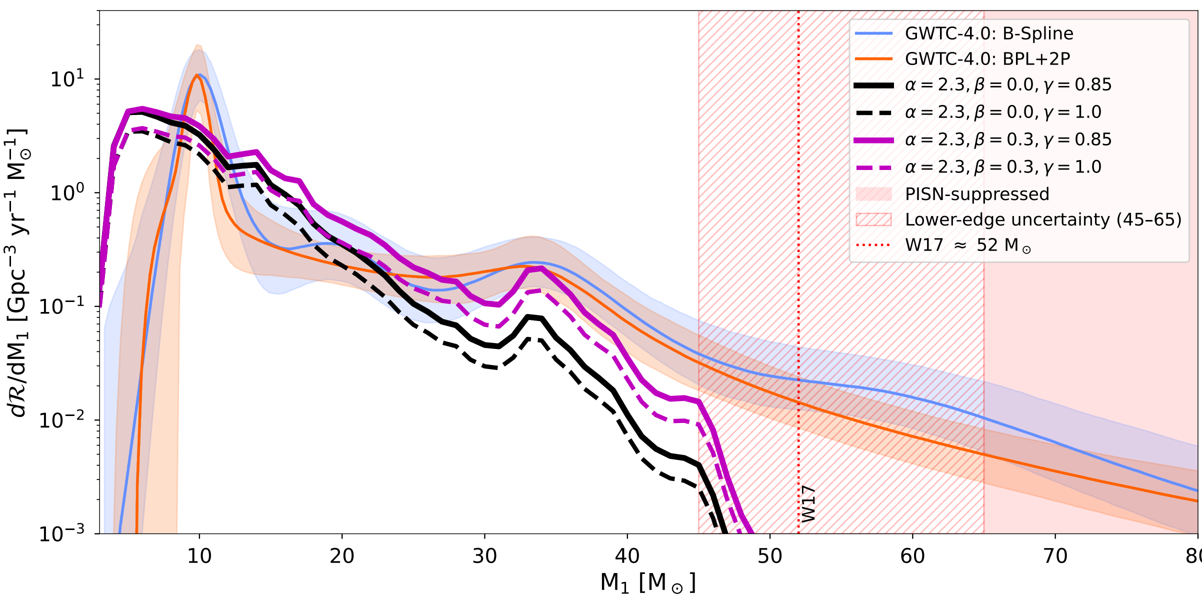

Figure 2.6: Comparison between the LIGO-Virgo-KAGRA *Broken Power Law + Two Peaks (BPL+2P)* and *B-Spline* models, shown in orange and blue, respectively, and our two IMF models: a constant-slope IMF with ($\alpha = 2.3$, $\beta = 0.0$) and a running-slope IMF with ($\alpha = 2.3$, $\beta = 0.3$). We also compare the impact of varying the delay-time distribution, $P(\tau) \propto \tau^{-\gamma}$, for $\gamma = 0.85$ and 1.0. The PISN threshold is given by the red band, reflecting uncertainty in the lower boundary ($\sim$ 45–65 $\mathrm{M}_\odot$). The SEVN M20 PISN implementation follows the prescription of Woosley (2017) (W17), whose threshold is also included. Our model produces very few BBH mergers within the pair-instability mass gap, consistent with standard stellar evolution theory, leading to a natural truncation around $\sim 50\,\mathrm{M}_\odot$.

merger-rate distribution. In our case, the distribution naturally tapers off around $\sim 50\,\mathrm{M}_\odot$, due to the M20 PISN model in SEVN Spera & Mapelli (2017), based on Woosley (2017). The true onset of the pair-instability process, however, may occur at slightly higher masses, up to $\sim 65\,\mathrm{M}_\odot$ (Woosley & Heger, 2021). This uncertainty could contribute to the observed discrepancy at the high-mass end. Black holes above this range may arise from alternative formation channels, such as the dynamical channel, population III stars (Farrell et al., 2021; Kinugawa et al., 2021; Liu & Bromm, 2020), hierarchical mergers (Mapelli, 2016; Farr et al., 2017; Mapelli et al., 2021; Gerosa & Fishbach, 2021; Sedda et al., 2023; Torniamenti et al., 2024; Bruel et al., 2025), primordial black holes (De Luca et al., 2021), or beyond the Standard Model (BSM) physics (Croon et al., 2020; Sakstein et al., 2020; Ziegler & Freese, 2021, 2024; Croon & Sakstein, 2024).

Delay-Time Distribution Effects

Included in the models in Figure 2.6 is a dependency on the delay-time distribution (DTD) power-law index. If we assume that the DTD takes the form $P(\tau) \propto \tau^{-\gamma}$ our simulations show that $\gamma = 0.85$ as shown in Figure 2.5. However, this can be influenced by many factors such as the true underlying distribution of parameters, e.g. eccentricity and orbit, as well as stellar evolution effects.

Here we briefly discuss one such possibility, that the variation in the common envelope (CE) parameters, α_{CE} and λ_{CE}, can lead to changes in the merger time and thus the DTD. We only examine one system here, and leave a full study to future work. In particular, we examine the case of a binary system with progenitor masses $M_1 = 51\,\mathrm{M}_\odot$ and $M_2 = 50\,\mathrm{M}_\odot$. In all cases where the binary survived to form black holes and merge through gravitational wave (GW) emission, the final remnant masses were fixed at $M_{1,\mathrm{BH}} = 22.7\,\mathrm{M}_\odot$ and $M_{2,\mathrm{BH}} = 22.0\,\mathrm{M}_\odot$. Depending on the CE parameters, systems either merged prematurely during the CE phase ($\mathcal{O}(1\,\mathrm{Myr})$) or survived as binary black holes with various delay times.

The CE efficiency parameter $\alpha_{\rm CE}$ controls the fraction of orbital energy used to unbind the envelope. A value greater than unity simply means that the orbital energy is not the only source removing the envelope (Ivanova et al., 2013; Nandez & Ivanova, 2016; Roepke & De Marco, 2023). Hydrodynamical studies have shown that additional energy sources, such as recombination energy, can contribute significantly to envelope ejection Nandez & Ivanova (2016); Sand et al. (2020); Roepke & De Marco (2023).

The envelope structure parameter $\lambda_{\rm CE}$ characterizes the binding energy of the envelope and depends on the stellar mass, evolutionary phase, radius, and envelope structure (van den Heuvel, 1976; Webbink, 1988). Different prescriptions for $\lambda_{\rm CE}$ have been proposed in the literature Xu & Li (2010); Claeys et al. (2014); Klencki et al. (2021). We refer the reader to Iorio et al. (2023a) for these implementations and further details on the common envelope prescription in `SEVN`.

The delay times and outcomes for various values of $\alpha_{\rm CE}$ and $\lambda_{\rm CE}$, varied within reason according to the `SEVN` prescriptions, are summarized in Table 2.1. We show that altering the α_{CE} parameter alone, by a factor of ~ 2, can shift the merger time by over an order of magnitude. This is in agreement with the findings in Claeys et al. (2014), who find that the slope of their DTD shifts when varying the common envelope efficiency. They find that for higher α_{CE} the DTD flattens and longer merger times are favored, while for smaller α_{CE} the DTD steepens, favoring shorter merger times. Decreasing α_{CE} leads to a lower binary black hole formation efficiency, potentially lowering the overall merger rate. While we do not quantify the full impact of such effects on the DTD in this work, we illustrate how sensitive the merger timescale is to even modest changes in stellar evolution assumptions, spanning outcomes from prompt stellar mergers to binaries with merger times exceeding the Hubble time by orders of magnitude. A comprehensive exploration of these effects is left for future study.

Parameter	Delay Time [Myr]	Outcome
$\alpha_{\rm CE} = 5$	234.4	BBH merger (GW)
$\alpha_{\rm CE} = 4$	111.9	BBH merger (GW)
$\alpha_{\rm CE} = 3$	52.1	BBH merger (GW)
$\alpha_{\rm CE} = 2$	15.7	BBH merger (GW)
$\alpha_{\rm CE} = 1$	–	Stellar merger (CE)
$\alpha_{\rm CE} = 0.1$	–	Stellar merger (CE)
$\lambda_{\rm CE} = -1$ (Claeys et al., 2014)	52.1	BBH merger (GW)
$\lambda_{\rm CE} = -4$ (Klencki et al., 2021)	9.4	BBH merger (GW)
$\lambda_{\rm CE} = -5$ (Xu & Li, 2010)	25.0	BBH merger (GW)
$\lambda_{\rm CE} = -12$ (Claeys et al., 2014)	$> 10^6$	No merger (BBH)
$\alpha_{\rm CE} = 1, \lambda_{\rm CE} = -12$	444.6	BBH merger (GW)
$\alpha_{\rm CE} = 5, \lambda_{\rm CE} = -5$	104.4	BBH merger (GW)

Table 2.1: Delay times and outcomes for a binary with $M_{p,1} = 51\,{\rm M}_\odot$ and $M_{p,2} = 50\,{\rm M}_\odot$, evolved under different common envelope (CE) parameter choices in `SEVN`. All systems that survive CE and undergo GW-driven inspiral result in BBH mergers with $M_{1,\rm BH} = 22.7\,{\rm M}_\odot$ and $M_{2,\rm BH} = 22.0\,{\rm M}_\odot$. The α-formalism for CE evolution in `SEVN` is adopted from Hurley et al. (2000). For the CE binding energy parameter $\lambda_{\rm CE}$, we follow `SEVN`'s convention of using integer flags to select different parameterizations and denote which study they are from in the Table. The default parameters in `SEVN` are $\alpha_{\rm CE} = 3$ and $\lambda_{\rm CE} = -1$. In each row, we indicate which parameter is varied while the other is held fixed at its default value, except in the last two rows, where both parameters are modified simultaneously.

Dynamical Single Black Holes

An intriguing avenue to explore the population of mergers within the upper mass gap lies in the role of single black holes. Although the fraction of massive stars born as singles is relatively low (around 10%), these single black holes evolve in isolation from ZAMS to BH and are not subject to the same efficiency factor constraints as shown in Figure 2.3. As a result, if they retain their hydrogen envelopes, they can form more massive black holes, even extending into the mass gap.

Once formed, single black holes can enter binaries through dynamical capture (Mapelli, 2016; Gerosa & Fishbach, 2021), a process that depends on several factors, including the cluster density, stellar interaction rates, and black hole retention within these environments. After formation, black holes tend to migrate toward the cluster core due to dynamical relaxation (Portegies Zwart & McMillan, 2000). Closer to the cluster core, the more frequent dynamical encounters can result in the formation of binaries or higher-order systems. While many black holes will experience multiple interactions or even ejection from the cluster, some may successfully become part of a binary system capable of merging within a Hubble time.

We sketch a simplified picture to estimate the efficiency of this process, η, which we define as the probability for a single black hole to capture a companion (i.e., form a binary) and merge within the age of the Universe. We can break up this efficiency into various factors to get a better idea; $\eta \sim f_{\rm sse} \times f_{\tau} \times f_{\rm dyn} \times f_{\rm GC} \times f_{\rm primary}$. Here, $f_{\rm sse}$ represents the fraction of massive stars born as singles, f_{τ} is the probability for dynamically formed binaries to merge within a Hubble time, $f_{\rm dyn}$ accounts for the fraction of dynamically interacting black holes that form binaries instead of being ejected that is likely dependent on metallicity (Kumamoto et al., 2020), $f_{\rm GC}$ quantifies the fraction of black holes residing in dense stellar environments like globular clusters (which we expect to be dependent on the stellar mass of the galaxy and perhaps its metallicity), and $f_{\rm primary}$ reflects the likelihood that the single black hole becomes the primary in the binary. These various factors depend on the mass of the single black hole, which would require detailed simulations to explore.

To explore the possibility of binaries formed dynamically contributing to the upper mass gap, we decompose the total differential merger rate into two components:

$$\frac{dR}{dM_1} = \frac{dR_{\rm BBH}}{dM_1} + \bar{\eta}(M_1)\frac{dR_{\rm SBH}}{dM_1}\,, \tag{2.11}$$

where $\frac{dR_{\rm BBH}}{dM_1}$ corresponds to the contribution from black holes formed in isolated binaries, while $\frac{dR_{\rm SBH}}{dM_1}$ accounts for the contribution from single black holes and $\bar{\eta}$ is the efficiency averaged over galaxy masses and metallicities. Given our discussion above, we should build in the efficiency calculation before averaging over the properties of galaxies but we are only interested in an order of magnitude answer here. If the black hole mass dependence of $\bar{\eta}$ is neglected, we can do a quick estimate of the efficiency needed to populate black holes in the mass gap. We estimate that, for an efficiency $\bar{\eta} \sim 10^{-4}$, the contribution from dynamically formed binaries can extend the black hole mass spectrum up to about $60\,\mathrm{M}_\odot$, about $\sim 10\,\mathrm{M}_\odot$ further into the mass gap than our model.

It is also worth keeping in mind that the lower mass when pair instability kicks in is not well known. The effects discussed in Woosley & Heger (2021) suggest that the pair instability could shift to higher masses ($\sim$ 65–70 $\mathrm{M}_\odot$) and this would offer an additional way to populate the mass gap. It would be interesting to include these effects along with the mass loss in binary systems to explore this further. Together, these considerations present pathways to populate the upper-mass gap that deserve further study.

2.4.2 Secondary Mass Spectrum

With over 150 confirmed BBH events detected up to O4a (Abac et al., 2025), the LVK dataset offers growing potential to probe the component mass spectra of merging binary black holes. While the primary mass distribution has received significant attention, the

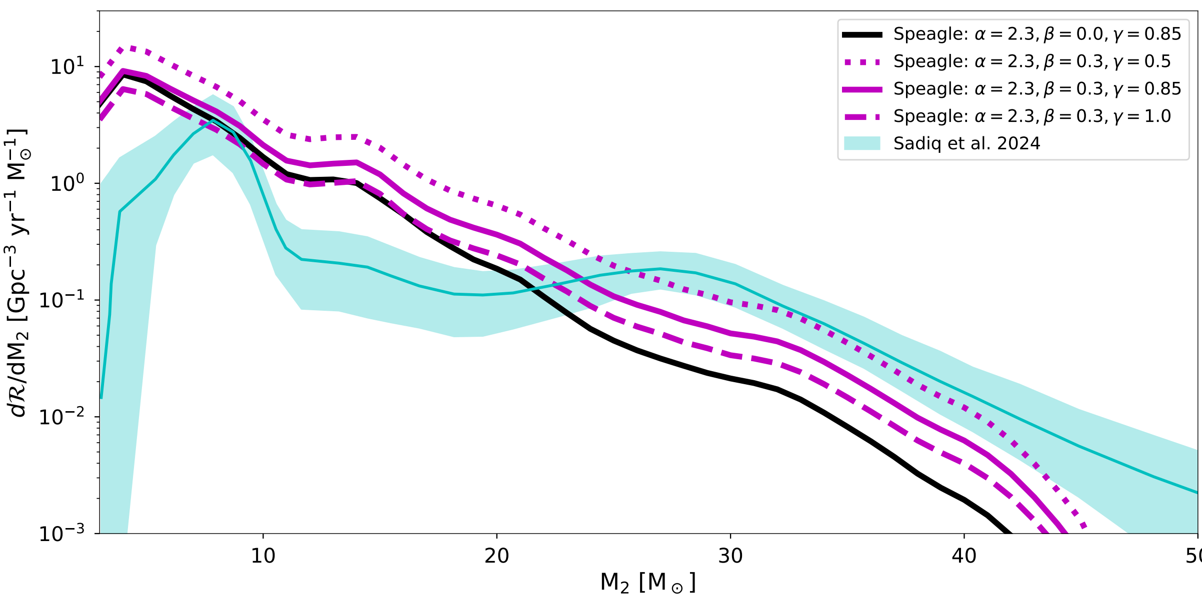

Figure 2.7: Differential merger rate distribution as a function of secondary black hole mass, dR/dM_2. We show both constant-slope ($\alpha = 2.3$, $\beta = 0.0$) and running-slope ($\alpha = 2.3$, $\beta = 0.3$) IMF models, where the running-slope model is evaluated under varying delay-time power-law indices. For comparison, we include the inferred distribution from Sadiq et al. (2024), shown as a cyan band. While our results are broadly consistent, our predicted shape follows a closer approximation to a single power-law with small bump-like features, rather than the more pronounced structure in their inferred distribution. This is particularly useful to compare against future LVK results.

secondary mass spectrum remains comparatively underexplored, though recent work has begun to address this gap (Sadiq et al., 2024; Farah et al., 2024). While the LVK *FullPop-4.0* model provides posterior distributions for both component masses, the published secondary-mass spectrum is limited to $m_2 \lesssim 15\,\mathrm{M}_\odot$ and represents the full compact-object population, including both neutron stars and black holes. As such, it does not isolate the secondary-mass distribution for binary black holes alone. Since a BBH-specific secondary spectrum has not yet been released, we compare our results to the inference-based study of Sadiq et al. (2024). Unlike the inference-based approaches of Sadiq et al. (2024); Farah et al. (2024), our framework models the astrophysical population by combining galaxy and cosmological evolution with stellar population synthesis, allowing for a self-consistent exploration of the origin and properties of the secondary mass distribution.

Starting from Equation 2.3, and omitting the integration over mass ratio q, we obtain the differential distribution $\frac{d^2N}{dM_1\,dq\,dt}(M_\star, z)$. To express this in terms of the secondary mass $M_2 = qM_1$, we perform a change of variables in phase space, applying the appropriate Jacobian transformation:

$$\frac{d^2N}{dM_2\,dt}(M_\star, z) = \int_{M_1=M_2}^{80\,\mathrm{M}_\odot} \left|\frac{dq}{dM_2}\right| \frac{d^3N}{dM_1\,dq\,dt}(M_\star, z)\,dM_1\,. \tag{2.12}$$

This transformation requires evaluating $\frac{d^2N}{dM_1\,dq\,dt}$ outside our simulated grid for values of (M_1, q), since M_2/M_1 will generally not lie exactly on the original simulation grid. To handle this, we use a bivariate interpolation across the simulated grid of M_1 and q to estimate the distribution at arbitrary values. We then follow the steps in Section 2.3.4 to obtain the secondary mass spectrum $d\mathcal{R}/dM_2$ at $z = 0.2$ which we show in Figure 2.7.

Compared to the primary mass distribution, the secondary mass spectrum features a larger number of low-mass events, $\lesssim 8\,\mathrm{M}_{\odot}$, but the rate declines more steeply at higher masses. Except for this overabundance at low masses the secondary spectrum remains below the primary mass spectrum in terms of overall events. The peak which is present in the primary spectrum at $\sim 35\,\mathrm{M}_{\odot}$ is flattened to a small bump like feature and shifted to slightly lower masses in the secondary spectrum. Within the peak mass range, $\sim 30-40\,\mathrm{M}_{\odot}$, the spectra are dominated by low metallicity systems. Contributions from metallicities in the range $10^{-3} \leq Z < 10^{-4}$ exhibit a mass ratio near unity, thus both the primary and secondary spectra have contributions here. However, the primary mass spectrum is unique in that it obtains a boost of systems in the intermediate metallicity regime between $10^{-2} < Z < 10^{-3}$, yet the mass ratio strays from unity and these systems in the case of the secondary mass spectrum get spread out to lower masses weakening the peak like feature.

Similar to the primary mass spectrum in Figure 2.6 we explore the effect of varying the delay-time distribution (DTD), $P(\tau) \propto \tau^{-\gamma}$, using our fit value of $\gamma = 0.85$, a steeper distribution with $\gamma = 1.0$, and a flatter value of $\gamma = 0.5$ leading to longer mergers. The effect of changing the DTD is effectively an amplitude shift. The impact or DTD variations on the spectral shape is negligible. On the other hand, the variation in the IMF can lead to order of magnitude differences at the high-end tail of the mass distribution when considering our top-heavy model (magenta lines).

For comparison, we include the results from Sadiq et al. (2024), who infer the secondary mass distribution directly from LVK data. Their results exhibit more pronounced features, including a stronger bump near $\sim 30\,M_{\odot}$. In fact Farah et al. (2024) show that the high mass peak can be stronger in the case of the secondary mass spectrum. While our model does not reproduce this structure, the overall slope of our predicted distribution aligns more closely with a single power-law, featuring only mild bump-like deviations. If such a feature is confirmed with higher significance in O4, it may indicate the contribution of a distinct formation channel, such as dynamical assembly in dense stellar environments, or possibly

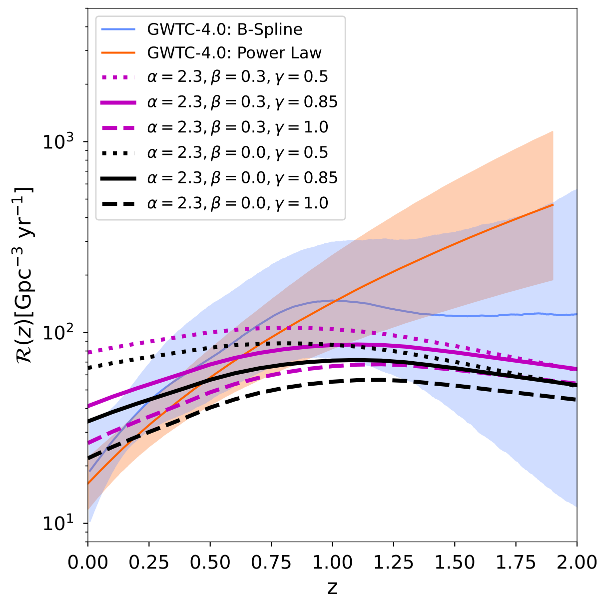

Figure 2.8: Total binary black hole merger rate density as a function of redshift for our constant- and running-slope IMF models ($\alpha = 2.3$, $\beta = 0.0$ and 0.7, respectively), evaluated under three different delay-time distribution power-law indices: $\gamma = 0.5$, 0.85, and 1.0. The shaded bands represents the 90% credible interval inferred from LVK observations. Models with longer delay-times ($\gamma = 0.85$ and 1.0) show better agreement with the inferred redshift evolution, while the $\gamma = 0.5$ model over-predicts the low-redshift rate.

hint at new physics beyond standard stellar evolution models. As with the primary mass spectrum, we observe a sharp decline in the secondary mass spectrum above $\sim 50\,\mathrm{M}_\odot$, consistent with the expectations from our prescription of pair-instability supernovae.

2.4.3 Redshift Dependence

To obtain the total merger rate at a given redshift, we first determine $d\mathcal{R}/dM_1$ at the corresponding cosmic time, $t_0 = t(z)$, using Equation 2.10. We then integrate over the black hole primary mass to find $\mathcal{R}(z)$:

$$\mathcal{R}(z) = \int \frac{d\mathcal{R}}{dM_1}(z) dM_1 \,. \tag{2.13}$$

The results of the redshift evolution of the total merger rate density are given in Figure 2.8 for both our constant- and running-slope IMF models, each evaluated at three different delay-time distribution (DTD) power-law indices: $\gamma = 0.5$, 0.85, and 1.0. These values bracket the distribution favored by our simulations ($\gamma = 0.85$) and explore both steeper and shallower scenarios, corresponding to shorter and longer delay-times. We choose $\gamma = 1.0$ as this is the value that is inferred from the LVK data in Fishbach & Kalogera (2021). The value of $\gamma = 0.5$ is chosen to highlight how the flattening of the DTD affects the redshift evolution of the merger rate, while extreme is not completely ruled out (Mukherjee & Moradinezhad Dizgah, 2022; Karathanasis et al., 2023).

The $\gamma = 0.5$ case produces a flatter evolution, over predicting the rate at low z and is overall a poor fit to the inferred spectrum. Models with shorter delay-times ($\gamma = 0.85$ and 1.0) show better agreement with the LVK-inferred trend remaining mostly in the 90% credible interval for their *B-Spline* model, whle also aligning with their peak location. The peak at lower redshift than the power-law model can be attributed in part due to our adopted star-formation rate density, which peaks at z ≈ 1.5 (Appendix 2A), and is further affected by the convolution with the delay-time distribution and the effects of stellar-evolution prescriptions. For instance, Boesky et al. (2024a,b) show that changes to common envelope efficiency, stable mass transfer efficiency, and supernova natal kick dispersions can shift the peak redshift of the total merger rate. There are also uncertainties surrounding the galactic and cosmological inputs which can impact the merger rate (Chruslinska et al., 2019; Neijssel et al., 2019a; Chruślińska et al., 2021; van Son et al., 2023a).

The differences between the constant- and running-slope models that we present are minimal. The similarity between the two models, for each γ, is because the total merger rate is primarily dominated by low-mass systems. The running-slope IMF is designed to enhance the high-mass end of the black hole mass spectrum, but has limited impact on the total rate due to the suppression of the high-mass tail. Lastly, we only present results on the isolated channel, it is feasible that the dynamical channel may supplement our total merger rate density, at higher redshifts, improving agreement with LVK.

We further explore redshift evolution in Figure 2.9, which shows how the primary mass spectrum, $d\mathcal{R}/dM_1$, evolves over cosmic time for both our constant-slope (left panel) and running-slope (right panel) models. The location of the peaks in $d\mathcal{R}/dM_1$ remains robust across redshift, which is to be expected as they arise from the same underlying physical processes. The primary change we find is in the slope between the low-mass and high-mass peaks: as redshift increases (moving to earlier times), the slope becomes flatter. This evolution is more pronounced in our constant-slope model (fiducial Kroupa IMF) than in the running-slope model ($\beta = 0.3$) which is constructed to increase the relative amount of low-metallicity high-mass black holes, leading to a growing contribution from high-metallicity low-mass black holes relative to low-metallicity high-mass black holes.

These trends, the static location of the peaks and slowly varying slope, are consistent with the observations of LVK data (Fishbach et al., 2021; Abbott et al., 2023; Lalleman et al., 2025). The study by Lalleman et al. (2025) show that location of the $\sim 35\,\mathrm{M}_\odot$ peak is constant across redshift observable by LVK and that the slope of the primary mass spectrum only slightly varies over this time, consistent with our findings in Figure 2.9. On the other hand, Rinaldi et al. (2024) report the overall peak and spectral shape of their primary mass distribution shifts substantially with redshift attributing this to contributions from two distinct BBH populations. Since our analysis is restricted to the isolated channel, other formation pathways could, in principle, shift the primary mass spectrum at earlier redshifts. In contrast, the evolution in our models is insufficient to produce the pronounced spectral changes reported in their work, the distribution would remain dominated by the low-mass peak at $\sim 8\,\mathrm{M}_\odot$ across redshift. Similarly, it was shown that if the PISN mass scale was allowed to vary with metallicity, the peak at $\sim 35\,\mathrm{M}_\odot$ would shift with redshift (Karathanasis et al., 2023). While this is an interesting approach, our models use a fixed PISN prescription

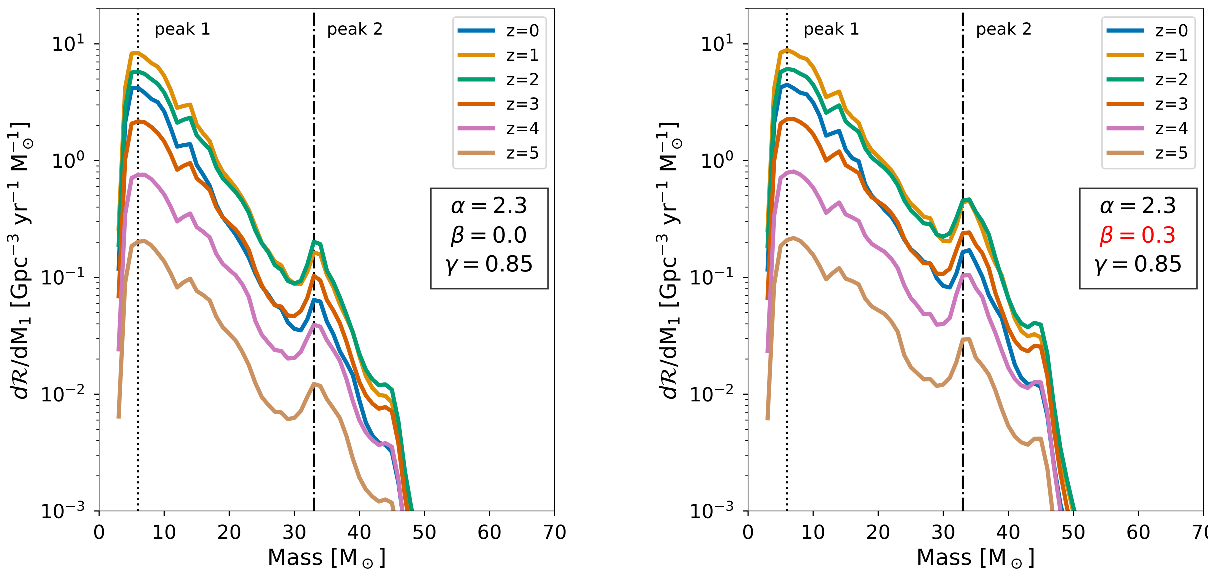

Figure 2.9: Binary black hole primary mass spectrum across cosmic time. Included are results from our constant-slope model (left) and running-slope model (right). The location of the features of the mass spectrum are largely robust against time as can be shown via the vertical lines. However, the peak's relative abundances does vary with time. This can be understood via the redshift evolving metallicity-mass relation (Ma et al., 2016) where high-metallicity low-mass black holes become increasingly abundant relative to low-metallicity high-mass black holes as time evolves.

at each metallicity, so we do not observe such a shift, inline with LVK findings.

2.4.4 Mass Ratio

We conclude our comparison with the LVK's binary black hole observations by analyzing how the merger rate depends on the mass ratio of the black holes, $q_{BH} = \frac{M_2}{M_1}$, as this has been inferred by LVK. To achieve this, we use Equation 2.3 integrating over the full range of primary masses, M_1, rather than mass ratio q:

$$\frac{d^2N}{dq\,dt}(M_\star, z) = \frac{S(z, M_\star)}{\langle M_{\rm ZAMS}\rangle(M_\star, z)} \int_{M_l}^{M_u} dM_1 \int_{Z_l}^{Z_u} dZ\, \rho(Z|M_\star, z) \times \int_{8M_\odot}^{200M_\odot} dM_{\rm ZAMS}\, p(M_{\rm ZAMS}|z, M_\star)\, p(M_1, q|Z, M_{\rm ZAMS})\,, \tag{2.14}$$

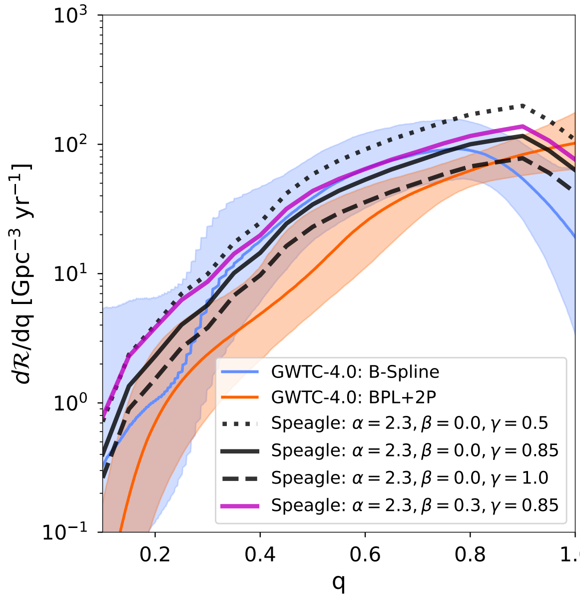

Figure 2.10: Comparison of the mass-ratio distribution for our constant- and running-slope models with the 90% credibility intervals from the LVK *BPL+2P* (orange) and *B-Spline* (blue) models (Abac et al., 2025). Both our models and the LVK non-parametric *B-Spline* analysis exhibit a rise toward equal-mass binaries with a turnover near $q \sim 0.8 - 0.9$, followed by a decline at higher q values.

where, $p(M_1, q \mid M_{\rm ZAMS}, Z)$ represents the joint probability of forming a black hole with primary mass M_1 and mass ratio q, , as described in Section 2.3.1. The limits for the metallicity integral match our simulation limits $Z_l = 1.5 \times 10^{-6}$ and $Z_u = 3.2 \times 10^{-2}$, and the limits for the primary black hole mass are taken from $M_l = 3\,M_\odot$ to $M_u = 80\,M_\odot$. Building on Equation 2.14, we integrate over the delay-time distribution (DTD) and galaxy stellar mass function (GSMF) to obtain:

$$\frac{d\mathcal{R}}{dq} = \int_0^{t_0} \int_{M_{min}=10^8\,\mathrm{M}_\odot}^{M_{max}=10^{12}\,\mathrm{M}_\odot} \phi(M_\star, z(t_0 - \tau')) \frac{d^2 N_{BBH}}{dq\,dt}(M_\star, t_0 - \tau') P(\tau') d\tau' dM_\star \,. \quad (2.15)$$

Our results are generally consistent with the LVK mass ratio spectrum at z=0.2, following the same overall trend of increasing merger rate density with q as shown as shown in Figure 2.10.

We show the 90% C.I. from from LVK's *BPL+2P* and *B-Spline* models. Our models exhibit the best agreement with the non-parametric *B-Spline* model. We include our constant- and running-slope models, with varying delay-time distribution power-law indices. Both reproduce the overall rise toward equal-mass binaries, though our distributions exhibit a turnover at slightly higher mass ratios, $q \sim 0.9$ rather than $q \sim 0.8$, inline with previous inferences Edelman et al. (2023); Farah et al. (2024); Rinaldi et al. (2025). This feature has been linked to stable Roche-lobe overflow systems (van Son et al., 2022) and variations in mass-transfer efficiency in `SEVN` could shift this to lower mass ratio. With the completion of the O4 observing run and an expanded catalog of detections, the shape of the high-q regime will be better constrained, providing an important test of these models.

To connect the primary and secondary mass spectra to the mass ratio distribution, Figure 2.11 shows the joint distribution of M_1 and M_2 for systems in our simulations that form BBHs merging within a Hubble time. This can be compared to the inferred two-dimensional underlying mass distributions from LVK data in Farah et al. (2023) (their Figure 2) and Callister (2024) (their Figure 7). While our simulation outputs do not represent the full astrophysical population, since they have not yet been convolved with our galactic and cosmological relations, they are illustrative for understanding how features in the simulated population propagate into the merger rate.

The simulated M_1–M_2 distribution shows over-densities near $10\,\mathrm{M}_\odot$, $20\,M_\odot$, and $35\,\mathrm{M}_\odot$ in the primary mass spectrum. The lower-mass excesses appear in both $d\mathcal{R}/dM_1$ and $d\mathcal{R}/dM_2$, whereas the $20\,\mathrm{M}_\odot$ excess is largely washed out, though there is a small peak remnant in the primary spectrum. The $35\,\mathrm{M}_\odot$ excess is observable in $d\mathcal{R}/dM_1$ but suppressed in $d\mathcal{R}/dM_2$, as discussed in Section 2.4.2. Furthermore, the M_1–M_2 space reveals how systems cluster at various q values. In particular, our simulations show a steady rise in the number of systems with increasing mass ratio up to $q = 1$. However, as shown in Figure 2.10, the actual merger rate peaks at slightly smaller values of $q \sim 0.9$, underscoring how the raw simulation outputs

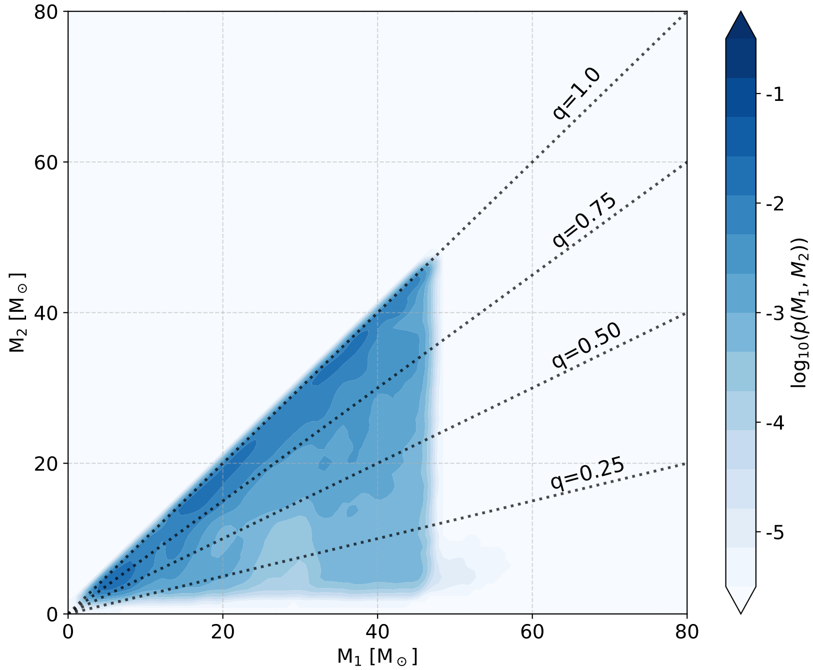

Figure 2.11: Joint distribution of primary and secondary black hole masses (M_1–M_2) for binaries that merge within a Hubble time in our simulations. The color scale indicates the normalized density $\log_{10}(\rho(M_1, M_2))$, and black dotted lines mark contours of constant mass ratio $q = M_2/M_1$. Binary systems cluster toward $q \to 1$, reflecting the tendency for near equal-mass mergers. Overdense regions appear near $M_1 \sim 10\,M_\odot$, $20\,M_\odot$, and $35\,M_\odot$, consistent with the primary mass features inferred by LVK.

are reshaped into an astrophysical population once our galactic and cosmological inputs are applied.

2.5 Conclusion

Our approach, grounded in observed galactic and cosmological relations, utilizes the `SEVN` stellar evolution code to construct a model of the binary black hole (BBH) population, which are found to merge within the age of the Universe. We folded in the redshift dependent relations for the metallicity mass relation (MZR), galaxy stellar mass function (GSMF), and star-forming main sequence to compute the binary black hole merger rate density. Our results show good agreement with the LVK observations, in particular with their *B-Spline* model, across the total and differential merger rates with respect to primary mass, secondary mass, and mass ratio. This work will help create better model predictions for comparison to the additional BBH mergers ($\gtrsim 200$) from observing run 4 (Abbott et al., 2020; Kiendrebeogo et al., 2023; Broekgaarden et al., 2024).

There are significant uncertainties in the various elements of our calculation, including the star formation history, the initial mass function, and the physical processes in stellar evolution theory, such as mixing, wind mass loss, and common envelope phase as modeled by the `SEVN` code (Iorio et al., 2023a). While these factors introduce variability, they also present opportunities to gain insight. Studies to systematically quantify the impact of these uncertainties have been made by Chruslinska et al. (2019); Neijssel et al. (2019a); Spera et al. (2019); Tang et al. (2020); Santoliquido et al. (2020); Broekgaarden et al. (2022); Briel et al. (2022); Iorio et al. (2023a); van Son et al. (2023a), who explored the effects of uncertain stellar evolution and galactic and cosmological relations across cosmic time on binary compact object merger rates. These uncertainties combined can alter the merger rate upwards of an order of magnitude (Neijssel et al., 2019a). Our aims in this paper are to elucidate how the various elements of the calculation come together to produce the distribution of binary black holes in a simple way, and to highlight the importance of the metallicity distribution in galaxies and other key features summarized below. We show that it is difficult to match the LVK inferred distribution through the isolated channel alone and what is needed of star formation and stellar physics to get close. Our work further shows the need for a population arising from the dynamical channel, especially at higher masses.

The predicted merger rate as a function of redshift is consistent with observations, as shown in Figure 2.8, where our model mostly lies within the 90% confidence interval across cosmic time. This is encouraging given the simplicity of the model and the reasonable value for the average core collapse supernova rate used to normalize the merger rate. Our predicted rate is robust to redshift-dependent variations in the GSMF and MZR, as the most significant uncertainties in these quantities arise at higher redshifts. This insensitivity stems from the relatively short delay-time distribution (DTD) in our model, with most mergers occurring within 1 Gyr (Figure 2.5) of BBH formation. As a result, the majority of systems contributing to the observed LVK merger are expected to originate from lower redshifts, where the GSMF and MZR are better constrained and less variable.

A more significant source of uncertainty lies in the form of the DTD itself. We adopt a power-law parameterization, $P(\tau) \propto \tau^{-\gamma}$, though the precise value of the spectral index γ remains unconstrained. While Fishbach & Kalogera (2021) demonstrate that the LVK data are consistent with $\gamma = 1.0$, our simulations favor a slightly shallower distribution with $\gamma = 0.85$. We also include a more extreme value of $\gamma = 0.5$ in our calculations which lies just outside the 90% credibility interval from Mukherjee & Moradinezhad Dizgah (2022) ($\sim 1.6\sigma$ for a Gaussian posterior). We explore the impact of varying γ and find that steeper values (e.g., 0.85 and 1.0) yield better agreement between our model and the LVK-inferred merger rate evolution. That our total merger rate peaks at smaller redshifts than the inferred distribution, and the star formation rate (see Appendix 2A), stems from our choice of GSMF and star-forming main sequence, the convolution with the DTD, and our choice of stellar prescriptions which have been shown to shift the peak of the total merger rate (Boesky et al., 2024a,b).

Furthermore, comparing our model against the inferred differential merger rates yields good agreement. The black hole mass spectrum, $d\mathcal{R}/dM_1$, has several features rich with information (Abbott et al., 2023; Tiwari, 2022; Farah et al., 2023). The smooth tapering at low masses is apparent, driven by binary effects and the decreasing probability of lower-mass black holes being the primary in a BBH system. We obtain a small peak at $\sim 20\,M_\odot$, consistent with the LVK's *B-Spline* model. However, the observed abundance of black holes in the $30-40\,M_\odot$ range presents a challenge for the model with a constant power-law stellar initial mass function.

We explored the possibility that the average IMF is top-heavy by allowing the power-law index of the IMF to vary with mass. We found that a running-slope model ($-2.3 + \beta \log_{10}(M_{\rm ZAMS}/8{\rm M}_\odot)$) is a better fit to the inferred primary mass distribution, with the best results for $\beta = 0.3$. Most high-mass binary black holes in our model arise in dwarf galaxies, which are metal-poor, and hence this result is essentially arguing for a more top-heavy IMF in dwarf galaxies. Furthermore, Tanikawa et al. (2021, 2022) find that coupling their top-heavy IMF with the assumption of inefficient convective overshooting can produce mass gap BHs. These results are sensitive to the evolutionary model they choose, for instance utilizing the so-called L model from Yoshida et al. (2019), the progenitors of mass gap BHs no longer retain their hydrogen envelopes and are not formed. Similar results have been found using SEVN where Iorio et al. (2023a) show that varying the core overshooting parameter and the pair instability model can nearly double the final maximum mass of a black hole. If the results in (Tanikawa et al., 2021, 2022) hold for higher metallicity systems, like those in this work, then it would be possible for the isolated channel to produce mass gap BHs.

These findings are qualitatively aligned with observations of the top-heavy IMFs inferred in metal-poor galaxies and globular clusters (Marks et al., 2012; Geha et al., 2013; Gennaro et al., 2018; Weatherford et al., 2021) and the work in Chruślińska et al. (2020) who highlight the importance of a galaxy/metallicity varying IMF. Alternatively, these black holes may predominantly form through dynamical capture mechanisms (Antonini et al., 2022; Banerjee, 2022; Torniamenti et al., 2024) or result from the possible redshift evolution of the PISN mass scale and maximum black hole mass (Mukherjee, 2022). The impacts of these disparate scenarios on the mass ratio distribution and the redshift distribution may offer clues to distinguish them.

Addressing the high-mass end of the merger-rate spectrum remains a challenge and is an active area of research. As mentioned above, the peak at $\sim 35\,{\rm M}_\odot$ could be populated by a subpopulation of black holes forming in globular clusters (Antonini et al., 2022; Banerjee, 2022; Torniamenti et al., 2024). In our analysis, this bump arises due to the pair instability mechanism and hence becomes more prominent when considering a top-heavy IMF (running-slope model). However, we have treated each galaxy as a monolithic entity, neglecting diverse components such as globular clusters, star clusters, bulges, and disks. These components exhibit varying densities, metallicities and stellar populations, leading to differences in merger rates and dynamical interactions. For instance, Chowdhury & Santra (2024) show that differences in metallicity between the primary and secondary components within stellar clusters lead to the formation of more massive compact objects and an increased frequency of eccentric, inspiraling binary systems. This kind of diversity could contribute to filling the mass gap through hierarchical mergers and to the formation of a more prominent peak at $\sim 35\,\mathrm{M}_{\odot}$. A more granular examination of these factors, set within the framework of galactic relations, is left for future work.

Beyond $50\,\mathrm{M}_{\odot}$, our model predicts a sharp decline in the number of BBH systems, revealing the well-known mass gap caused by the pair-instability mechanism and mass loss in binary systems. Incorporating mergers in dynamically-formed BBHs originating from the single black hole population could help fill this gap. We estimate that an efficiency for these dynamical processes of $\sim 0.01\%$ could impact the high-mass end, making it more comparable to observations up to about 60 $M_{\odot}$ in primary black hole mass. This coupled with the uncertainties in the location of the pair-instability mass gap can shift our results further more.

We also direct the reader to previous studies on possible explanations for the high-mass end of the distribution, including population III stars (Farrell et al., 2021; Kinugawa et al., 2021; Liu & Bromm, 2020; Tanikawa et al., 2021; Safarzadeh & Ramirez-Ruiz, 2021), hierarchical mergers (Sedda et al., 2023; Gerosa & Fishbach, 2021; Mapelli et al., 2021; Farr et al., 2017; Antonini et al., 2022; Torniamenti et al., 2024), primordial black holes (De Luca et al., 2021), and beyond the Standard Model (BSM) physics (Croon et al., 2020; Sakstein et al., 2020; Ziegler & Freese, 2021, 2024; Croon & Sakstein, 2024).

The redshift evolution of the primary mass spectrum is investigated and we find that the positions of the $\sim 6\,\mathrm{M}_\odot$ and $\sim 35\,\mathrm{M}_\odot$ peaks in our model remain fixed over cosmic time. The primary change is a gradual flattening of the slope between the peaks toward higher redshifts, reflecting the larger contribution from metal-poor environments at earlier times. Such conditions favor the formation of more massive black holes relative to the low-mass high-metallicity systems that dominate at lower redshifts. This stability in peak location, coupled with metallicity-driven slope evolution, is consistent with LVK observations (Fishbach et al., 2021; Abbott et al., 2023; Lalleman et al., 2025), suggesting that the underlying mass scales set by stellar physics remain constant in the isolated channel.

We also explore how our model compares to the secondary mass spectrum and mass ratio distribution. The secondary spectrum has been largely unexplored and a BBH-specific inferred rate from LVK has not been presented. Recent studies by Farah et al. (2024); Sadiq et al. (2024) explore inferences from LVK observations concerning the properties of secondary components. We find that our calculated secondary mass spectrum based on the astrophysical population has weaker peak structure when compared to the primary mass spectrum. For example, the peak that is prominent at $\sim 35\,\mathrm{M}_\odot$ in $d\mathcal{R}/dM_1$ has been suppressed in $d\mathcal{R}/dM_2$. We find that both spectra receive comparable contributions from low-metallicity systems $10^{-3} \leq Z < 10^{-4}$ associated with mass ratios near unity. However, $d\mathcal{R}/dM_1$ receives enhanced contributions from intermediate metallicity systems $10^{-2} \leq Z < 10^{-3}$ with mass ratios distributed more broadly, leading to the suppression in $d\mathcal{R}/dM_2$.

For the mass ratio distribution, we compare to LVK's Broken Power Law + Two Peaks (BPL+2P) and B-Spline models (Abac et al., 2025). We find that our calculation yields a much richer mass ratio distribution than the previous GWTC-3 results when this paper was submitted, and find structure similar to the new GWTC-4 *B-Spline* model and previous non-parametric studies (Edelman et al., 2023; Rinaldi et al., 2025). For instance, peaking at q below unity, $q \sim 0.8 - 0.9$ and slight kinks throughout the spectrum. Further O4 results will shed light on the high mass ratio regime.

In summary, our study demonstrates that a population model grounded in galactic relations, such as the galaxy stellar mass function, metallicity distribution, and star formation, can get the correct number of mergers with peaks in the mass spectrum near $\sim 10\,\mathrm{M}_{\odot}$ and $\sim 35\,\mathrm{M}_{\odot}$ and redshift evolution of BBHs inferred by the LVK collaboration. Our modeling suggests that for the isolated channel alone, a top-heavy IMF in low-mass metal-poor galaxies would be required to match the primary mass spectrum up to about $50\,\mathrm{M}_{\odot}$. Alternately, we would need to include dynamical capture mechanisms such as those that could happen in globular clusters. Finally, our mass ratio distribution does not follow a strict power-law assumed by previous LVK O3 results, but is in agreement with the newly released LVK O4a *B-Spline* model. Although the uncertainties in stellar evolution and star formation history are still large, our findings demonstrate that the bulk of the BBH mergers with primary mass below $\sim 30\,\mathrm{M}_{\odot}$ observed by LVK can be explained by a simple model that averages over the observed properties of galaxies. This allows for a more robust assessment of the mismatch between the predictions and observations of black holes at the high-mass end.

Appendix 2A: Star Formation Rate Density

In this appendix, we discuss how the star-forming main sequence from Speagle et al. (2014), when integrated with the star-forming galaxy stellar mass function (GSMF) from Weaver et al. (2023), produces the evolution of the cosmic star formation rate density (SFRD). This provides a consistency check for our adopted parameterizations of both the main sequence and the GSMF across cosmic time.

Figure 2.12 shows the resulting SFRD (red curve) compared with the empirical fits from

Madau & Dickinson (2014) (MD14; gray dashed), Madau & Fragos (2017) (MF17; black solid), and Driver et al. (2018) (D18; orange dash-dot). The vertical gray lines indicate transitions between regimes in our adopted GSMF: (1) the low-redshift double-Schechter fit from Baldry et al. (2012) applied at $z < 0.2$, (2) the switch from the double- to single-Schechter fit from Weaver et al. (2023), and (3) the transition to the high-redshift extrapolation based on our power-law fits to the evolution of the Schechter parameters.

Our star-formation rate density peaks at $z \sim 1.5$, shifted from MD14 and M17 but in close agreement with D18, based on the GAMA/G10–COSMOS/3D–HST compilation, and with the modeled SFRD of Behroozi et al. (2019). This indicates a slightly later turnover than the canonical $z \sim 2$. These analyses suggest that the global SFRD maximum is broader and shifted to lower redshift than previously assumed, consistent with our results and supporting our adopted framework.

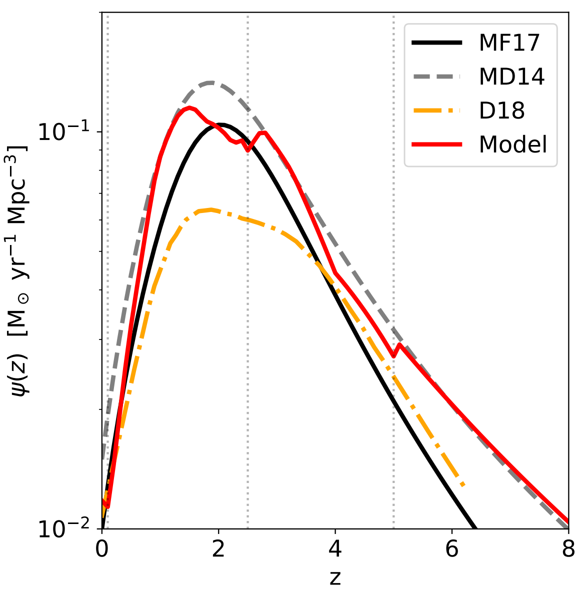

Figure 2.12: Cosmic star formation rate density (SFRD) derived from the convolution of the Speagle et al. (2014) star-forming main sequence and the Weaver et al. (2023) galaxy stellar mass function. The red curve shows our model prediction, compared with empirical fits from Madau & Dickinson (2014) (MD14; gray dashed), Madau & Fragos (2017) (M17; black solid), and Driver et al. (2018) (D18;orange). Vertical gray lines mark the transitions between GSMF regimes: (1) low-redshift Baldry et al. (2012), (2) the switch from double- to single-Schechter fits in Weaver et al. (2023), and (3) the onset of our high-redshift power-law extrapolation. Our model reproduces the broad SFRD peak around $z \approx 1.5$, consistent with recent analyses by D18 and Behroozi et al. (2019).

Chapter 3

Massquerade: Impacts of Mass Ratio Reversals on Binary Black Hole Merger Rates and Mass Distributions

3.1 Abstract

Mass ratio reversal (MRR), in which the initially less massive star in a binary ultimately forms the more massive compact object, is a well-established outcome of interacting binary evolution. We investigate its role in shaping the binary black hole (BBH) merger rate and mass distribution inferred by LIGO–Virgo–KAGRA, using population synthesis frameworks `COMPAS` and `SEVN`. The observational imprint of MRR differs qualitatively: in `COMPAS`, MRR systems dominate the merger rate density at high primary masses ($M_1 \gtrsim 20\ \mathrm{M}_\odot$), secondary

This chapter is as it was submitted to the Astrophysical Journal and is currently under review, with expected publication later this year. The coauthors on this work are Floor Broekgaarden, Sasha Levina, Amedeo Romagnolo, Manasvini Komandur, Melanie Santiago, and Kyle A. Rocha.

masses ($M_2 \gtrsim 12$ M$_\odot$), and mass ratios ($q > 0.6$), whereas in `SEVN` they remain subdominant across the BBH mass distribution. This implies that the initially less massive star can *massquerade* as the observed primary black hole, such that the primary-mass distribution is not a direct tracer of the initially more massive stars, but instead a superposition of physically distinct evolutionary populations. We identify three evolutionary pathways leading to MRR: *core-growth*, where stable mass transfer increases the helium-core of the secondary; *PPISN-shrinking*, where pulsational pair-instability episodes reduce the primary remnant mass; and *asymmetric-CCSN*, where differential supernova mass loss drives the reversal. When weighted by the local merger-rate density, the core-growth channel dominates. MRR systems predominantly originate from massive ($\gtrsim 50$ M$_\odot$) low-metallicity progenitors ($\lesssim 0.1$ $Z_\odot$). MRR is a physically distinct, potentially observable feature of isolated binary evolution that must be accounted for to robustly connect future gravitational-wave observations to massive binary evolution and compact-object formation.

3.2 Introduction

The LIGO–Virgo–KAGRA (LVK) collaboration has detected over 150 confirmed gravitational-wave events through the first part of observing run 4 (O4a), primarily from binary black hole (BBH) mergers (Abbott et al., 2019a; Abbott et al., 2021; Abbott et al., 2021a; Abac et al., 2025). These observations are commonly summarized in terms of the total merger rate density (MRD), $\mathcal{R}(z)$, as well as differential merger rates with respect to intrinsic binary parameters, such as the component BH masses. In particular, LVK population analyses characterize the BBH population using the primary- and secondary-mass distributions, $d\mathcal{R}/dM_1$ and $d\mathcal{R}/dM_2$, where the primary (secondary) black hole is defined as the more (less) massive component of the system.

From the perspective of binary evolution, however, the two compact objects often experience systematically different evolutionary histories: the initially more massive star typically evolves first and acts as the donor during the first mass-transfer phase, while the initially less massive star may accrete mass and angular momentum before collapsing (Renzo & Götberg, 2021). As a result, one may expect the observed M_1 and M_2 distributions to approximately trace two distinct evolutionary populations of donors and accretor stars.

We show in this work that this mapping can break down in the presence of *mass ratio reversal* (MRR), where the initially less massive star ultimately forms the more massive black hole. This outcome is not unusual, but rather a well-established outcome of short-period binary evolution. In particular, systems such as Algol binaries exhibit clear observational evidence of MRR (e.g. Crawford, 1955; Morton, 1960; Smak, 1962; Plavec, 1968; Paczyński, 1971; Budding, 1986; Harries et al., 2003; de Mink et al., 2007; Sen et al., 2022, 2023, 2026). This phenomenon has also been observed in compact-object studies, e.g. in neutron star–white dwarf (NS–WD) binaries (Tauris & Sennels, 2000; Toonen et al., 2018), BH–NS binaries (Sipior et al., 2004; Broekgaarden et al., 2021), and BBH studies (Gerosa et al., 2013, 2018; Olejak & Belczynski, 2021; Broekgaarden et al., 2022; Zevin & Bavera, 2022; Olejak et al., 2024; Banerjee & Olejak, 2024; Smith & Kaplinghat, 2026), demonstrating that MRR is a generic consequence of binary evolution. Recently, MRR has been studied in the context of gravitational-wave observations where characteristic mass-spin correlations have been used to constrain the fraction of LVK sources which arise from MRR (Mould et al., 2022). The majority of BBH studies focus on MRR as related to the spins of BBH systems. For instance, Broekgaarden et al. (2022) found that MRR can occur in $\gtrsim 70\%$ of observable LVK sources and between 11–82% of the astrophysical population, depending on stellar, binary, and star formation prescriptions used. In this work, MRR systems preferentially populate the primary-mass spectrum above $\sim 20\,\mathrm{M}_{\odot}$, and the second-born BH acquires non-negligible spin in up to $\sim 25\%$ of BBH systems. The latter finding is consistent with expectations that the initially more massive star produces a weakly spinning remnant, while the secondary star may be tidally spun up prior to collapse (Kushnir et al., 2016; Hotokezaka & Piran, 2017; Qin et al., 2018; Fuller & Ma, 2019; Bavera et al., 2022; Belczynski et al., 2020).

Building on these results, we investigate how MRR shapes the BBH merger-rate and mass distributions predicted by population-synthesis models. We quantify the MRR contribution to the BBH merger rate as a function of redshift, primary mass, secondary mass, and mass ratio, and assess the implications for interpreting BBH observations. We further examine the progenitor properties and formation channels of MRR BBH systems, identifying the physical mechanisms driving MRR and their BBH mass distribution signatures. Finally, we compare predictions from two fundamentally different population-synthesis frameworks: `COMPAS` and `SEVN`, to investigate how assumptions in stellar and binary evolution shape the prevalence and observational imprint of MRR BBH systems.

3.3 Methodology

3.3.1 Population Synthesis Models

In order to determine the MRR contribution to the astrophysical merger rate, we first construct a population of merging BBHs using two distinct population synthesis frameworks, `COMPAS` and `SEVN`, which employ contrasting approaches to stellar evolution (analytic fitting formulae versus interpolating tracks *on-the-fly*, respectively). Specifically, we include the fiducial and single-parameter variation `COMPAS` models from Broekgaarden et al. (2022), in which MRR has been studied in Broekgaarden et al. (2022) finding that MRR comprises 11–82% of the astrophysical population depending on the model choice. We compare this to the results from `SEVN`, where MRR systems were previously found to comprise $\sim 20\%$ of the astrophysical population (Smith & Kaplinghat, 2026). These studies indicate that there can be upwards of 60% difference in their calculated MRR contributions.

We note that the two codes will have inherent differences as they approach stellar evolution in contrasting ways, i.e. `COMPAS` employs analytic fitting formulae from Hurley et al. (2000) based on the tracks from Pols et al. (1998) (with masses up to 50 $\mathrm{M}_{\odot}$ at 7 distinct metallicities) vs interpolating these quantities *on-the-fly* from PARSEC tracks (Bressan et al., 2012; Tang et al., 2014; Chen et al., 2015; Marigo et al., 2017; Nguyen et al., 2022) (with masses up to 600 $M_\odot$ at 15 distinct metallicities) in `SEVN`. This effect alone can largely impact where MRR contributes across the component mass distributions, however we perform a systematic comparison with a consistent cosmic integration, i.e. star-formation history, which results in a similar overall MRR contribution to the astrophysical population ($\sim 30\%$ level) across the two codes. We further extend the two studies mentioned above by including not only the primary- and chirp-mass distributions, but the secondary-mass and mass ratio distributions with a full MRR contribution decomposition.

COMPAS simulations

`COMPAS` (Stevenson et al., 2017; Barrett et al., 2018; Vigna-Gómez et al., 2018; Broekgaarden et al., 2019; Neijssel et al., 2019b) utilizes parameterized models of binary evolution from Hurley et al. (2002). Specifically, we use the publicly available models from Broekgaarden et al. (2022), which include 20 single parameter variation on massive (binary) star evolution and are publicly available at Broekgaarden (2021). In particular, we use the fiducial dataset and variations (B,C,D,G,I,J,L,P,Q,S, and T) from Broekgaarden et al. (2022). The initial parameters are sampled according to Sana et al. (2012) within the following ranges: $\log_{10}(Z/Z_\odot) \in [-2.3, -0.38]$, $M_{1,\mathrm{ZAMS}}/\mathrm{M}_\odot \in [5, 150]$, and orbital separation $a/\mathrm{AU} \in [0.01, 1000]$.

SEVN simulations

We compare these results with `SEVN`, which calculates stellar properties via the interpolation of precomputed evolutionary tracks (Spera et al., 2015; Spera & Mapelli, 2017; Spera et al., 2019; Mapelli et al., 2020; Iorio et al., 2023b). We use version `SEVN v2.13.0` with the default parameters adopted in the `PARSEC` tracks with overshooting parameter $\lambda = 0.5$ (Bressan et al., 2012; Tang et al., 2014; Chen et al., 2015; Marigo et al., 2017; Nguyen et al., 2022). Similar to COMPAS, the initial parameters were sampled following Smith & Kaplinghat (2026) based on Sana et al. (2012) in the following ranges: $\log_{10}(Z/Z_\odot) \in [-4.0, 0.18]$, $M_{1,\mathrm{ZAMS}}/\,\mathrm{M}_\odot \in [3, 200]$, and $a/\mathrm{AU} \in [0.07, 670]$.

3.3.2 Reanalyzing the Simulations Under Consistent Star Formation History Modeling

In order to mitigate discrepancies due to the chosen cosmic integration, i.e. the assumptions for the metallicity-dependent star formation history to calculate cosmological BBH merger rates, we use the SSPC (Hendriks et al., 2023) code and recalculate the BBH merger properties across cosmic time for both the SEVN and COMPAS simulations under similar parameter settings, an exploration of single-parameter variations is provided in the results section. This allows for a standardized cosmic integration based on ingredients such as the star formation rate and metallicity-mass relation. In particular, in population synthesis studies the BBH merger rate can be calculated by a convolution of the BBH formation efficiency and a cosmic star formation history following

$$\begin{aligned} \mathcal{R}_{\mathrm{BBH}}(z_{\mathrm{merge}},\,\zeta) = \int_{Z_{\min}}^{Z_{\max}} dZ \int_{0}^{t^*_{\mathrm{first\,SFR}} - t^*_{\mathrm{merge}}} & dt_{\mathrm{delay}} \\ \times \mathcal{N}_{\mathrm{form}}(Z, t_{\mathrm{delay}}, \zeta)\, \mathcal{S}(Z, z_{\mathrm{birth}})\,. & \end{aligned} \tag{3.1}$$

The inner integral is over the delay time, t_{delay}, the time it takes from birth of the binary system to compact object merger. The upper limit is the difference between the time of the first star formation ($t^*_{\mathrm{first\,SFR}}$), based on the specific SFRD assumptions, and the merger time (t^*_{merge}). We follow the notation from Hendriks et al. (2023) where a $*$ superscript indicates lookback time, while times without a $*$ superscript correspond to a duration. The outer integral over metallicity is evaluated over the metallicity range spanned by our simulations, see Section 3.3.1 and Section 3.3.1 for exact values. The function $\mathcal{N}_{\text{form}}(Z, t_{\text{delay}}, \zeta)$ is the BBH formation efficiency, i.e. the number of BBH systems, N_{BBH}, formed per solar mass:

$$\mathcal{N}_{\text{form}}(Z, t_{\text{delay}}, \xi) = \frac{N_{\text{BBH}}(Z, t_{\text{delay}}, \xi)}{M_{\text{total,formed}}(Z)} \tag{3.2}$$

and is dependent on the metallicity, Z, delay time, t_{delay}, and system properties, ζ, e.g. component masses and orbital parameters. We assume a binary fraction of unity throughout.

The star-formation rate, $\mathcal{S}(Z, z_{\text{birth}})$ is defined as a function of Z and birth redshift, z_{birth}. It is the combination between the star-formation rate density and the redshift-dependent metallicity distribution:

$$\mathcal{S}(Z, z) = \text{SFRD}(z) \times \frac{dP}{dZ}(Z, z)\,, \tag{3.3}$$

where we use the functional form from Madau & Dickinson (2014) for our SFRD(z):

$$\text{SFRD}(z) = a\frac{(1+z)^b}{1 + [(1+z)/c]^d}\,, \tag{3.4}$$

and the following parameters $(a, b, c, d) = (0.02, 1.48, 4.45, 5.90)$ from van Son et al. (2022). The metallicity distribution, $\frac{dP}{dZ}(Z, z)$, follows van Son et al. (2023a):

$$\frac{dP}{dZ}(Z, z) = \frac{2}{Z}\,\phi\left(\frac{\ln Z - \xi(z)}{\omega(z)}\right)\Phi\left(\alpha\,\frac{\ln Z - \xi(z)}{\omega(z)}\right), \tag{3.5}$$

with ϕ the standard log-normal distribution, and Φ the cumulative distribution function of the standard log-normal distribution. The width is taken to be $\omega(z) = \omega_0 \cdot 10^{\omega_z \cdot z}$, with $\omega_0 = 1.125$ and $\omega_z = 0.05$. The function $\xi(z)$ is given as:

$$\xi(z) = \ln\left(\frac{\mu_0 \cdot 10^{\mu_z \cdot z}}{2\Phi(\beta\, \omega(z))}\right) - \frac{\omega(z)^2}{2}\,, \tag{3.6}$$

with $\mu_0 = 0.025$ and $\mu_z = -0.05$ and the parameter $\beta = \alpha/\sqrt{1+\alpha^2}$ with $\alpha = -1.77$.

This calculation is carried out in `SSPC` for both `COMPAS` and `SEVN`, providing a standardized calculation of the merger rate density and its derivatives.

3.3.3 Selecting Mass Ratio Reversal Candidates

Throughout, we will use the criteria for MRR that the mass of the black hole descending from the initially less massive star (B) needs to be larger than the initially more massive star's black hole (A), i.e. $\mathrm{M_B} > \mathrm{M_A}$. While this includes systems which barely pass the threshold, we show in Appendix 3A that our results remain largely the same when imposing a conservative threshold of $\mathrm{M_B} > 1.05 \times \mathrm{M_A}$. As we will show, MRR predictions differ substantially across the two population synthesis results, with discrepancies going beyond those found by varying this threshold.

Observations

We compare our population synthesis results with population models based on gravitational-wave observations from the LVK collaboration's GWTC-4 (Abac et al., 2025). In particular, we compare to the inferred primary-mass and mass ratio distributions from the non-parametric B-Spline model and provide an additional comparison in the joint M_1–M_2 space from LVK based on their FullPop-4.0 model. The LVK collaboration does present an inferred secondary-mass distribution, however it is limited to $M_2 \lesssim 15\ \mathrm{M}_\odot$. As such we compare our secondary-mass distribution with the results from Sadiq et al. (2024) whose non-parametric population model is based on GWTC-3 (Abbott et al., 2023).

3.4 Results

3.4.1 Imprint of Mass Ratio Reversal on BBH Merger Rates and Mass Distributions

We investigate how MRR imprints on both the overall BBH merger rate and the distribution of BBH masses. Across both population synthesis frameworks, MRR constitutes a substantial fraction of merging systems, contributing to $\sim 1/3$ of the total merger rate. However, its impact on observable BBH properties differs qualitatively between models: `COMPAS` predicts that MRR systems dominate the high-mass and high-mass-ratio regimes, whereas `SEVN` predicts a more diffuse contribution that remains subdominant across the whole mass range. In this section, we first examine the redshift evolution of the merger rate and then quantify how MRR shapes the primary mass, secondary mass, and mass ratio distributions.

MRR contribution to the merger rate

Figure 3.1 shows the expected total BBH merger rate density $\mathcal{R}_{\mathrm{BBH}}(z)$ for the fiducial `COMPAS` and `SEVN` models, decomposed into MRR and non-MRR contributions and compared to the LVK O4a B-spline population model (Abac et al., 2025). In both frameworks, the total merger rate rises toward $z \sim 1$–2, broadly tracking the cosmic star formation density (SFRD), while the fractional contribution from MRR systems remains approximately constant with redshift for `SEVN`, `COMPAS` exhibits a mild decrease in the MRR fraction toward higher redshift. In both models, MRR systems contribute roughly one-third of the total merger rate across cosmic time. The near-constant MRR fraction with redshift suggests that the physical processes driving mass ratio reversal operate efficiently across a broad range of metallicities and cosmic epochs, rather than being confined to specific formation environments.

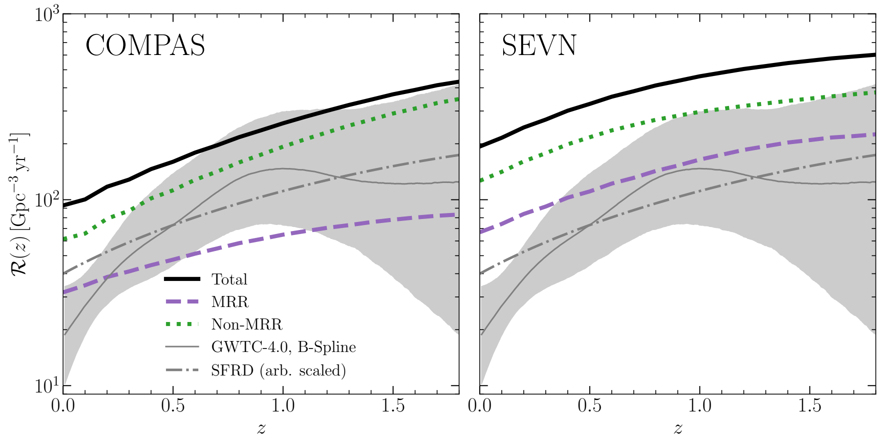

Figure 3.1: Redshift evolution of the BBH merger rate density $\mathcal{R}(z)$ expected by the fiducial `COMPAS` (left) and `SEVN` (right) models. The total merger rate (black) is decomposed into contributions from MRR (purple) and non-MRR (green) systems, and compared to the LVK O4a B-spline population model (gray band; Abac et al. 2025). In both models, MRR systems contribute approximately one-third of the total BBH merger rate across cosmic time, with only weak redshift evolution. While the overall normalization differs from LVK constraints, both models broadly track the rise and fall of the cosmic star formation history (gray dashed-dotted line).

At $z \sim 0$, `COMPAS` expects a local merger rate of $R_{\rm BBH} = 93\ {\rm Gpc}^{-3}\,{\rm yr}^{-1}$, with 34% arising from MRR systems, while `SEVN` predicts $R_{\rm BBH} = 200\ {\rm Gpc}^{-3}\,{\rm yr}^{-1}$, with a similar MRR fraction of $\sim 33\%$. These values exceed the LVK-inferred range of $R_{\rm BBH} = 14$–$26\ {\rm Gpc}^{-3}\,{\rm yr}^{-1}$ for the local BBH merger rate, built from the union of their FullPop-4.0 and BGP models Abac et al. (2025). This reflects known sensitivities of population synthesis predictions to assumptions such as the cosmic star formation history, metallicity evolution, and binary fraction that can impact the *magnitude* of the BBH merger rate by orders of magnitude (e.g. Neijssel et al., 2019a; Broekgaarden et al., 2022; van Son et al., 2023a; Santoliquido et al., 2020; Chruślińska, 2024; Mandel & Broekgaarden, 2022; van Son et al., 2023b; Romagnolo et al., 2023; Sgalletta et al., 2025; Romagnolo et al., 2025; Sgalletta et al., 2026, and references therein). Our focus in this work is therefore primarily on the *relative* contribution of MRR systems and their imprint on the BBH mass distributions, rather than on the absolute normalization of the merger rate.

MRR contribution to BBH masses

Figure 3.2 shows the intrinsic BBH merger rate density as a function of component masses (M_1, M_2) and mass ratio $q = M_2/M_1$ at $z \approx 0.2$, decomposed into MRR and non-MRR contributions for the fiducial `COMPAS` and `SEVN` models. The primary and secondary masses are defined following the LVK convention, with $M_1 \geq M_2$. Figure 3.3 shows the corresponding two-dimensional distribution of the MRR fraction in the (M_1, M_2) plane. The two population synthesis frameworks predict qualitatively different roles for MRR in shaping the BBH mass distributions.

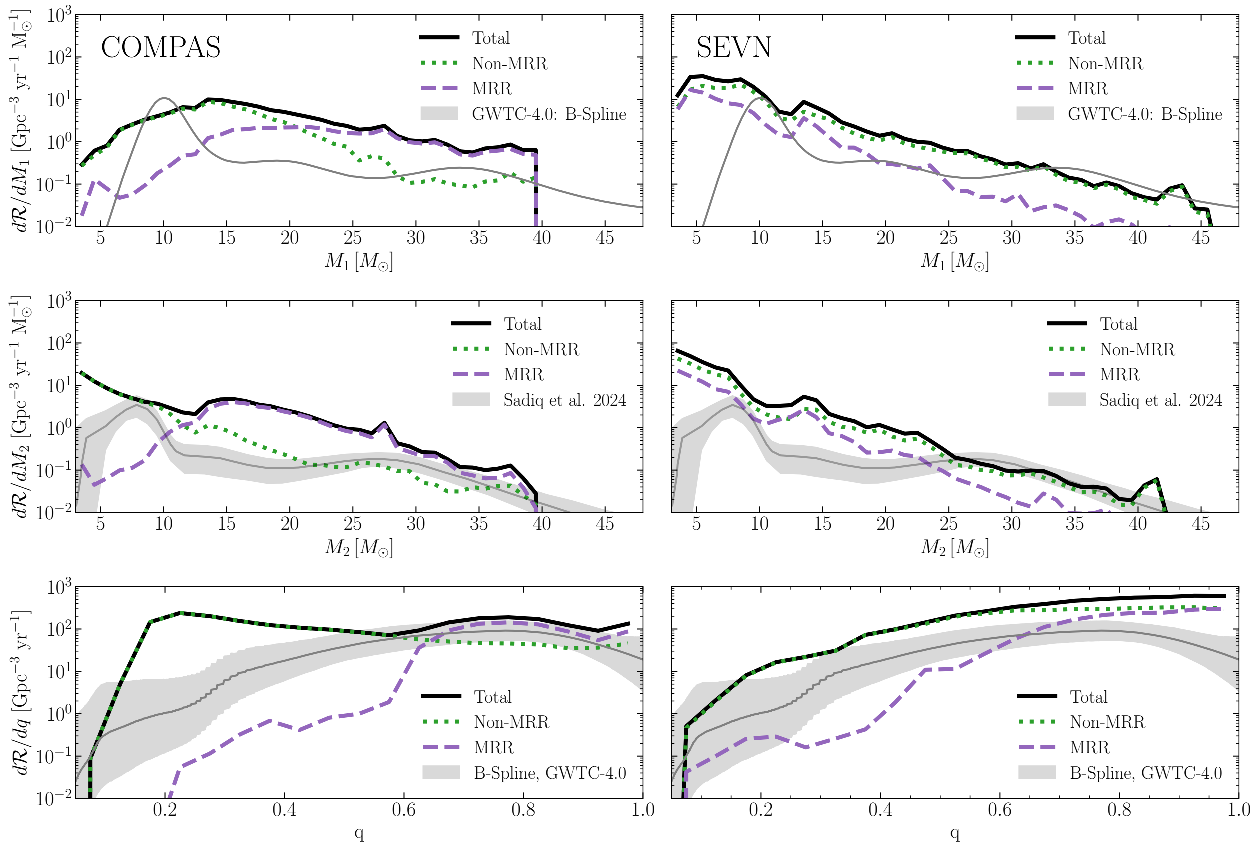

Figure 3.2: Intrinsic BBH merger rate density as a function of primary mass M_1 (top), secondary mass M_2 (middle), and mass ratio q (bottom) at $z \approx 0.2$, for `COMPAS` (left) and `SEVN` (right). The total population (black) is decomposed into MRR (purple) and non-MRR (green) contributions, and compared to LVK-inferred constraints (gray shaded regions; Abac et al. 2025; Sadiq et al. 2024). `COMPAS` predicts that MRR systems dominate the high-mass regime ($M_1 \gtrsim 20\,\mathrm{M}_\odot$, $M_2 \gtrsim 12\,\mathrm{M}_\odot$), whereas `SEVN` predicts a more diffuse contribution that remains subdominant across the mass range. In contrast, both models show that MRR preferentially populates the high mass-ratio regime ($q \gtrsim 0.6$).

In the fiducial COMPAS model, MRR systems dominate the high-mass regime and reshape both component-mass distributions. The primary-mass distribution exhibits a peak near $15\,\mathrm{M}_\odot$ dominated by non-MRR systems, while a secondary enhancement at 25–$35\,\mathrm{M}_\odot$ is driven by MRR systems, which dominate above $M_1 \gtrsim 20\,\mathrm{M}_\odot$. A similar trend is present in the secondary-mass distribution, where MRR systems dominate for $M_2 \gtrsim 12\,\mathrm{M}_\odot$. This behavior is even more apparent in the two-dimensional (M_1, M_2) plane (Figure 3.3), where the MRR contribution is concentrated at high masses and high mass ratios ($q \gtrsim 0.6$).

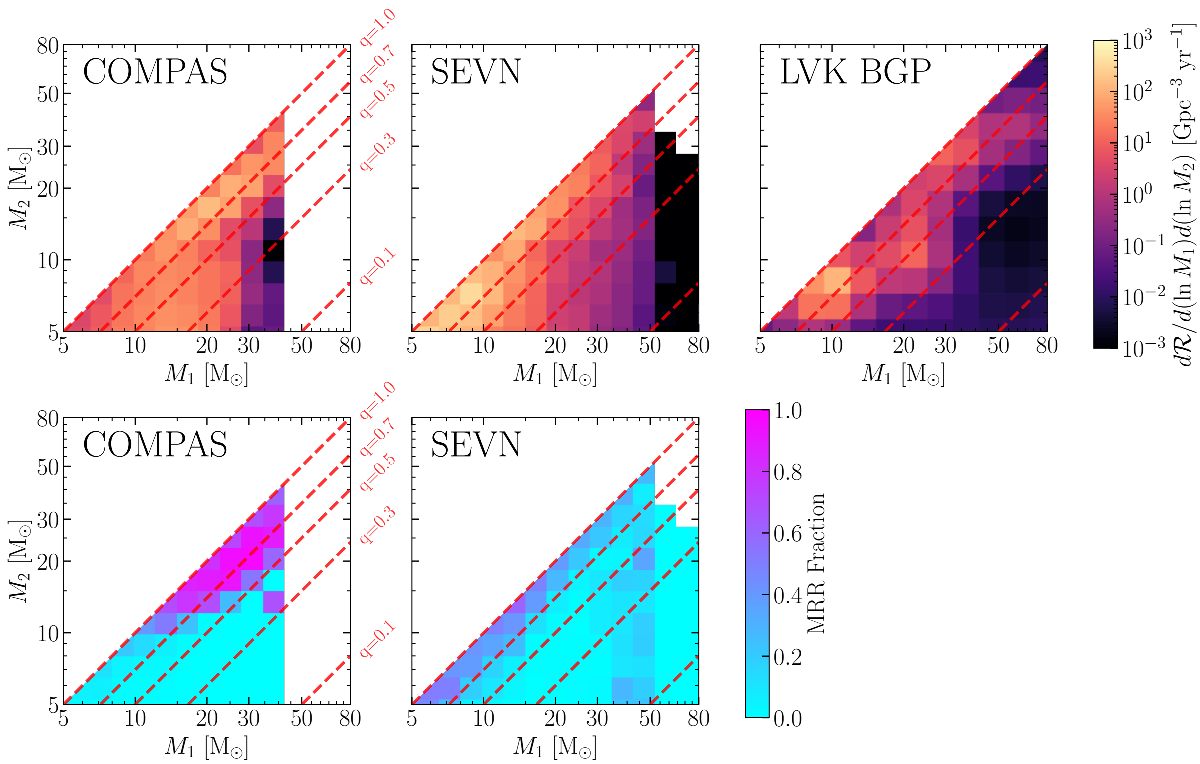

Figure 3.3: Two-dimensional BBH merger rate density ($z \approx 0$) in the (M_1, M_2) plane (top row) and corresponding MRR fraction (bottom row) for COMPAS (left) and SEVN (middle), compared to LVK inference (right). Contours indicate lines of constant mass ratio q. COMPAS exhibits a concentrated region at high masses and $q \gtrsim 0.6$ where MRR systems dominate, while SEVN shows no comparable region of dominance and instead displays a more diffuse MRR contribution across parameter space. This contrast highlights that MRR can either restructure the high-mass BBH population (COMPAS) or remain a subdominant component (SEVN), despite similar overall MRR fractions.

In contrast, the fiducial SEVN model predicts that MRR systems remain subdominant across the full mass range and do not strongly shape the component-mass distributions. Instead, the MRR contribution peaks at lower masses ($M \lesssim 20\,\mathrm{M}_\odot$) and is more diffusely distributed. In the (M_1, M_2) plane, the MRR fraction is largest in the low-mass, high-q region, without a clear regime where MRR dominates the overall population, aside from a preference for $q \gtrsim 0.7$. Combined, this demonstrates that, in `SEVN`, MRR contributes to but does not restructure the observable BBH mass distribution. These results demonstrate that MRR either restructures the high-mass BBH population (`COMPAS`) or remains a secondary contribution (`SEVN`), highlighting a fundamental model-dependent uncertainty in linking massive BBHs to their formation channels.

Despite these differences, both models predict that MRR systems preferentially populate the high mass-ratio regime. In `COMPAS`, MRR systems dominate for $q \gtrsim 0.6$, while in `SEVN` they become most prominent at even higher mass ratios ($q \gtrsim 0.7$), though never dominating. This reflects the tendency for MRR to arise in binaries that evolve toward near-equal-mass configurations through stable mass transfer (see Section 3.4.2 for more details). As a result, the high-q regime could provide a model-independent signature of MRR formed systems, even when its impact on the mass distributions differs substantially between population synthesis frameworks.

Taken together, these results show that while the overall fraction of MRR systems is similar between models, their observational imprint depends sensitively on how MRR systems populate the BBH mass distribution. This is reflected in the comparison to LVK-inferred mass distributions: `COMPAS` overpredicts intermediate-mass systems, while `SEVN` exhibits an excess at low masses. These discrepancies highlight that the imprint of MRR on the BBH population is intertwined with broader modeling uncertainties in binary evolution.

Impact of uncertainties in massive binary evolution in COMPAS

To assess the robustness of the trends identified above, we examine how uncertainties in key binary-evolution prescriptions affect the simulated BBH primary-mass distribution. We focus on single-parameter variations within `COMPAS`, where extensive public model grids enable a controlled exploration of physical assumptions (see Appendix 3B for a complementary analysis of the chirp-mass distribution). The explored variations span physically motivated ranges relevant for BBH formation (Broekgaarden et al., 2022), including the common-envelope efficiency $\alpha_{\rm CE}$, supernova engine prescription, black-hole natal kick dispersion $\sigma_{\rm BH}$, mass-transfer efficiency β, and Wolf–Rayet wind strength $f_{\rm WR}$.

Figure 3.4 shows that these parameters affect the primary-mass distribution in qualitatively different ways, reflecting the distinct physical processes they regulate. Variations in $\alpha_{\rm CE}$, $\sigma_{\rm BH}$, and the supernova engine primarily impact the low-mass end of the distribution ($M_1 \lesssim$ 15–20 $\mathrm{M}_\odot$), below the regime in which MRR systems dominate. These parameters regulate binary survival and compact-object formation pathways, altering the normalization and shape of the low-mass spectrum while leaving the high-mass (MRR-dominated) regime largely unchanged. An exception occurs for extreme common-envelope efficiencies (e.g. $\alpha_{\rm CE} = 10$), where the non-MRR contribution is significantly reduced while the MRR fraction remains comparable. Higher values of $\alpha_{\rm CE}$ lead to larger orbital separations, thus systems which already had wide separations end up not being able to merge in a Hubble time; considering MRR systems require close separations to initiate mass transfer, they may take longer to merge overall but are largely unaffected in this model.

In contrast, variations in the mass-transfer efficiency β and Wolf–Rayet wind strength $f_{\rm WR}$ directly affect the growth and retention of the secondary's mass and helium core prior to collapse. As a result, these parameters modify the MRR contribution across the full mass range: lower mass-transfer efficiency or stronger winds suppress MRR, while more efficient mass transfer enhances it. This highlights the central role of envelope stripping and core growth in setting the prevalence of MRR.

Despite these quantitative differences, all explored variations consistently show that MRR systems preferentially populate the high-mass regime. This indicates that the emergence

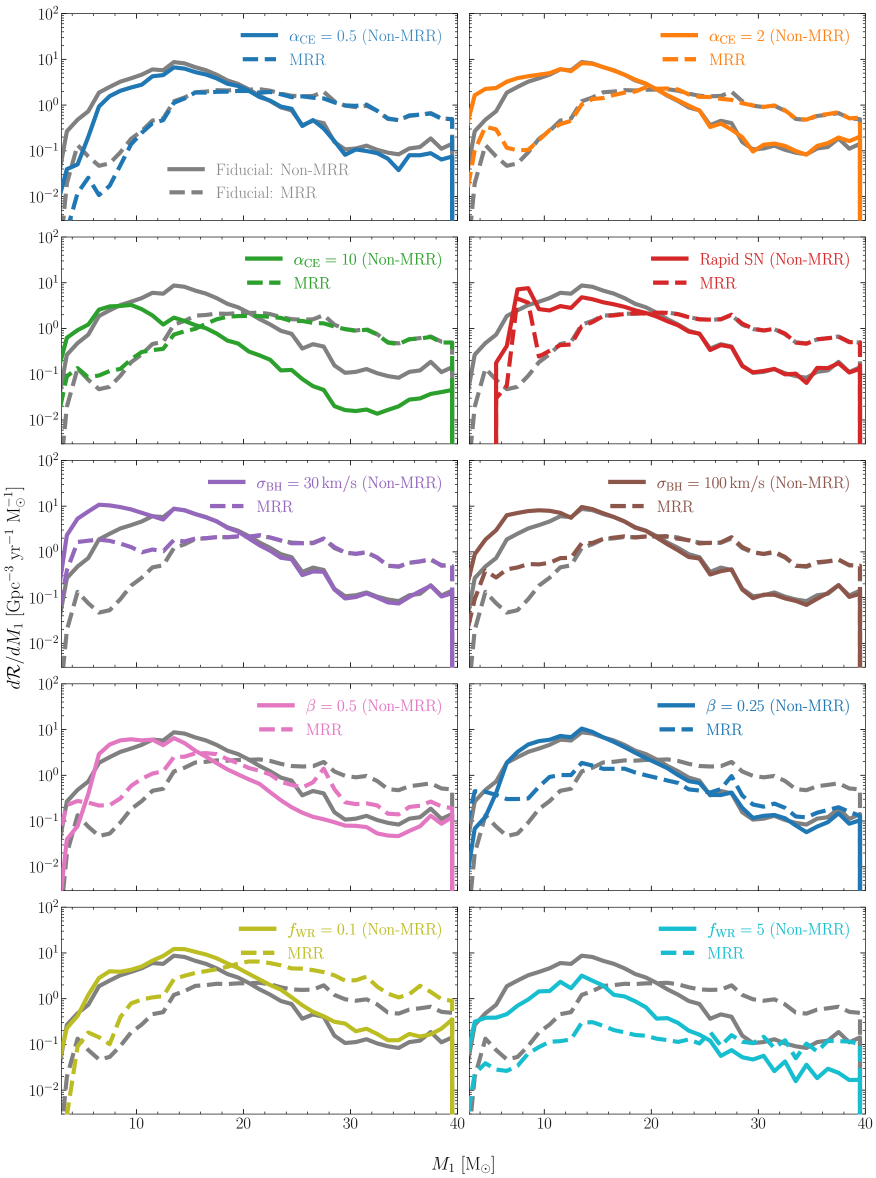

Figure 3.4: Sensitivity of the primary black hole mass distribution, M_1, to key uncertainties in massive binary evolution within `COMPAS`. Each panel shows a single-parameter variation relative to the fiducial model, with the M_1 distribution decomposed into MRR (dashed) and non-MRR (solid) contributions. Variations in the common-envelope efficiency α_{CE}, supernova engine prescription, and black-hole natal kick dispersion σ_{BH} primarily affect the low-mass regime ($M_1 \lesssim 15$–$20\,\mathrm{M}_\odot$) by altering binary survival and compact-object formation pathways. In contrast, changes in the mass-transfer efficiency β and Wolf–Rayet wind strength f_{WR} directly impact core growth and modify the MRR contribution across the full mass range. Despite these differences, all models consistently show that MRR systems preferentially populate the high-mass regime, demonstrating that this is a robust feature of the simulations.

of MRR at high primary masses is not driven by a specific parameter choice, but instead represents a robust feature of isolated binary evolution in `COMPAS`.

3.4.2 What it takes to MASSquerade

The results of Section 3.4.1 show that MRR can either dominate the high-mass BBH population (`COMPAS`) or remain a subdominant contribution (`SEVN`), despite contributing a similar overall fraction of systems in both models. This raises a key question: which binary systems undergo mass ratio reversal, and what physical processes determine where they appear in the BBH mass distribution? To address this, we connect the observable imprint of MRR to its progenitor properties by examining both the initial (ZAMS) parameter space and the evolutionary pathways that lead to MRR systems. We first map where MRR systems originate in ZAMS mass–metallicity space and how the less massive star can map to the more massive BH, or *massquerade* as the primary BH in observed space. We then identify the dominant formation channels that produce these systems.

Origin of MRR systems in initial (ZAMS) parameter space

Figure 3.5 shows the MRR fraction as a function of component masses and metallicity, for both zero-age main sequence (ZAMS) binaries and their resulting black hole masses. Here $\mathrm{M}_{1,\mathrm{BH}}$ is defined as the more massive black hole at merger, consistent with LVK's convention. A complementary version of this figure is presented in Appendix 3F with the direct mapping between the ZAMS masses and their corresponding BH masses. The fraction is weighted by the intrinsic BBH merger rate to reflect the progenitor properties of the underlying BBH merger population at $z \approx 0.2$. In this subsection, we focus on the intrinsic formation pathways of MRR systems within the simulations.

In the fiducial `COMPAS` model, MRR systems are concentrated primarily in binaries with $50\,\mathrm{M}_\odot \lesssim \mathrm{M}_{1,\mathrm{ZAMS}} \lesssim 90\,\mathrm{M}_\odot$ and metallicities $\log_{10}(Z/Z_\odot) \simeq -2.5$ to -1, with a secondary

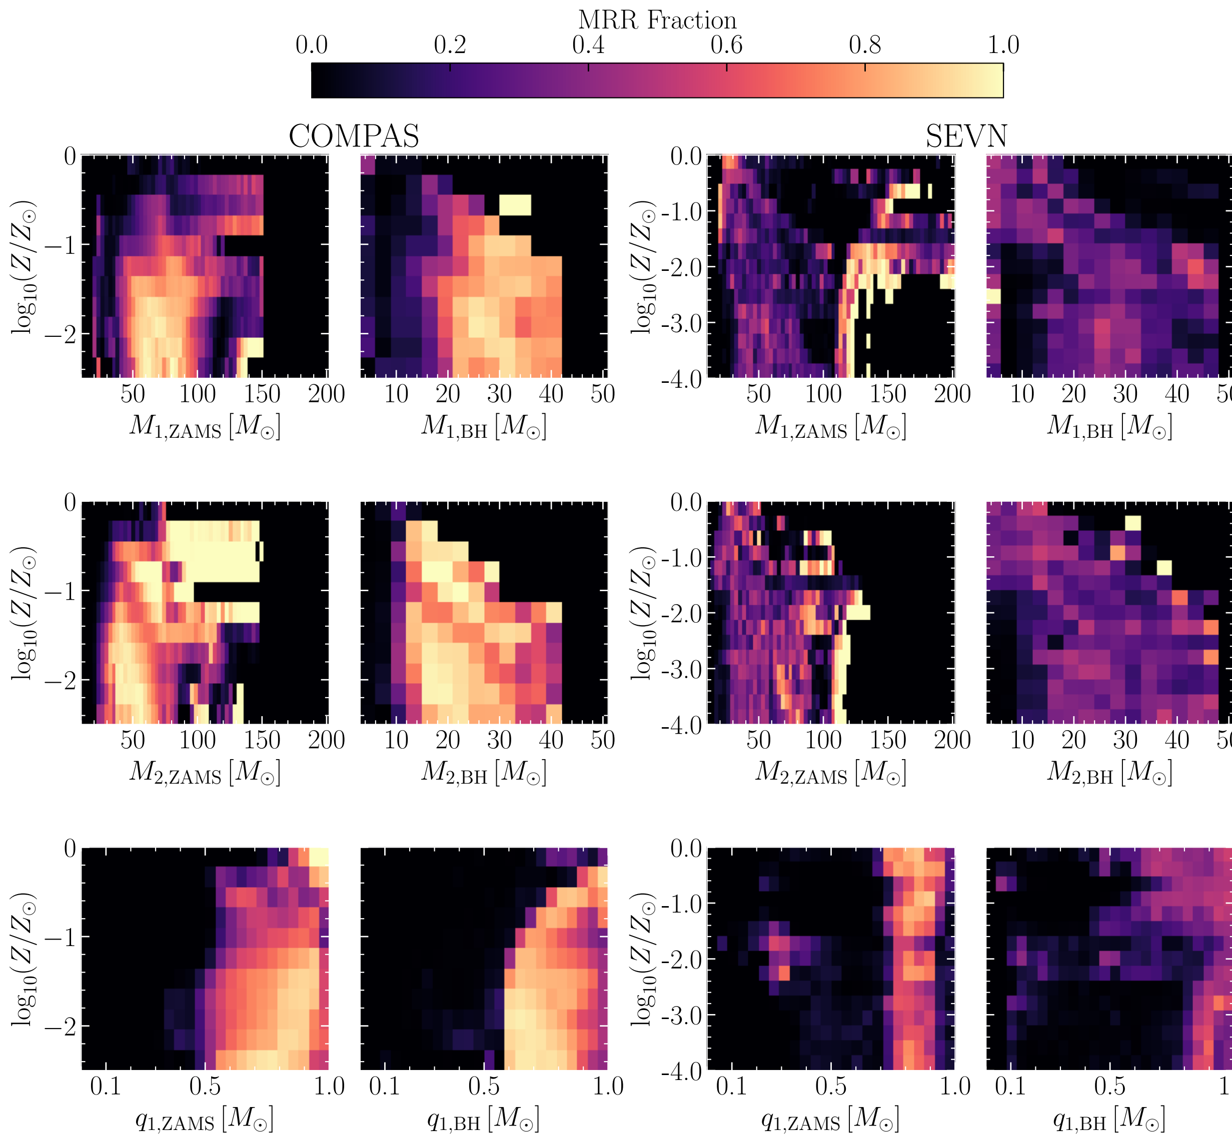

Figure 3.5: Fraction of systems undergoing mass ratio reversal (MRR) for `COMPAS` (left two columns) and `SEVN` (right two columns), weighted by their merger rate contribution at $z \approx 0.2$. For each population synthesis code, the leftmost column shows the MRR fraction in the plane of ZAMS quantities and metallicity, while the right panels show the corresponding BH masses in the "LVK" defined view, i.e. M_1 defined as the most massive BH and M_2 defined as the least massive. See Figure 3.16 for the mapping in the "binary" defined view, i.e. M_1 is the most massive star at ZAMS.

peak around $M_{1,\mathrm{ZAMS}} \sim 130\,\mathrm{M}_\odot$ at lower metallicities $\log_{10}(Z/Z_\odot) \lesssim -1.6$. These systems span a broad range of secondary masses, $M_{2,\mathrm{ZAMS}} \sim 30$–$140\,\mathrm{M}_\odot$, but preferentially originate from binaries with initial mass ratios $q_{\mathrm{ZAMS}} \gtrsim 0.5$, reflecting that MRR is favored in systems that begin with comparable stellar masses.

A mild trend toward higher initial mass ratios at higher metallicity is visible in Figure 3.5. This reflects the impact of stronger stellar winds at higher metallicity, which widen the binary orbit prior to interaction. As a result, successful MRR formation requires more efficient orbital tightening during subsequent binary interactions. This is preferentially achieved in systems with more massive secondaries and higher mass ratios, which provide more tightly bound envelopes during common-envelope evolution and therefore enable sufficient orbital shrinkage.

The MRR binaries predominantly produce black holes in the mass range 20–40 $\mathrm{M}_{\odot}$, consistent with the regime where MRR dominates the BBH mass distribution (Section 3.4.1). Although higher metallicity systems tend to favor more equal-mass binaries, their overall contribution to the BBH population remains small due to the suppressed BBH formation efficiency at $\log_{10}(Z/Z_{\odot}) \gtrsim -0.5$.

In contrast, `SEVN` predicts that MRR arises from two distinct regions of ZAMS parameter space. A high-mass population is concentrated at $M_{1,\mathrm{ZAMS}} \gtrsim 100\,\mathrm{M}_{\odot}$ across a broad metallicity range $-4 \lesssim \log_{10}(Z/Z_{\odot}) \lesssim -0.4$, corresponding to systems that approach or enter the pair-instability regime, where strong pulsational mass loss significantly alters the final remnant mass (see Section 3.4.2). We note that at these high masses, there are less overall systems (IMF suppression) and a higher MRR fraction is thus easier to achieve. A second population of MRR systems is present at lower primary masses and preferentially higher initial mass ratios, similar to systems found in `COMPAS`.

However, unlike in `COMPAS`, these two populations map differently to remnant masses. The high-mass systems experience substantial mass loss due to pulsational pair-instability and have lower overall contribution due to IMF suppression, while the lower-mass, high-q systems produce BHs over a broad mass range. As a result, the initially distinct regions in ZAMS space disperse in BH mass, such that no single BH-mass range is dominated by MRR systems. This explains why MRR does not restructure the BBH mass distribution in `SEVN`, despite contributing a non-negligible fraction of mergers.

Formation channels of MRR systems

The trends and structures in initial and remnant parameter space indicate that MRR arises from multiple evolutionary pathways and a range of initial conditions rather than a single formation channel. We identify three distinct channels through which the initially less massive star can ultimately form the more massive black hole. We include the temporal evolution of the total and core masses for the evolutionary channels in Figure 3.6. These channels occupy distinct regions of initial parameter space, as shown in Figure 3.7.

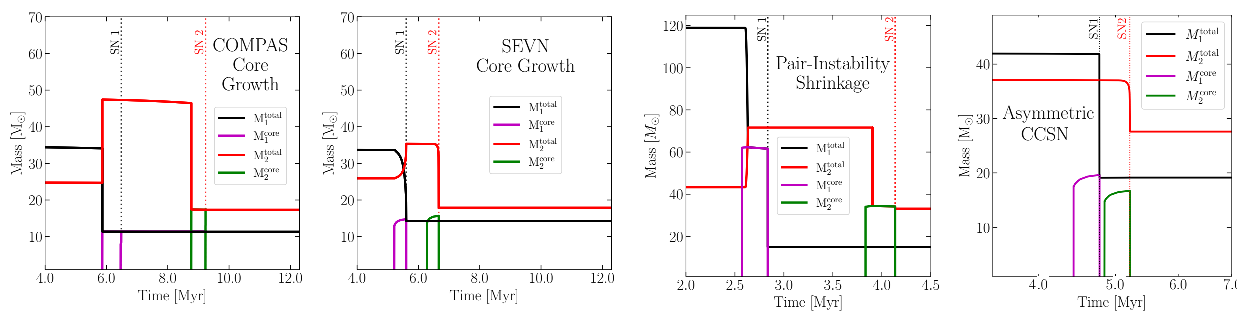

Figure 3.6: Representative evolutionary pathways leading to MRR in binary systems. Shown are the total stellar masses of the initially more massive star (black) and initially less massive star (red), together with their corresponding core masses (magenta and green), as a function of time. Vertical dotted lines mark the supernova events of the primary (SN1) and secondary (SN2). **First:** Core-growth channel (`COMPAS`). Stable mass transfer during post–main-sequence evolution increases the helium-core mass of the secondary, leading to a more massive black hole despite its initially lower ZAMS mass. The reversal here is driven by secular core growth rather than impulsive mass loss. **Second:** Core-growth channel (`SEVN`). The same as the first panel, with a similar system, but evolved in `SEVN`. We note that in `SEVN` the core growth is a gradual process with the primary losing less mass overall. This highlights the differences in stellar evolution mentioned in the text. **Third:** PPISN-shrinkage channel. Strong pulsational pair-instability episodes strip the primary near the PPISN–PISN boundary, reducing its final remnant mass below that of the secondary. **Fourth:** Asymmetric-CCSN channel. The primary becomes stripped while the secondary retains its Hydrogen envelope, leading to the primary star becoming the less massive black hole.

Core growth via stable mass transfer channel The dominant pathway in both `COMPAS` and `SEVN` is a channel we identify as *core growth via stable mass transfer*. In this channel, shown in Figure 3.6, the initially more massive star transfers mass after leaving the main sequence, while the secondary accretes (part of) the transferred material while still on the main sequence. This accretion phase allows the secondary to significantly grow its core mass, such that by the time both stars have undergone core collapse, the initially less massive star can form the more massive black hole. The mass-ratio reversal is therefore primarily driven by the core growth during stable mass transfer. While similar initial parameters were chosen for the systems in Figure 3.6, it is clear that the two population synthesis codes treat stellar evolution differently. In `SEVN` the cores are grown gradually with the secondary growing to a similar size as in `COMPAS`, while the primary loses less overall mass leading to a higher BH mass ratio. Overall, when weighted by the intrinsic BBH merger rate, this pathway accounts for $\sim$ 99% of all MRR systems in both codes, reflecting its dominant contribution to the MRD at $z \approx 0.2$ shown in Figure 3.7.

Pair Instability shrinkage channel A second channel opens in `SEVN`, in which the primary undergoes strong pulsations near the PPISN-PISN boundary, removing a substantial fraction of its mass prior to collapse and reducing its final remnant mass below that of the secondary (see the third panel of Figure 3.6). The secondary, which avoids this strong pulsational mass-loss phase, experiences weaker PPISN stripping and therefore retains a larger final remnant mass, ultimately forming the more massive black hole. We refer to this pathway as *PPISN-induced shrinking*. Although this channel contributes less than 1% of the MRD-weighted MRR systems, it dominates the MRR contribution at high $M_{1,\mathrm{ZAMS}}$ and low mass ratios. In contrast to the core growth channel, the mass ratio reversal here is driven primarily by asymmetric mass loss at late evolutionary stages.

Asymmetric CCSN channel The final pathway identified in `SEVN` is driven by asymmetric mass loss during the CCSN phase and contributes 1.5% of the MRD-weighted MRR systems. This channel is distinct from the core growth channel, where the reversal is encoded in the pre-SN core mass hierarchy, and from the PPISN shrinkage channel, where the reversal is driven by enhanced pulsational stripping of the primary's mass. In the asymmetric-CCSN channel, some primary stars experience envelope stripping prior to collapse, such that the primary is stripped while the secondary retains its envelope, leading to substantially different CCSN mass loss between the two stars (see right panel of Figure 3.6). However, the channel also includes systems in which both stars are stripped, but the non-linear mapping between pre-supernova structure and remnant mass causes the secondary to lose relatively less mass during collapse and ultimately form the more massive remnant. Overall, this channel remains subdominant across the full initial parameter space.

Pulsational-pair instability assisted core-growth channel The small remaining fraction of systems in `COMPAS` follow an evolutionary pathway similar to the core-growth channel described above, but with one important difference: the secondary's core mass never exceeds that of the primary. Instead, stable mass transfer drives the system toward nearly equal core masses, placing both stars in the pulsational-pair instability supernova (PPISN) regime. Because the primary still retains a slightly larger core, it loses marginally stronger PPISN mass loss than the secondary, ultimately producing the less massive black hole. We refer to this pathway as *PPISN-assisted core-growth*. Because the evolution is a mix between core-growth and PPISN-shrinking (with a much smaller mass-loss differential), we do not include a plot of its temporal evolution. This channel contributes only a small fraction ($\lesssim 1\%$) of MRR systems and is found primarily at the lowest mass ratios.

Throughout both codes, metallicity plays a secondary role with core-growth dominating entirely. The contributions of the other channels are minimal: PPISN-assisted core-growth truncates above $Z \gtrsim -1.3$ in `COMPAS`, a similar truncation occurs for PPISN-shrinking in `SEVN`, and asymmetric-CCSN follows the core-growth shape in `SEVN` but contributes minimally.

The dependence of MRR on orbital separation for both `COMPAS` and `SEVN` (including the

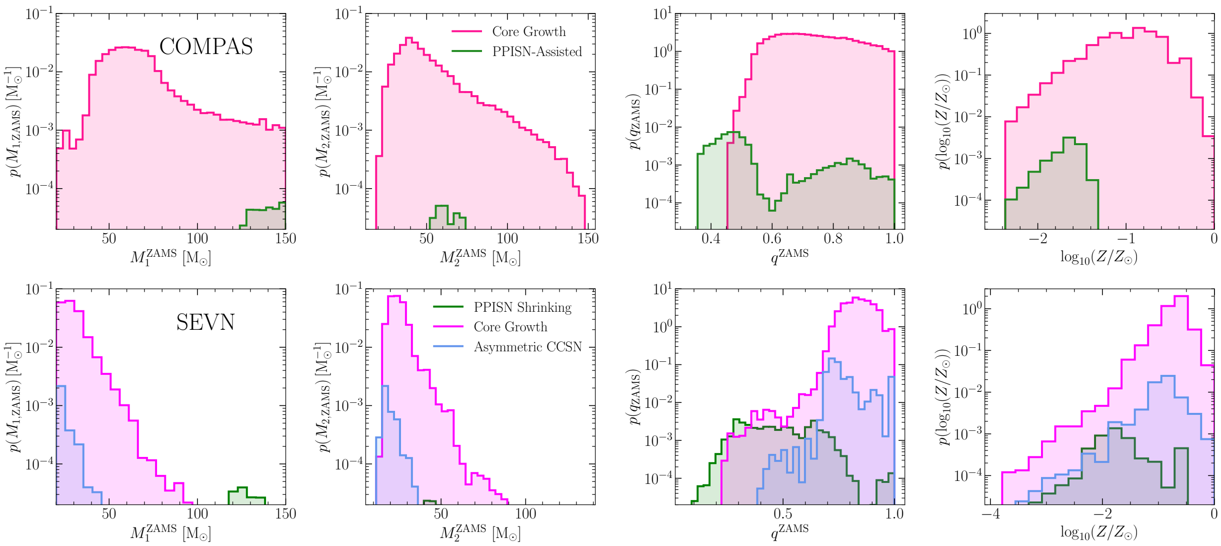

Figure 3.7: Initial parameter distributions of the distinct MRR channels in COMPAS (top-row) and SEVN (bottom-row). The distributions are normalized jointly across each codes channels, so that the area under the sum of all curves equals unity per panel. From left to right the columns are: ZAMS primary mass, ZAMS secondary mass, initial mass ratio, and metallicity. The channels for COMPAS are core-growth (pink) and PPISN-assisted core-growth (green), while for SEVN they are the PPISN-shrink (green), core-growth (magenta), and asymmetric-CCSN (blue). The core-growth channel accounts for 99% of MRR systems dominating throughout the initial parameter space for both codes. In COMPAS the PPISN-assisted core-growth dominates only at low-mass ratios. While in SEVN, a small dominant contribution for PPISN-shrink arises at high $M_{1,\rm ZAMS}$ and small mass ratio. The asymmetric-CCSN channel in SEVN remains subdominant, peaking at high component masses and mass ratio. Appendix 3C shows how the relative contributions of these channels change when weighted by the formation efficiency (Eq. 3.2) instead of the merger rate density.

evolutionary channels) is included in Appendix 3D.

Taken together, these results demonstrate that the presence and observational imprint of MRR systems are set by a combination of initial mass, mass ratio, and the efficiency of mass transfer and pair-instability processes, providing a physical explanation for the contrasting MRR populations predicted by COMPAS and SEVN.

3.5 Discussion

3.5.1 Massquerading in LVK: when a secondary masquerades as a primary

A central implication of our results is that mass ratio reversal (MRR) fundamentally complicates the interpretation of observed binary black hole (BBH) populations. In particular, MRR enables black holes formed from the initially less massive star (the secondary) to appear as the more massive component (M_1) in the observed BBH system. We refer to this effect as *massquerading*: the secondary star "masquerades" as the primary in the observed mass distribution.

Considering the massquerading phenomena is important to understand features in the BBH space and map back to formation pathways of their progenitor stars. From the perspective of binary stellar evolution, the initially more massive star (the donor in the first mass-transfer phase) and the initially less massive star (the accretor) follow systematically different evolutionary pathways in the typical formation channel of a BBH source (van Son et al., 2023b). The donor evolves first, loses mass through winds and mass transfer, and forms the first compact object, whereas the accretor can gain mass, angular momentum, and build a more massive helium core prior to collapse. As a result, the two stars leave distinct imprints on the resulting black hole population, a phenomenon well studied in binary evolution studies (e.g. Kippenhahn & Weigert 1967; Wellstein et al. 2001; de Mink et al. 2013; Renzo et al. 2019, 2023; Renzo & Götberg 2021; Bavera et al. 2021; Wagg et al. 2024; Sen et al. 2026; Zapartas et al. 2021, 2026).

This distinction is reflected in our simulations through the different mappings from ZAMS mass to black hole mass for primaries and secondaries (Figure 3.5; see also Appendix 3F). Secondaries that undergo efficient accretion can develop disproportionately large cores and form black holes that exceed the mass of the primary remnant, despite their initially lower mass.

However, gravitational-wave observations do not retain information about the progenitor identities of the two compact objects. Instead, the LVK collaboration defines M_1 and M_2 purely by rank ordering ($M_1 \geq M_2$), independent of their evolutionary origin. This creates a fundamental mismatch between theoretically meaningful labels (donor vs. accretor) and observational labels (more massive vs. less massive black hole) (e.g. Broekgaarden et al., 2022; Zevin & Bavera, 2022; Mould et al., 2022).

Figure 3.8 illustrates this mismatch directly. When black holes are classified according to their mass ordering at ZAMS ("binary" definition), the mass distribution separates into contributions from remnants of initially more massive stars and those formed from initially less massive stars. These two populations occupy partially distinct regions of parameter space, reflecting their different evolutionary histories. However, when the same systems are re-labeled according to the observational convention ("LVK" definition), the two populations become mixed: black holes originating from secondaries are reassigned into the M_1 distribution whenever they exceed the mass of the primary remnant. See Figure 3.15 in Appendix 3E for the effects of single parameter variations in `COMPAS` on massquerading.

As a result, the observed primary-mass distribution dR/dM_1 is not a direct tracer of the evolution of initially more massive stars, but instead represents a superposition of two physically distinct populations. In particular, features produced by accretor evolution—such as enhanced core growth through stable mass transfer—can appear as excess structure in the observed M_1 distribution. This effect is clearly visible in Figure 3.8, where the "LVK" M_1 distribution shows contributions at intermediate masses that originate from the secondary population in the binary-defined view.

This massquerading effect provides a natural explanation for several discrepancies between population-synthesis predictions and LVK-inferred mass distributions discussed in Section 3.4.1.

In particular, in `COMPAS`, the excess of systems in the ~ 10–$30\,\mathrm{M}_\odot$ range arises in part from MRR systems in which the accretor forms the more massive black hole and is therefore counted as M_1 observationally. Conversely, the donor remnants are shifted into the M_2 distribution, altering both component-mass spectra simultaneously.

More broadly, this implies that the commonly used interpretation of M_1 as tracing the evolution of the initially more massive star is not valid in the presence of MRR. Instead, M_1 must be interpreted as a mixed population whose composition depends on the efficiency of mass transfer, wind mass loss, and core growth. Neglecting this mixing can therefore bias inferences on key aspects of massive binary evolution, including the strength of stellar winds, the efficiency of envelope stripping, and the mapping from stellar core mass to compact remnant mass or understanding which of the BHs is spinning (e.g. Stevenson et al. 2017; Mandel & Broekgaarden 2022; Zevin et al. 2021).

The strength of the massquerading effect is model-dependent. In `COMPAS`, where MRR systems dominate the high-mass regime, secondaries contribute substantially to the observed M_1 distribution, leading to a pronounced reshaping of the BBH mass spectrum. In contrast, in `SEVN`, where MRR systems are more diffusely distributed in remnant mass due to the combined effects of core growth and pulsational pair-instability mass loss, the impact of massquerading is less localized but still broadens the observed distributions.

Taken together, these results highlight that interpreting gravitational-wave observations requires careful consideration of the mapping between progenitor properties and observed quantities. In particular, forward-modeling approaches that retain information about progenitor identity, or inference frameworks that explicitly marginalize over formation channels, will be essential to disentangle the contributions of donor and accretor evolution in the observed BBH population.

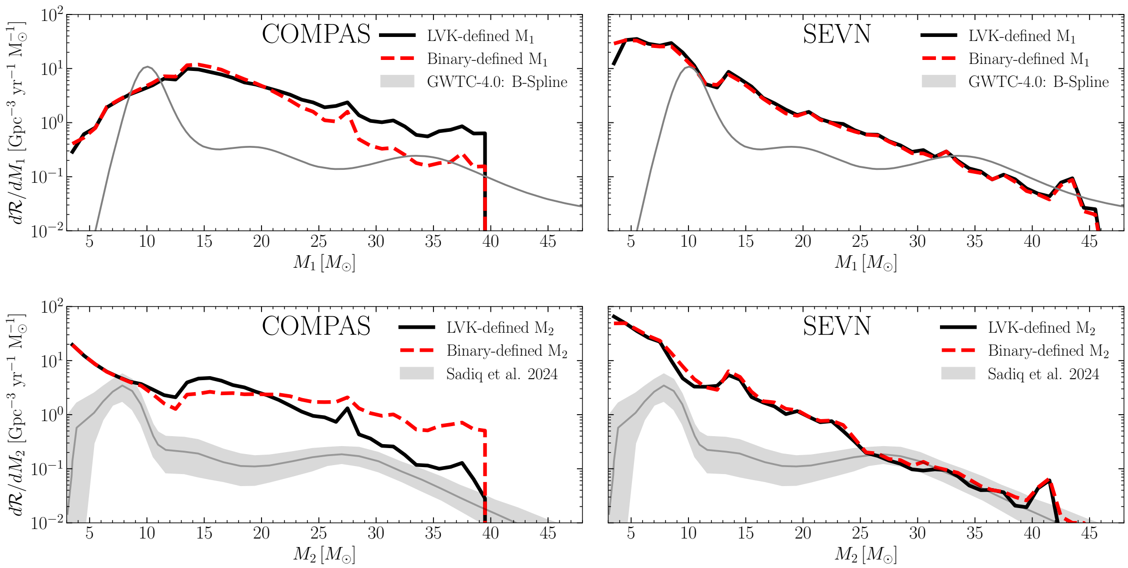

Figure 3.8: Massquerading in the component-mass distributions for `COMPAS` (left) and `SEVN` (right). The results from `COMPAS` show that depending on whether the primary black hole is classified based on being the larger black hole or stemming from the initially larger progenitor can have significant effects on the high-end regime in the mass distributions. Conversely, this definition has little effect on the results from `SEVN` because MRR never dominates a specific region in the mass distributions. The gray bands correspond to LVK-inferred constraints (Abac et al., 2025; Sadiq et al., 2024).

3.5.2 MRR fraction across cosmic time

Figure 3.9 shows the fraction of binary black hole mergers undergoing mass ratio reversal as a function of redshift. The fiducial `COMPAS` model is shown in black, while colored curves correspond to variations in key binary-evolution prescriptions, including the common-envelope efficiency ($\alpha_{\rm CE}$), black-hole natal kick dispersion ($\sigma_{\rm BH}$), the mass-transfer efficiency parameter (β), and the Wolf–Rayet wind mass-loss strength ($f_{\rm WR}$). Across all models, the MRR fraction decreases with increasing redshift. Early stars are metal poor, as such they lose less mass to stellar winds and ultimately end their lives as more massive black holes than their metal-rich counterparts. In order to undergo MRR, the less massive star needs to overcome the mass of the primary, which becomes increasingly difficult at early times.

Notably, the $\alpha_{\rm CE} = 10$ variation has the highest MRR fraction throughout cosmic time. This can be understood in conjunction with Figure 3.4, where the MRR contribution to

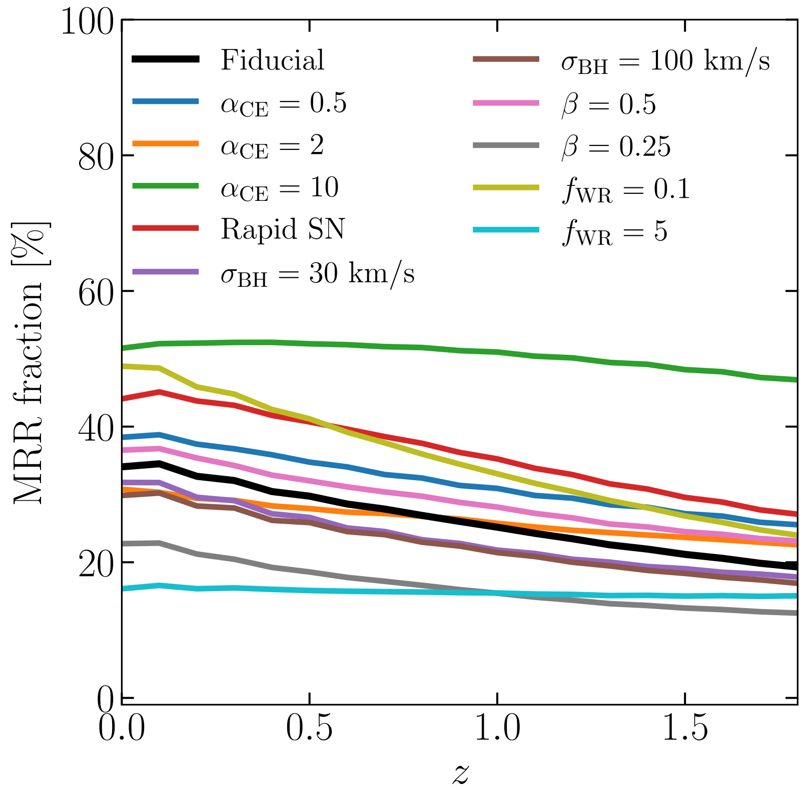

Figure 3.9: The mass ratio reversal contribution to the merger rate density across cosmic time. The black line is the fiducial `COMPAS` model, while the colored lines correspond to single parameter variations. Throughout cosmic evolution all variations show a slight decline in MRR fraction, hinting at a dependence on metallicity and/or delay-time.

the primary BBH mass distribution is largely unchanged, but the non-MRR contribution is greatly reduced leading to an overall increase in the total MRR fraction. For other values of $\alpha_{\rm CE}$ the MRR fraction doesn't deviate largely from the fiducial model.

The overall MRR fraction is particularly sensitive to the Wolf–Rayet wind mass-loss strength. In fact, the lowest MRR fraction across cosmic time is seen for the model with $f_{\rm WR} = 5$. This is also clear from Figure 3.4, where the MRR contribution is largely reduced compared to the non-MRR contribution. The fiducial model, with $f_{\rm WR} = 1$, doubles the MRR fraction at early times, and further lowering this parameter to $f_{\rm WR} = 0.1$ increases the MRR fraction another $\sim 70\%$ at early times. At early times, $z \gtrsim 1.0$, these discrepancies lower to less than a factor of 2 following the metallicity dependence of WR formation and mass loss rates.

The remaining model variations have modest departures from the fiducial model in terms of MRR fraction. The rapid supernova engine increases the overall MRR fraction $\sim 30\%$ over the redshift range shown. The BH natal kick dispersion $\sigma_{\rm BH}$ affects the MRR fraction at the sub $\sim 15\%$ level, increasing the non-MRR and MRR fraction in tandem as can be seen in Figure 3.4. Finally, the shift to a constant mass-transfer efficiency of $\beta = 0.5$ slightly increases the MRR fraction, while lowering it to $\beta = 0.25$ suppresses the MRR fraction substantially. Overall, the MRR fraction's evolution consistently increases as time evolves and is mainly affected by winds, envelope stripping, and mass transfer.

3.5.3 Comparison to earlier work

Our results on MRR in BBH systems are broadly consistent with previous studies detailed here. Much of the existing literature has focused on the connection between MRR and BH spins, whereas here we emphasize its imprint on the mass distribution of BBHs.

Using the same set of `COMPAS` simulations adopted in this work, Broekgaarden et al. (2022) found that MRR occurs in $\gtrsim 70\%$ of BBHs observable by LVK and between 11–82% of the astrophysical population, depending on model assumptions. They showed that MRR occurs for masses $\gtrsim 20\,\mathrm{M}_\odot$ and mass ratios of ~ 0.7 in agreement with our findings. They further demonstrated that the second-born BH can acquire non-negligible spin in up to 25% of BBH systems and that the MRR systems can account for the anti-correlation in $\chi_{\rm eff}$–q reinforced by later studies, e.g. Banerjee & Olejak (2024).

Similar conclusions were reached by Zevin & Bavera (2022), who emphasized the role of MRR in producing asymmetric-mass BBH systems with a spinning primary BH. They found that the accretion efficiency during mass transfer is the dominant factor determining whether a system undergoes MRR. This is in direct agreement with our results (Section 3.4.1), where we find that higher mass transfer efficiency leads to a significantly larger fraction of MRR systems.

Consistent with this picture, Olejak & Belczynski (2021); Olejak et al. (2024) showed that stable mass transfer naturally leads to MRR in BBH systems, typically producing mass ratios in the range $q \sim 0.4$–0.7. These systems can also give rise to primary black holes with non-negligible spins, driven by accretion and tidal spin-up during binary evolution, further reinforcing the connection between MRR and observable spin signatures.

Recently, on the observational side, Sen et al. (2026) use samples of Algol-type binaries, i.e. systems having underwent MRR, to constrain mass transfer efficiency and angular momentum loss using an analytic framework. Their results provide empirical support for the range of mass transfer efficiencies required for MRR to occur, an observational complement to our work and the other population synthesis studies mentioned above.

Our analysis focuses exclusively on the BBH population, however MRR has been noted in the context of NS–BH binaries (Sipior et al., 2004), though we note that Broekgaarden et al. (2021) find the occurrence rate to be in less than 1% of such systems. MRR has also been shown to be a generic consequence of binary evolution for neutron star–white dwarf binaries (Tauris & Sennels, 2000; Toonen et al., 2018). Similarly, MRR may play a role in shaping the NS–NS binary populations. A systematic study of MRR across all compact object binary classes would be a valuable direction for future work.

3.5.4 COMPAS vs SEVN

An important outcome of our work is the discrepancy between the `COMPAS` and `SEVN` results. These differences largely stem from the contrasting implementation of stellar evolution in each code. `COMPAS` utilizes the parameterized models for stellar and binary evolution from Hurley et al. (2000); Hurley et al. (2002) respectively, in which fitting formulae are employed to estimate stellar evolution. This is in contrast to the methodology employed by `SEVN`, which uses an updated set of tracks and interpolates between neighboring tracks *on-the-fly* to evolve stars.

While the fiducial models in this chapter utilize different stellar tracks, it has been shown that discrepancies still arise between rapid and interpolated codes even when using identical tracks, as shown by Agrawal et al. (2020, 2023). Their work shows discrepancies arise in the core and remnant masses, lifetime, and radii with their interpolating code (`METISSE`) retaining finer details in reproducing the stellar tracks.

Similarly, Romagnolo et al. (2023) find discrepancies in the maximum radius–mass relationship between the Hurley et al. (2000) (`SSE`) fitting formulae and `METISSE`. This work also reports differences in the total and component mass distributions between codes, becoming most pronounced for component masses of $\gtrsim 25\,\mathrm{M}_{\odot}$, depending on the choice of tracks. The authors explain that this is also due to the fact that uncertainties increase with stellar masses, in particular for the regime $M_{\mathrm{ZAMS}} \gtrsim 50\,\mathrm{M}_{\odot}$, where `SSE`-based models require extrapolations of the Hurley et al. (2000) fitting formulae.

Other explorations have been carried out, for instance by Belczynski et al. (2022) who compared rapid population synthesis codes (`StarTrack` and `COMPAS`) to detailed evolution code `MESA`. In particular, they attempt to match the observational properties of binary system Melnick 34 finding a wide variety of remnant outcomes ranging from a close BBH system to a totally disrupted system, showcasing that evolutionary conclusions from codes must be interpreted with care.

We note that the discrepancies arising from the stellar treatment in the fiducial models (e.g. Figure 3.2) is larger than those arising from single parameter variations (Figure 3.4). More plainly, no matter which `COMPAS` model is employed, there is always a regime in which MRR dominates in the mass distributions. While we did not explore single parameter variations in

SEVN, the fiducial model contains no region where MRR is dominant in the mass distributions.

Within COMPAS, the effects of MRR can lead to the *massquerading* effect in the mass distributions most clearly shown in Figure 3.8. However, the effects of MRR in the SEVN results are hardly noticeable in this figure and are shown to be subdominant across the mass distributions (Figure 3.2). A one-to-one comparison with the same tracks and same set of stellar/binary parameters would be worthwhile to quantify these differences robustly in COMPAS and SEVN.

3.5.5 Uncertainties in high-mass primary ZAMS systems

Stellar evolution represents one of the key uncertainties in the formation of BBH systems. Following for instance Iorio et al. (2023b), the stellar tracks of which are used in our analysis with SEVN, at high metallicity and high initial masses stars do not significantly expand due to the ejection of their hydrogen-rich envelopes from strong winds (this is also highlighted by other recent studies such as Romagnolo et al. 2024; Hirschi et al. 2025 and Romagnolo et al. 2026b), with rotation potentially extending this phenomenon also down to Small Magellanic Cloud metallicities (Boco et al., 2025). Non-expanding stars are not likely to initiate mass transfer events, and even if they do there is a limit on how close they can get from mass transfer (Klencki et al., 2026), which means that, i) such binaries are unlikely to contribute to the population of GW sources since their stellar components may evolve in isolation, and ii) the only available MRR for such systems is through PPISN-shrinking. However, PPISN can only happen for massive cores that are usually produced from the evolution of low-metallicity stars. This suggests that the high-$M_{1,\mathrm{ZAMS}}$ population of PPISN MRR binaries may represent only a small sample of low-metallicity systems, with no contribution from the high-Z and high-$M_{1,\mathrm{ZAMS}}$ binaries (see also Romagnolo et al. 2025). This is however only something that is clearly visible from our SEVN simulations, since codes that were derived from Hurley et al. (2000); Hurley et al. (2002) tend by design to expand considerably even after strong wind-driven mass loss (Agrawal et al., 2020; Romagnolo et al., 2023).

3.6 Conclusions

In this work, we have investigated mass ratio reversal (MRR), the evolutionary outcome in which the initially less massive star in a binary produces the more massive black hole, and assessed its imprint on the BBH population inferred by LVK. By analyzing population synthesis predictions from `COMPAS` and `SEVN` with merger rate modeling, we find that while MRR systems constitute approximately one-third of the local merger rate, they play a disproportionate role in shaping the high-mass and high-mass-ratio regions of the BBH parameter space. Our results are summarized below:

1. Mass ratio reversal comprises $\sim 1/3$ the total merger rate density across redshift evolution, with `COMPAS` exhibiting a slightly decreasing fraction over z while `SEVN` does not (Figure 3.1).

2. The MRR fraction is plotted for the component and mass ratio distributions in Figure 3.2, in which we find that MRR dominates the high mass ($M_1 \gtrsim 20\,\mathrm{M}_\odot$, $M_2 \gtrsim 12\,\mathrm{M}_\odot$) and high mass ratio ($q > 0.6$) regime in `COMPAS`. For `SEVN`, MRR peaks at low masses and high mass ratio ($q > 0.7$), but never dominates the distributions. The MRR mapping is further emphasized in Figure 3.3 with a clear overabundance in the high mass / high mass ratio regime for `COMPAS`, but only a slight peak in MRR fraction near the $q = 1$ boundary in `SEVN`.

3. We analyze 10 single parameter variations from the fiducial `COMPAS` model (Figure 3.4), finding that though there are differences in the overall MRR structure, these systems are found preferentially in the high-mass regime across all models, with 57–87% of MRR systems lying above $15\,\mathrm{M}_\odot$.

4. MRR systems are mapped from their initial parameter space to their LVK defined black hole space, i.e. M_1 being the larger BH, in Figure 3.5 (and to their ZAMS definition, M_1 the initially more massive star in Figure 3.16). We find that in `COMPAS`, MRR fractions are highest for $M_{1,\mathrm{ZAMS}} \sim 50$–$90\,\mathrm{M}_\odot$, $M_{2,\mathrm{ZAMS}} \sim 30$–$140\,\mathrm{M}_\odot$, $\log_{10}(Z/Z_\odot) \simeq -2.5$ to -1, and $q_{\mathrm{ZAMS}} \gtrsim 0.5$. These get mapped to primary BHs with 20–40 $\mathrm{M}_\odot$ and mass ratios of $q > 0.6$. In `SEVN`, MRR fractions are highest for $M_{1,\mathrm{ZAMS}} \gtrsim 100\,\mathrm{M}_\odot$, $M_{2,\mathrm{ZAMS}} \gtrsim 80\,\mathrm{M}_\odot$, $\log_{10}(Z/Z_\odot) \simeq -4.0$ to -0.4, and $q_{\mathrm{ZAMS}} \gtrsim 0.6$. However, for `SEVN` the mapping is dispersed across the BH mass space and no clear region has a dominating MRR fraction consistent with the findings in Figures 3.2 and 3.3.

5. The primary evolutionary channel in which a binary becomes an MRR system is through the core-growth channel (Figure 3.6), where the secondary core is grown larger than the primary star's through stable mass transfer. We find this channel accounts for 99% of MRR systems contributing to the MRD at $z \sim 0.2$. We additionally identify three other MRR pathways that contribute $\lesssim 1\%$ to the MRR BBH merger rate density at $z \sim 0.2$. This includes:

    - `SEVN`: Two subdominant channels are found, the first is PPISN-shrinking in which asymmetric pulsational mass loss leads to the primary star losing a substantial amount of its mass before collapse while the secondary retains a larger fraction of its envelope and becomes the larger BH. The last channel, asymmetric-CCSN, is attributed to uneven mass loss at the CCSN event, sometimes due to the secondary retaining its envelope.

    - `COMPAS`: A secondary channel follows the core-growth route, however the secondary core doesn't surpass the primary and instead assistance is required from PPISN pushing the primary to a slightly lower remnant mass than the secondary's.

6. Finally, we find that the labels given by LVK, i.e. M_1 and M_2 referring to the more

and less massive BH component do not map directly to the ZAMS definitions of donor and accretor. This *massquerading* where the the secondary star "masquerades" as the primary in the observed BH mass distribution can lead to excess structure in the mass distributions which needs to be carefully considered when mapping between observed quantities and progenitor properties.

Taken together, our results demonstrate that MRR is not a rare outcome of stellar evolution, but a structurally important feature of isolated binary evolution, particularly at high masses and mass ratios. These systems must be carefully integrated into inference studies to ensure proper mapping between progenitor properties and observed quantities.

Data Availability

For the `COMPAS` models the authors made use of the binary black hole simulations from (Broekgaarden et al., 2022), which are publicly available at Broekgaarden (2021). The `SEVN` model was based on the simulations from Smith & Kaplinghat (2026), using the same analysis pipeline we used an updated version of `SEVN (v2.13.0)` Iorio et al. (2023b). The `SEVN` data is publicly available along with necessary datasets to recreate all figures in Smith (2026). The observational data from LVK was obtained from the publicly available dataset at LIGO Scientific Collaboration et al. (2025). The data from Sadiq et al. (2024) was digitized using `PlotDigitizer` (PlotDigitizer, 2026).

Software Acknowledgment

This work made use of the following software packages: `astropy` (Astropy Collaboration et al., 2013, 2018, 2022; Robitaille et al., 2023), `Jupyter` (Perez & Granger, 2007; Kluyver et al., 2016), `matplotlib` (Hunter, 2007), `numpy` (Harris et al., 2020), `pandas` (Wes McKinney, 2010; pandas development team, 2026), `python` (Van Rossum & Drake, 2009), `SSPC` (Hendriks et al., 2023), and `scipy` (Virtanen et al., 2020; Gommers et al., 2024).

Software citation information aggregated using `The Software Citation Station` (Wagg & Broekgaarden, 2024; Wagg et al., 2025).

Appendix 3A: Mass ratio criterion

In this work we have used the simple criteria that the mass of the black hole descending from the initially less massive star (B) needs to be larger than the initially more massive star's black hole (A), i.e. $\mathrm{M_B} > \mathrm{M_A}$. However, the criteria for MRR can be shifted to be more conservative, for instance it can be set to $\mathrm{M_B} > 1.05\ \mathrm{M_A}$. However, this can unfairly impact massive stars where the difference can be multiple $\mathrm{M_\odot}$. It is also possible to present the criteria as a minimum mass difference between the two, e.g. $\mathrm{M_B} > \mathrm{M_A} + 1\ \mathrm{M_\odot}$. This on the other hand can be more difficult for smaller mass systems to overcome. Thus, we simply use $\mathrm{M_B} > \mathrm{M_A}$ in this work, but highlight one possible alternative in this appendix.

Comparing the fractional contribution from MRR systems in the total MRD at $z \approx 0.2$, we find that in our current approach `COMPAS` sits at 33%, while `SEVN` at 32%. If we implement a threshold of $\mathrm{M_B} > 1.05 \times \mathrm{M_A}$ the fractional contributions are: `COMPAS` - 29% and `SEVN` - 27%. These are relatively negligible differences, and this level persists across the full redshift evolution.

While, these are small differences in the overall MRD, the differential MRD tells a slightly different story. Figure 3.2 shows the case with no threshold and we include in Figure 3.10 the case for $\mathrm{M_B} > 1.05\ \mathrm{M_A}$. For `COMPAS` the differences are overall minimal, there exists a peak in the primary- and secondary-mass distributions at $\sim 25\ \mathrm{M_\odot}$, the non-MRR contribution dominates at the very end of the secondary distribution ($\sim 36 - 40\,\mathrm{M_\odot}$), and the high-mass ratio regime falls off for MRR systems (as expected due to the cut being placed on mass ratio itself). Similarly, the results from `SEVN` remain consistent with MRR never dominating, however the MRR contribution drops off much steeper when implementing a conservative cut. Overall, the results are largely consistent including a conservative cut, the total MRD remains nearly the same, and the main behavior in dR/dX remains consistent.

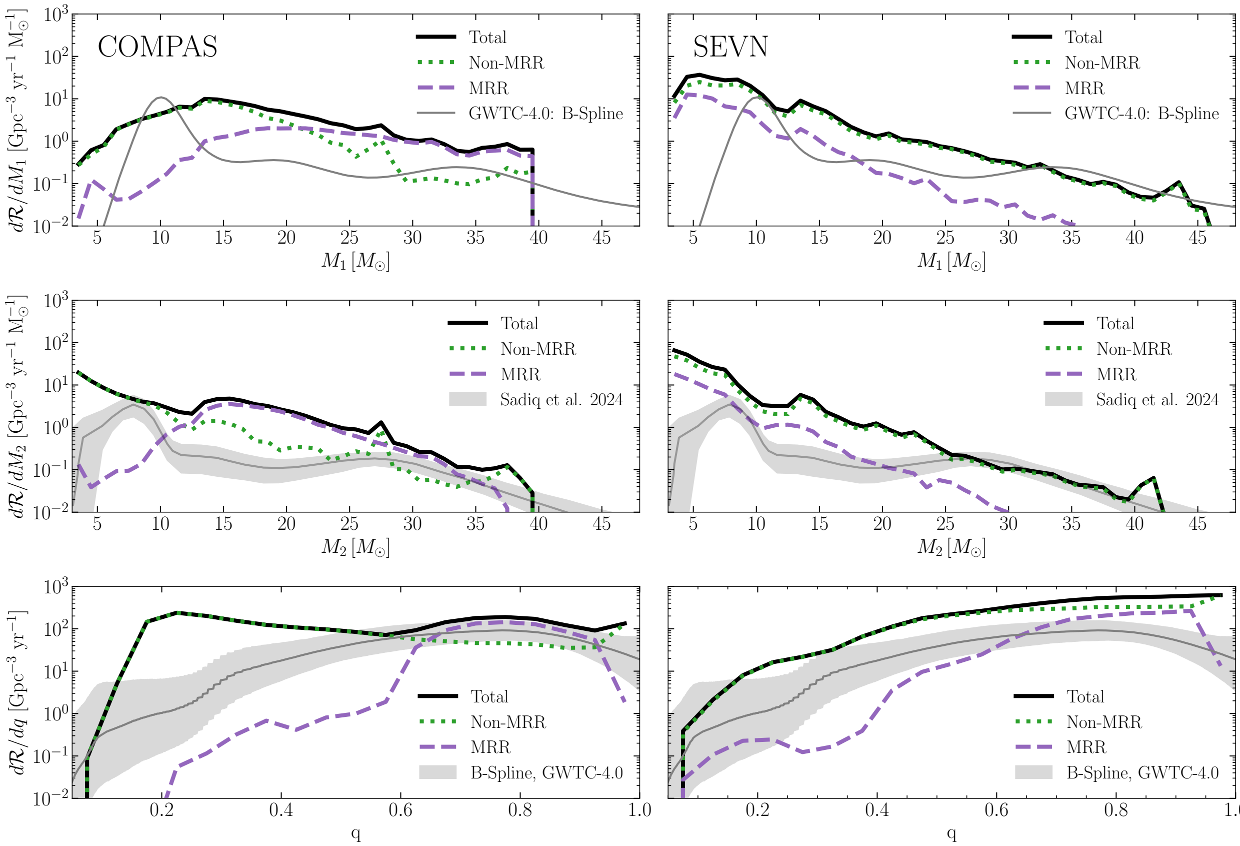

Figure 3.10: The same as Figure 3.2, but with an imposed threshold that the secondary star's remnant needs to be at least 5% more massive than the primary star's remnant (M_B ¿ $1.05 \times M_A$). For `COMPAS` (left-column) the main difference in the primary- and secondary-mass distributions are that the non-MRR contribution increases with a peak between 25–30 $M_\odot$ and a dip at mass ratios near unity (since this is where the cut-off lives. In `SEVN` the same behavior is seen in the mass ratio distribution, and in the primary- and secondary-mass distributions we note that there is a steeper drop off from the MRR contribution tending towards higher masses.

Appendix 3B: Variations in the chirp mass distribution

We present an analysis on single-parameter variations in `COMPAS` for the chirp-mass distribution in Figure 3.11. The majority of the models exhibit a pronounced peak at $\mathcal{M}_c \simeq 5\,\mathrm{M}_\odot$, which is notably more robust than the peak of the BH primary-mass distribution in Section 3.4. This low-chirp-mass peak is dominated by non-MRR systems and remains largely unchanged across all explored prescriptions considered here.

The largest variations in the chirp-mass spectrum arise from changes to the mass-transfer efficiency β and the Wolf-Rayet wind strength f_{WR}. Increasing β enhances the relative contribution of MRR systems at lower chirp masses while suppressing their contribution at higher chirp masses, with the non-MRR population remaining comparatively stable. Variations in f_{WR}, particularly for stronger winds, flatten the MRR contribution across the chirp-mass range and reduce the non-MRR rate, producing the largest overall change among the tested parameters. In contrast, variations in the common-envelope efficiency α_{CE}, supernova engine prescription, and the BH natal kick dispersion σ_{BH} primarily affect the low-mass end of the chirp-mass distribution, akin to the primary-mass distribution, with limited impact on the relative MRR contribution at higher chirp masses.

While much of our discussion in Section 3.5 focused on the component-mass and mass-ratio distributions, chirp mass provides a complementary and observationally robust diagnostic. We find that the low-chirp-mass peak, which is dominated by non-MRR systems, is stable across a wide range of binary-evolution prescriptions. In contrast, MRR systems preferentially contribute at higher chirp masses, with their relative importance varying modestly under changes to mass transfer efficiency and Wolf-Rayet wind strength. These results demonstrate that while the overall normalization and detailed shape of the chirp-mass distribution can vary, the qualitative separation between MRR- and non-MRR-dominated regimes is robust to reasonable changes in binary-evolution physics. This robustness strengthens the case for using chirp mass, in conjunction with component masses and mass ratio, to assess

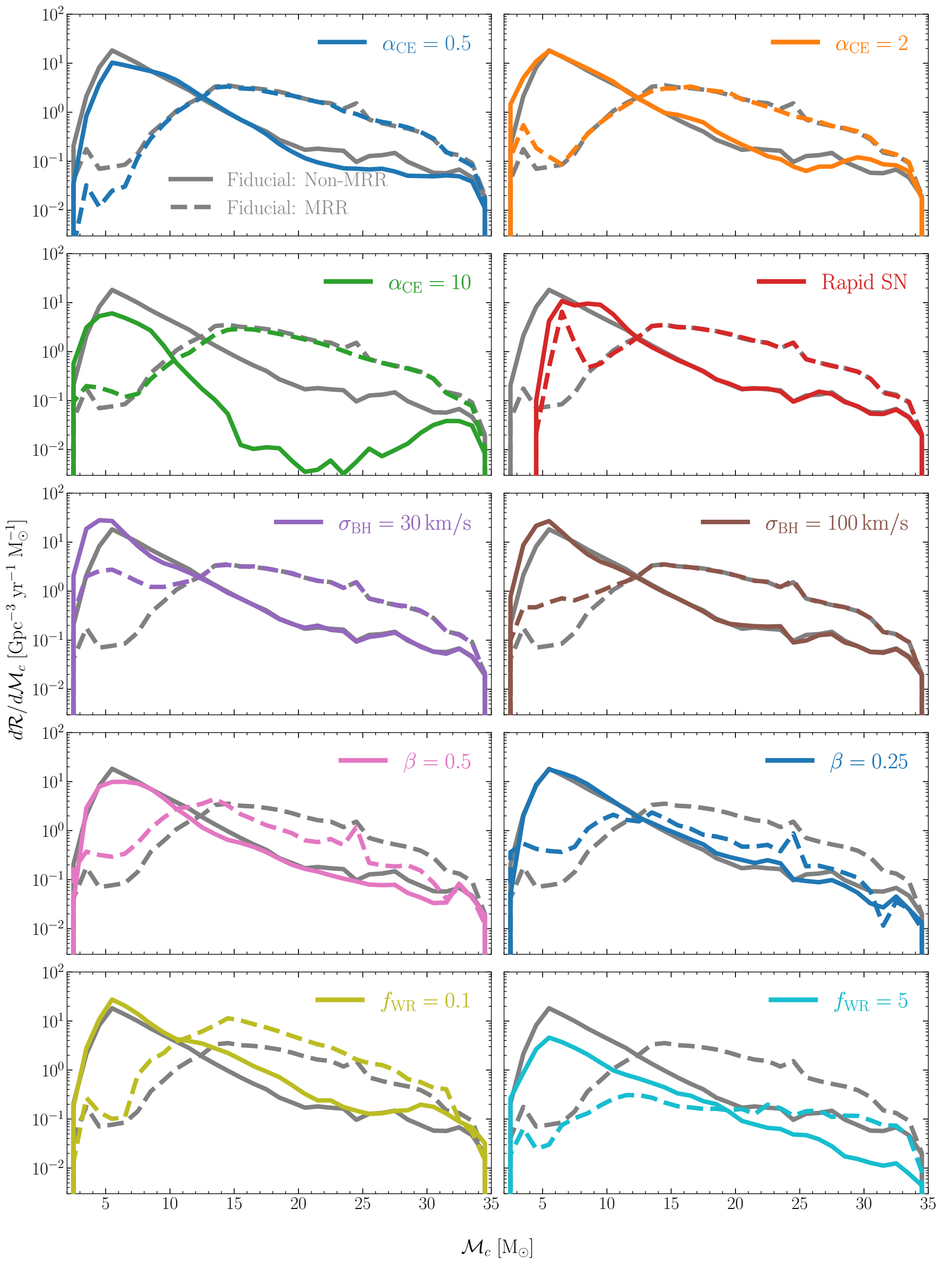

Figure 3.11: The chirp-mass distribution is shown for variations in binary parameters relative to the fiducial model, following Broekgaarden et al. (2022). Each panel compares the MRR and non-MRR contributions for the fiducial model and a single parameter variation. While individual prescriptions modify the overall normalization and relative MRR contribution, the low-mass peak, dominated by non-MRR systems, remains robust across each variation.

the potential contribution of MRR in LVK observations.

Appendix 3C: MRR channels - formation efficiency yields

We showed in Section 3.4.2 that MRR can arise from several distinct evolutionary channels. In particular, that core-growth is the dominant channel in both `COMPAS` and `SEVN` when the distributions are weighted by the BBH merger rate at $z \approx 0.2$. We discuss in this appendix the initial parameter space when weighted instead by the formation efficiency (Equation 3.2).

In `COMPAS`, there are minimal differences with the core-growth contribution dropping to 98% of the total MRR contribution. However, the PPISN-assisted core-growth channel does dominate in small regions of initial parameter space in particular at high primary masses of $M_{1,\rm ZAMS} \gtrsim 120\,\rm M_\odot$, secondary masses of 110–125 $\rm M_\odot$, and low mass ratios (Figure 3.12).

In `SEVN` we find that despite both PPISN-shrinking and asymmetric-CCSN almost never being the dominant MRR GW channel in `SEVN` (0.5% and 1.5% respectively), their formation yields are fundamentally more efficient (4.5% and 3.5% respectively) than what can be observed due to their formation likelihood.

The contribution from the PPISN-shrinking channel has the most noticeable difference between the two weighting schemes. For the primary mass distribution (left most column in Figure 3.12), the core-growth and asymmetric-CCSN channels are largely similar to that of the MRD weighted distributions. On the other hand, while the PPISN-shrinking channel contributed minimally to the MRD weighted distribution, it has a much stronger contribution to the overall formation-efficiency (which is marginalized over metallicity). We note that due to the metallicity evolution of the Universe, the contribution of the PPISN-shrinking channel may vary with cosmic time, i.e. for the MRD at earlier times we expect the fraction to increase.

The effects on the secondary-mass and mass ratio distribution are more subtle with PPISN-shrinking dominating for $\rm M_{2,ZAMS} \gtrsim 100~\rm M_\odot$ and $q \lesssim 0.7$.

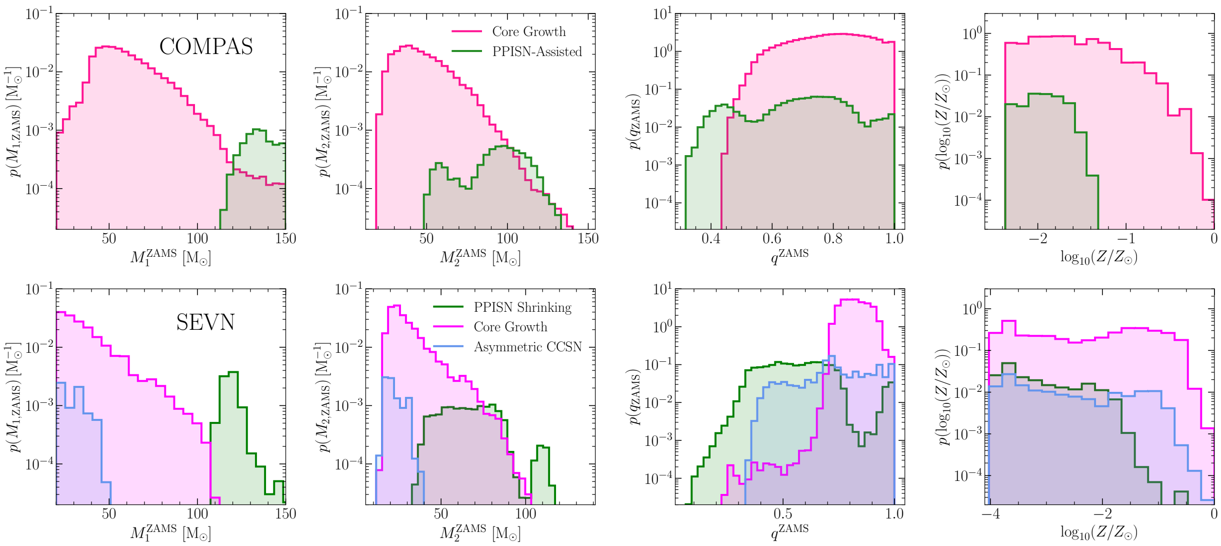

Figure 3.12: Same as Figure 3.7, but instead weighted by the formation efficiency.

Lastly, the metallicity distribution shows that the relevant metallicities for the MRD at $z \approx 0.2$ peaks at higher metallicities (Figure 3.7) with asymmetric-CCSN and PPISN-shrinking remaining subdominant over the full metallicity range. On the other hand, the formation-efficiency weighted distribution favors low metallicity. This is a reflection of the well-known result that formation-efficiency is highest for low metallicities and decreases as metallicity approaches solar metallicity (e.g. Belczynski et al., 2010b; Mapelli et al., 2010; Stevenson et al., 2017; Chruslinska et al., 2019; Giacobbo & Mapelli, 2018; Neijssel et al., 2019a; Spera et al., 2019; Broekgaarden et al., 2022; Schiebelbein-Zwack & Fishbach, 2024; Smith & Kaplinghat, 2026).

To summarize, the formation-efficiency weighted distributions have larger contributions from ppisn-assisted core-growth in `COMPAS` and asymmetric-CCSN and PPISN-shrinking in `SEVN`. While these differences mostly arise from the marginalization over metallicity/redshift, they can lead to these channels dominating in particular regions of initial parameter space.

Appendix 3D: Mass ratio reversal dependence on orbital separation

The orbital separation is sampled in `COMPAS` from [0.1, 1000] AU and from [0.07, 670] in `SEVN`. Figures 3.13 and 3.14 show contributions from MRR and non-MRR systems weighted by the formation efficiency and merger rate yield respectively for both `COMPAS` (left) and `SEVN` (right). Included for both codes is a deconstruction of the MRR contribution into constituent channels.

In both `COMPAS` and `SEVN`, the majority of merged BBH systems and the overall MRR fraction have a preference for short semimajor axes. The results by Schneider et al. (2015) show that orbital separation is a good proxy for determining which mass transfer case a system underwent. For instance, case A typically dominates below 10–100 AU, above this case B takes over up to several thousand solar radii followed by case C with a strong dependence on stellar mass. In these approximations, the majority of MRR systems undergo case A mass transfer. It would be interesting for future work to quantify the dependence and fraction of MRR on various mass transfer cases.

In `SEVN`, the PPISN-shrink channel peaks at slightly higher orbital separations, as slightly wider separations are less likely to disrupt the binary. The total fraction of MRR systems originating from the PPISN-shrink channel contribute a negligible fraction to the merger rate density. Asymmetric-CCSN follows a similar pattern to core-growth, but remains subdominant throughout both formation and MRD weighting. Similarly, in `COMPAS` the overall contribution from the PPISN-assisted core-growth channel is small compared to the dominant core-growth channel.

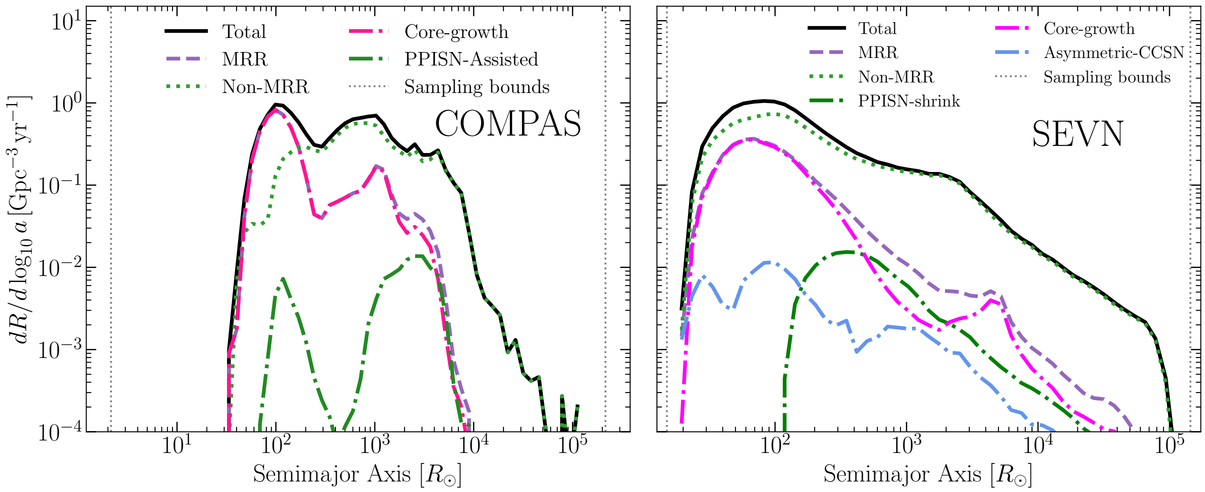

Figure 3.13: MRR and non-MRR contributions as a function of initial progenitor semimajor axis for `COMPAS` (left) and `SEVN` (right). Included is the evolutionary channel decomposition of the MRR contribution. This distribution is weighted by the formation efficiency (Equation 3.2).

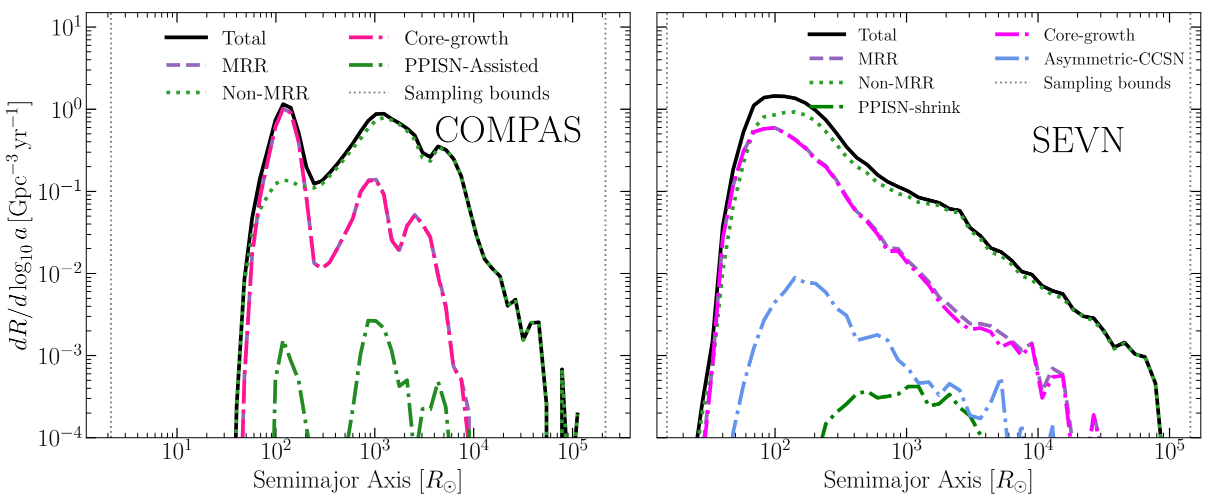

Figure 3.14: Same as Figure 3.13, but weighting by the contribution to the MRD at $z \approx 0.2$.

Appendix 3E: Do variations alter the massquerading effect?

We highlight in Figure 3.15 the effects of single parameter variations of the fiducial `COMPAS` model for the primary mass distribution. The results show that for each variation the *massquerading* effect is strongest in the high-mass regime, typically around $\mathrm{M} \gtrsim 20\,\mathrm{M}_{\odot}$. This highlights that this effect is not dependent on any single stellar/binary parameter, but is a robust outcome of stellar evolution found in `COMPAS`. If MRR plays a large role in redefining the primary and secondary mass according to LVK definition, then the component masses will contain a mixture of donor and accretor remnants.

Appendix 3F: Initial star properties mapping to binary defined BBH properties

We show in Section 3.4.2 how the initial parameter space is populated with MRR systems and the mapping, across metallicities, to the "LVK" defined mass space, i.e. M_1 is the more massive BH and M_2 the less massive. In this appendix, we also include the mapping between the initial parameters and the "binary" defined BHs (Figure 3.16), i.e. initially more massive star is the primary.

Interestingly, `COMPAS` remains fairly consistent in the overall 2D distributions, while for `SEVN` a region of MRR abundance becomes clear in the low-metallicity and low-mass regime of the binary-defined primary BH and high-mass binary-defined secondary BH spaces. These systems are a small subset of the high-mass contribution and are still overcome by the non-MRR systems in the merger rate distributions (Figure 3.2).

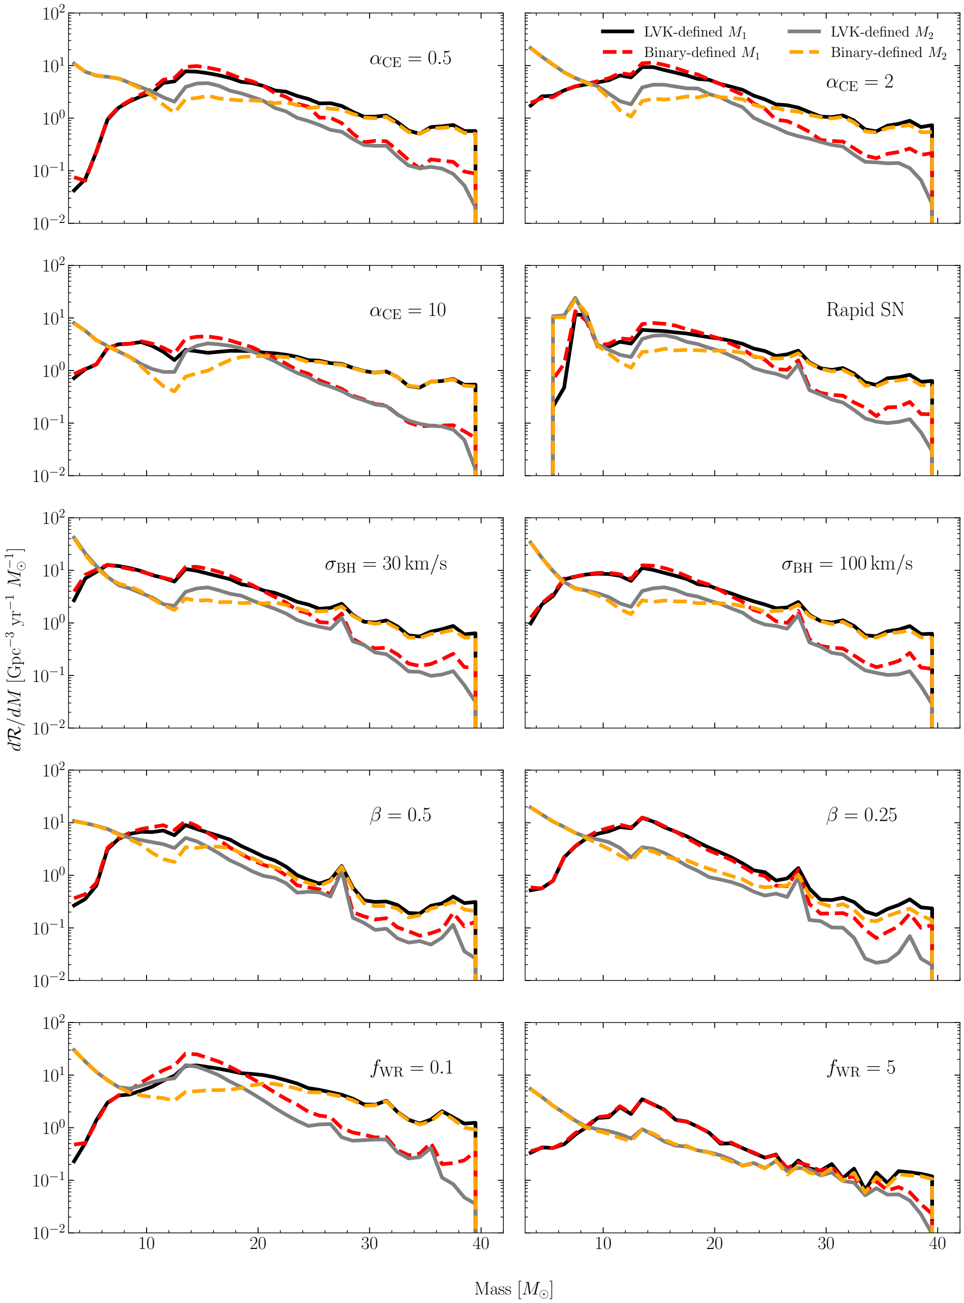

Figure 3.15: Variations of the massquerading effect based on single-parameter variations in `COMPAS`. The effect occurs mostly for $\gtrsim 20\,\mathrm{M}_\odot$ as this is where MRR constitutes a dominant fraction of systems. See Figure 3.8 for massquerading in the fiducial model.

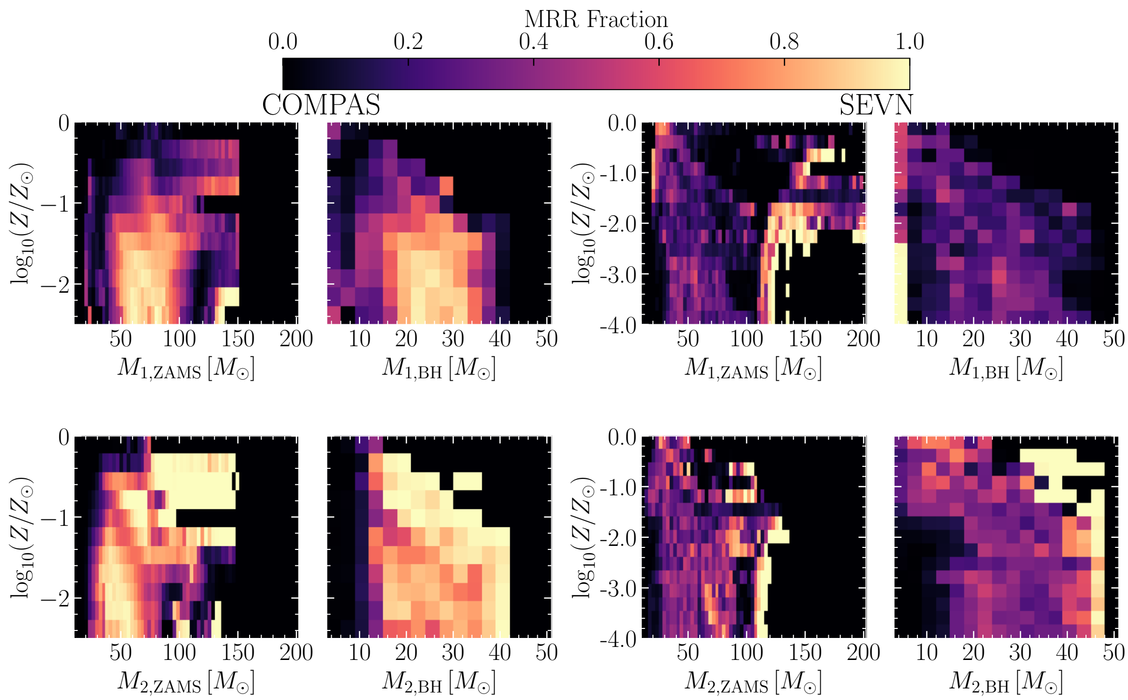

Figure 3.16: Same as Figure 3.5, but for the remnant masses defined with M_1 the most massive star at ZAMS and M_2 the less massive. For `COMPAS`, the MRR fractions do not shift by a large amount. However, for `SEVN` an abundance of MRR systems becomes clear at low masses for $M_1 \lesssim 10\,\mathrm{M}_\odot$ and $Z \lesssim -2.5$, additionally MRR becomes dominant for $M_2 \gtrsim 30\,\mathrm{M}_\odot$.

Chapter 4

Self-Interacting Dark Matter in the First Billion Years

4.1 Abstract

Self-interacting dark matter (SIDM) is a well-motivated extension of the cold dark matter (CDM) paradigm, and its effects on halo structure have been extensively studied and constrained using the low-redshift Universe — dwarf galaxy rotation curves, strong and weak lensing, and the observed subhalo population. Recent JWST observations of massive, rapidly assembling galaxies at high redshift open a complementary regime for this work: at these early epochs, halos are younger and dynamically distinct from their low-redshift counterparts, offering novel sensitivity to SIDM models previously unexplored via established low-redshift probes. Motivated in part by the possibility that efficient gas accretion from merged satellites could help explain these early galaxies, we use this high-redshift regime

This chapter is an adaptation of work to be submitted to the Astrophysical Journal later this year. The coauthors on this work are Manoj Kaplinghat, Michael Ryan, and Maya Silverman.

to test SIDM directly, with a suite of cosmological N-body simulations spanning multiple resolutions, evolved to $z = 6$, comparing CDM to constant and velocity-dependent SIDM cross sections. We find the halo mass function is nearly identical across all models at these redshifts, indicating that SIDM does not alter the abundance of massive host halos capable of hosting the brightest high-redshift galaxies, establishing that any differences between models arise from internal halo structure rather than host abundance. We find that SIDM's structural effects are already present at early times: halos are systematically rounder in SIDM, with clear separation from CDM already present by $z \approx 15$, and surviving subhalo counts are suppressed relative to CDM by up to $\sim 75\%$ in our most massive hosts by $z = 6$, even though the total number of objects accreted is comparable across models, i.e. self-interactions enhance post-infall disruption, not the accretion rate itself. We use TNG50-1 to estimate the gas fraction liberated by this enhanced disruption and thus available for potential star formation. Finally, despite a non-negligible population of analytically-predicted gravothermal collapse candidates, we find no realized collapse in our simulated density profiles by $z = 6$; we attribute this to the merger-dominated environment of early halo assembly, in which the great majority of massive hosts experience mergers frequent and large enough to delay the onset of collapse. Taken together, our results show that SIDM remains consistent with CDM while producing measurable structural signatures, establishing the high-redshift Universe as a genuine, independent complement to low-redshift constraints on the SIDM cross section.

4.2 Introduction

The James Webb Space Telescope (JWST) has identified a large sample of galaxies well into the epoch of reionization (Naidu et al., 2022; Castellano et al., 2022; Finkelstein et al., 2022; Harikane et al., 2023b; McLeod et al., 2024; Donnan et al., 2024; Conselice et al., 2024; Castellano et al., 2025; Pérez-González et al., 2025), including a growing number of spectroscopically confirmed systems at $z_{\rm spec} > 10$ (Robertson et al., 2023; Curtis-Lake et al., 2023; Wang et al., 2023; Harikane et al., 2024; Carniani et al., 2024; Napolitano et al., 2025). Several of these galaxies exhibit extreme properties: stellar masses $\gtrsim 10^{10}\,M_\odot$ within the first $\sim$ 500–700 Myr (Labbé et al., 2023; Xiao et al., 2024), unusually high star-formation efficiencies ($>$ 5–10%; Boylan-Kolchin 2023; Harikane et al. 2024; Pérez-González et al. 2025; Xiao et al. 2024), and sizes or luminosities unexpected at such early times (Wang et al., 2023; Carniani et al., 2024). Taken together, these findings point toward rapid, efficient baryon assembly and present challenges to standard galaxy formation models in the $\Lambda_{\rm CDM}$ paradigm.

A range of explanations have been proposed within standard ΛCDM astrophysics. Bursty star formation histories can boost the observed UV luminosity of a galaxy well above what its underlying stellar mass would suggest, particularly in the low-mass systems that dominate at high redshift (Sun et al., 2023). A stellar initial mass function that becomes increasingly top-heavy at low metallicity and high gas density has also been shown to reproduce the observed UV luminosity functions (Trinca et al., 2024; Ventura et al., 2024; Hutter et al., 2025; Lu et al., 2025; Durrant et al., 2026). Alternative explanations explore the background cosmology itself, with studies showing that early dark energy, independently motivated as a resolution to the Hubble tension (Poulin et al., 2019; Escudero et al., 2022), systematically enhances the abundance of massive halos at these redshifts (Shen et al., 2024). Still other work argues the tension is less severe than initially reported once systematics in photometric mass and redshift estimation are properly accounted for (Krishnan & Abazajian, 2026). Absent a clear consensus, it remains worthwhile to explore whether physics beyond the standard collisionless dark matter model could also play a role.

Self-interacting dark matter (SIDM) is one such candidate. Elastic dark matter self-scattering has long been proposed as a resolution to a series of small-scale structure puzzles in CDM and its effects on halo structure and evolution have been studied extensively at low redshift (Tulin & Yu, 2018; Adhikari et al., 2025). Far less explored is how self-interactions operate on galaxy assembly in the early Universe. Because subhalo disruption is enhanced under SIDM (Spergel & Steinhardt, 2000), and because JWST's most extreme high-redshift galaxies appear to require unusually efficient conversion of accreted material into stars, a natural question is whether the increased merger and disruption rate of subhalos in SIDM could itself help explain the elevated star-formation efficiencies inferred from JWST. In this chapter we investigate this question directly using a suite of cosmological N-body simulations, and provide the first study of SIDM's impact on massive host halo populations in large cosmological volumes at these redshifts ($z \gtrsim 6$), complementing existing SIDM work that has largely focused on present-day halos or on zoom-in simulations of individual systems.

Another, distinct, consequence of self-interactions is gravothermal collapse of dark matter halos. SIDM halos initially resemble their CDM counterparts, since scattering rates are low at the densities present at early times, preserving the same predictions for early-Universe physics such as the CMB (Dave et al., 2001). However, halos are not isolated, static systems: velocity dispersion rises outward from the halo center, so the dense central regions are initially cooler than their surroundings. Self-interactions conduct heat inward, heating and expanding the core and thermalizing it into an isothermal profile. Eventually, continued expansion leaves the core hotter than the surrounding envelope, reversing the temperature gradient and driving heat flow outward instead. The core then contracts, densities rise, and scatterings become more frequent, triggering a runaway loss of energy and the onset of gravothermal collapse (Lynden-Bell & Wood, 1968; Balberg et al., 2002).

Analytic models predict that, for isolated halos, the time to core collapse can fall well below the Hubble time (Balberg et al., 2002; Outmezguine et al., 2023), with direct implications for early galaxy evolution. This has motivated the use of gravothermal collapse as a channel for producing the supermassive black hole seeds invoked to explain the recently discovered "little red dots" (LRDs), though the proposed mechanisms differ in how collapse is triggered. One avenue considered is baryon-assisted collapse, in which a condensed central baryonic mass deepens the gravitational potential and accelerates the onset of gravothermal instability relative to a dark-matter-only (DMO) halo (Feng et al., 2021; Feng et al., 2025). Roberts et al. (2025) instead invoke a subdominant, ultra-strongly self-interacting fraction of the dark matter, distinct from and far more strongly coupled than the bulk population, which collapses on a correspondingly shorter timescale. Jiang et al. (2025), by contrast, consider unassisted collapse: standard, single-component SIDM halos with elastic scattering, collapsing under gravothermal evolution alone, with no additional baryonic potential or exotic subcomponent invoked. All three studies, however, rely predominantly on analytic or semi-analytic frameworks that do not directly resolve a halo's gravothermal evolution within a live, hierarchically assembling cosmological environment. Our DMO simulations can test whether unassisted collapse of this kind is realized once a halo is embedded in such an environment, subject to realistic, ongoing mergers and accretion at these early times.

In this chapter, we present a suite of cosmological, dark matter-only N-body simulations spanning multiple box sizes, resolutions, and self-interaction cross sections, evolved from $z = 99$ to $z = 6$, to directly test both proposed channels by which SIDM could affect the earliest, most extreme galaxies uncovered by JWST: enhanced star formation via subhalo disruption, and gravothermal core collapse. We describe our simulation suite and the analytic collapse prescription we compare against in Section 4.3; present our results on the halo mass function, halo shapes, subhalo survival and merger histories, and evidence for gravothermal collapse in Section 4.4; discuss the physical origin of these results and their implications for gas accretion, star formation, and LRDs in Section 4.5; and summarize our conclusions in Section 4.6.

4.3 Simulations

We perform several large scale cosmological simulations using the publicly available `GIZMO` code (Hopkins, 2015) with initial conditions made using Multi-scale Initial Conditions (`MUSIC`) (Hahn & Abel, 2011) where we assume a flat $\Lambda_{\rm CDM}$ cosmology, with parameters $\Omega_m = 0.315$,

$\Omega_\Lambda = 0.685$, and $h = 0.674$. We analyze two simulation volumes: a `BigBox (BB)` with comoving side length $L = 35$ Mpc/h and a `LittleBox (LB)` with $L = 7$ Mpc/h. These runs are designed to probe complementary regimes: the larger box captures rare, massive halos, while the smaller box provides higher mass resolution for analyzing small-scale structure. `BigBox` and `LittleBox` are simulated with a DM particle mass resolution of $m_{\rm DM} = 5.2 \times 10^6\,{\rm M}_\odot$ and $4.0 \times 10^4\,{\rm M}_\odot$, respectively.

Each box was simulated for a cold dark matter (CDM) case as well as two self-interacting dark matter (SIDM) models: one with a constant cross section of $\sigma_0 = 70$ cm^2/g,[1] and another with a velocity-dependent cross section, implemented via the viscosity cross section:

$$\sigma_V(v) = \frac{6\sigma_0 w^6}{v^6}\left[\left(2 + \frac{v^2}{w^2}\right)\ln\left(1 + \frac{v^2}{w^2}\right) - \frac{2v^2}{w^2}\right], \tag{4.1}$$

with $\sigma_0 = 147$ cm^2/g and the characteristic velocity $w = 24$ km/s (Yang & Yu, 2022; Nadler et al., 2023). The characteristic velocity, w, is defined as the ratio of the mediator mass, m_ϕ to the dark matter mass, m_χ, i.e. $w \equiv m_\phi c/m_\chi$.

Dark matter halos are identified using the `Rockstar` halo finder (Behroozi et al., 2012a) and merger trees are constructed with `Consistent Trees` (Behroozi et al., 2012b). To minimize noise and ensure robust statistics we place a resolution cut of > 200 particles per halo, unless otherwise stated. This threshold yields a smooth transition between the low-mass limit of the `BB` and the high-mass end of the `LB`, ensuring completeness across the halo mass function.

4.3.1 Analytic prescription for gravothermal collapse

The timescale of gravothermal collapse for an isolated halo can be calculated following the analytic framework of Balberg et al. (2002) and Outmezguine et al. (2023) in which the time

[1] This is roughly 100× the cross section per unit mass of a standard billiard ball (0.64 cm^2/g) though see note in Robertson et al. (2017).

to collapse is given by:

$$t_{\rm coll} = \frac{1}{c_1} \frac{6-\alpha}{(3\alpha-2) - n\,(\alpha-2)}\, t_{c,0} \approx 400\, t_{c,0}\,, \tag{4.2}$$

where $\alpha = 4/\sqrt{\pi}$, $n = 4.5 \times 10^{-5}$, and $c_1 = 1.9 \times 10^{-3}$ are constants calibrated from gravothermal evolution models, and the numerical prefactor evaluates to approximately 400. The collisional timescale $t_{c,0}$ is:

$$t_{c,0} = \frac{2}{3\,\alpha\, C} \left(\sigma_0\, \rho_{c,0}\, v_{c,0}\right)^{-1}\,, \tag{4.3}$$

where C is an order-one numerical constant fit to simulations. We are not able to fit for C from these simulations because, as we discuss in Section 4.4.4, halos do not reach gravothermal collapse. We take $C = 0.6$, the conservative end of the values used in the literature (Essig et al., 2019; Outmezguine et al., 2023; Nishikawa et al., 2020; Palubski et al., 2024; Koda & Shapiro, 2011; Yang et al., 2023). In Section 4.4.4, we discuss how our results are impacted by varying C. The core density $\rho_{c,0}$ and core velocity dispersion $v_{c,0}$ are evaluated at the onset of gravothermal evolution, and Outmezguine et al. (2023) show these can be derived directly from the NFW profile of the halo:

$$\rho_{c,0} = 2.4\, \rho_s\,, \qquad v_{c,0} = 0.64\, V_{\rm max}\,, \tag{4.4}$$

where ρ_s is the NFW scale density and $V_{\rm max}$ is the maximum circular velocity. We compute both of these values by fitting an NFW profile to CDM halos using their virial mass ($M_{\rm vir}$) and virial concentration ($c_{\rm vir} = r_{\rm vir}/r_{\rm s}$) using `Colossus` (Diemer, 2018).

We apply this estimate exclusively to the constant cross-section model ($\sigma/m = 70\,{\rm cm}^2\,{\rm g}^{-1}$), as this model has the largest effective cross section at the halo velocities relevant to our sample and therefore represents the most optimistic scenario for gravothermal collapse.

4.4 Results

4.4.1 Halo Mass Function

As a first test of whether SIDM alters the population of halos capable of hosting JWST's most extreme high-redshift galaxies (Section 4.2), we present the halo mass functions (HMFs) measured from both our `BigBox` and `LittleBox` in Figure 4.1. Results are given at four redshifts, $z = 12$, 10, 8, and 6. The vertical line marks a conservative mass threshold of $M_{\mathrm{vir}} = 10^9\,\mathrm{M_\odot/h}$, above which the histograms correspond to the `BigBox` runs and below they correspond to `LittleBox` runs. We show the HMFs in this way to maintain consistent resolution and minimize noise across nearly six dex in mass. The HMF falls off at $10^{6.5}\ M_\odot/h$, below which halos are not resolvable under the 200 particle resolution limit. For comparison, we include analytic fits from Yung et al. (2025) and Fernández-García et al. (2026), as implemented in `Colossus` (Diemer, 2018); both are calibrated to high redshifts ($z \gtrsim 20$) and are thus well-suited for comparison with our simulation suite.

Across all redshifts, the simulated HMFs agree within a factor of $\sim 2-4$ with the analytic predictions above the resolution limit, validating that both the CDM and SIDM simulations produce halos aligned with these studies. Notably, we find minimal differences between CDM and SIDM models, consistent with the expectation that self-interactions do not alter host halo abundances.

4.4.2 Halo Shapes

While the total abundance of halos as a function of mass is consistent across CDM and SIDM models, we find noticeable differences in their internal properties. One such difference is the overall shape of halos. Figure 4.2 shows the ratios of the semi-major axes for the top 100 most massive halos in the `LittleBox` run at $z = 15$ ($M_{\mathrm{vir}}/h > 5.6 \times 10^7\,M_\odot$) and $z = 6$ ($M_{\mathrm{vir}}/h > 1.4 \times 10^9\,M_\odot$). Halo shapes are characterized by the principal axes of the best-fit

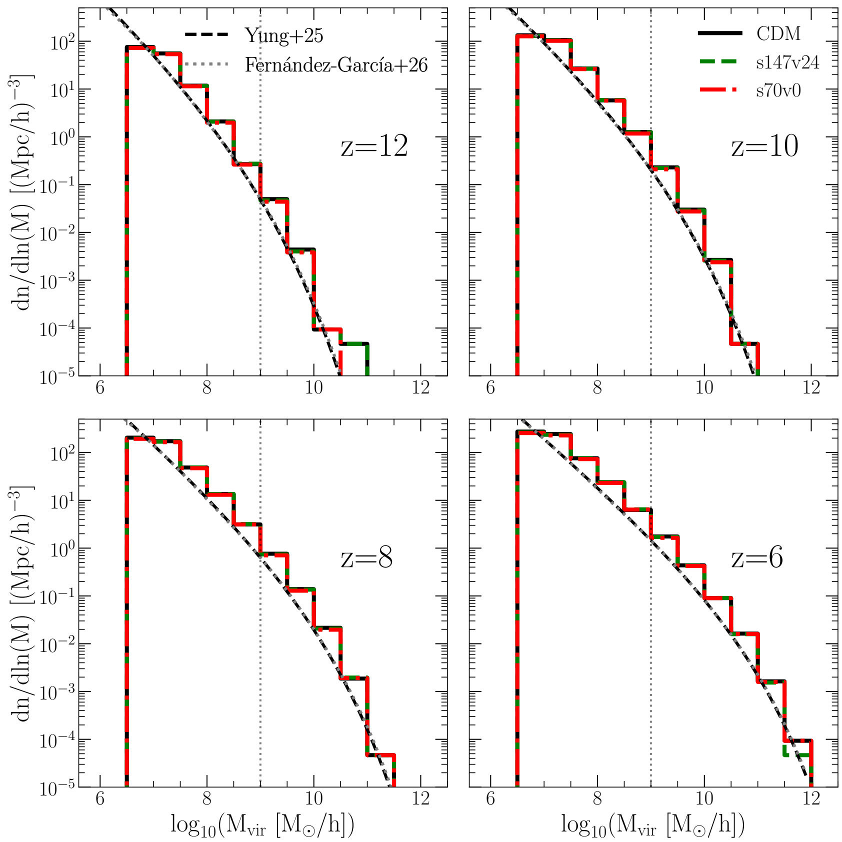

Figure 4.1: Halo mass functions from the `BigBox (BB)` and `LittleBox (LB)` simulations at $z = 12$, 10, 8, and 6. The grey vertical line marks our adopted transition mass of $M_{\rm vir} = 10^9\,M_\odot/h$, above which we use `BB` and below which we use `LB`. This threshold was chosen to provide a clean overlap between the two boxes as `BB` captures larger modes more reliably and `LB` provides the finer resolution to resolve low-mass halos. For comparison, we show analytic fits from recent works Yung et al. (2025) and Fernández-García et al. (2026), as implemented in `Colossus` (Diemer, 2018). Our simulated halo mass functions are in good agreement with these analytic predictions across the full mass range. CDM (black) and SIDM models with constant ($\sigma/m = 70\,{\rm cm}^2/{\rm g}$; red) and velocity-dependent ($\sigma/m = 147\,{\rm cm}^2/{\rm g}, v_0 = 24\,{\rm km\,s}^{-1}$; green) cross sections yield nearly indistinguishable halo abundances across all redshifts shown.

ellipsoid to the halo's mass distribution, $a \geq b \geq c$, computed following Zemp et al. (2011) in `Rockstar`. In the plane, b/a vs c/a, spherical halos live in the upper right, $\sim (1, 1)$, whereas more elongated or triaxial systems lie toward the lower left.

At both redshifts, the CDM sample (black points) spans a wide range of axis ratios, while increasing the effective cross section (CDM $\rightarrow$ s147v24 $\rightarrow$ s70v0) leads to higher sphericity,

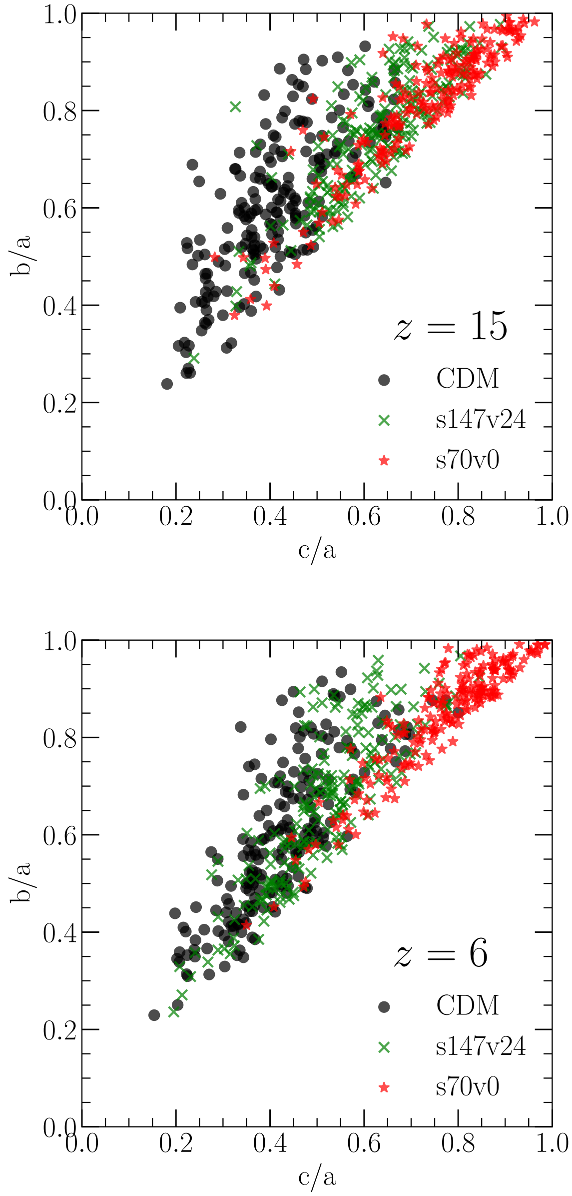

Figure 4.2: Halo shapes in the `LittleBox` simulations at $z = 15$ ($M_{\rm vir}/h > 5.6 \times 10^7\,M_\odot$, top) and $z = 6$ ($M_{\rm vir}/h > 1.4 \times 10^9\,M_\odot$, bottom). Each point shows the median-to-major (b/a) versus minor-to-major (c/a) axis ratio. Black circles correspond to CDM halos, green crosses to the velocity-dependent SIDM model with $v_0 = 24\,{\rm km\,s^{-1}}$, and red stars to the constant cross section model with $\sigma_0 = 70\,{\rm cm^2\,g^{-1}}$. SIDM halos are systematically rounder than their CDM counterparts at both redshifts, reflecting how self-interactions isotropize the particle velocities.

with the constant cross section model exhibiting clustering at $c/a \gtrsim 0.8$. This separation between CDM and the SIDM models is already apparent at $z = 15$ and persists to $z = 6$, indicating that self-interactions isotropize halo shapes within the first few hundred Myr of assembly and maintain this signature throughout the redshift range probed here. While the abundance of halos is comparable between DM models, the halo shapes can provide a clear discriminator between our models.

4.4.3 Subhalo Survival and Merger Histories

We define $N_{\rm sub}$ as the number of gravitationally bound subhalos identified within a host halo at a given redshift, and $N_{\rm total}$ as the cumulative number of objects that have either merged into or currently reside within the host halo over its history — that is, $N_{\rm total} = N_{\rm sub} + N_{\rm mergers}$, where $N_{\rm mergers}$ counts all resolved progenitors accreted along the main branch (the halo's single most massive progenitor at each prior snapshot, traced back through cosmic time).

Figure 4.3 shows the evolution of $N_{\rm sub}$ and $N_{\rm total}$ of host halos with $M_{\rm vir}/h > 5 \times 10^9\,M_\odot$ at $z = 6$, a threshold chosen to ensure each host has a sufficient number of resolved subhalos for $N_{\rm sub}$ to be a statistically meaningful quantity per halo. Solid lines show the mean across hosts and shaded bands indicate the 1σ scatter. Differences in $N_{\rm sub}$ between CDM and SIDM already emerge by $z \sim 12$ and grow progressively toward lower redshift, demonstrating that the suppression of substructure is not a sudden late-time effect but rather the cumulative result of enhanced disruption acting over an extended period of cosmic time. In contrast, $N_{\rm total}$ remains consistent across all models at all redshifts, confirming that host halos in all models experience comparable large-scale accretion histories. Self-interactions therefore enhance the disruption of subhalos after infall, not the rate at which they are accreted.

Figure 4.4 quantifies the mass dependence of this effect at $z = 6$. As host halo mass increases, the gap in the average number of subhalos per host, $\langle N_{\rm sub} \rangle$, between CDM and SIDM widens, with CDM retaining nearly 6x the number of subhalos that survive in the constant cross section model in the highest mass bin. The trend is clear across all models: larger effective cross sections produce fewer surviving subhalos at fixed host mass, consistent with enhanced

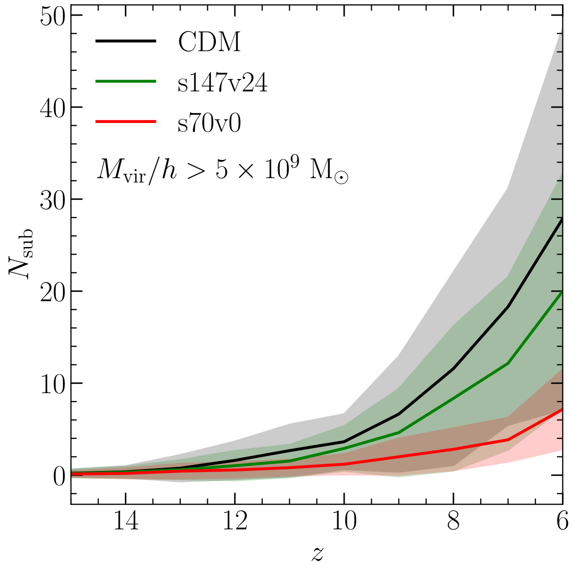

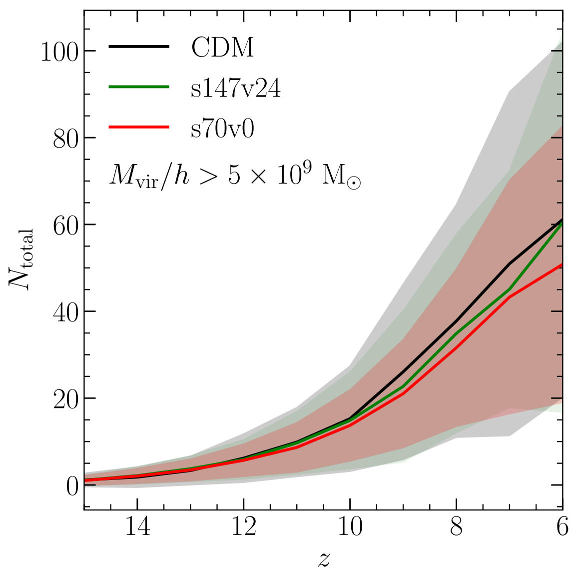

Figure 4.3: Redshift evolution of $N_{\rm sub}$ (top) and $N_{\rm total}$ (bottom) along the main progenitor branch of host halos with $M_{\rm vir}/h > 5 \times 10^9\, M_\odot$ at $z = 6$, for CDM (black), $\sigma/m = 147\,{\rm cm^2/g}$, $v_0 = 24\,{\rm km\,s^{-1}}$ (green), and $\sigma/m = 70\,{\rm cm^2/g}$ (red). Solid lines show the mean across all hosts in the sample and shaded bands indicate the 1σ scatter. $N_{\rm sub}$ counts the number of surviving bound subhalos at each snapshot, while $N_{\rm total} = N_{\rm sub} + N_{\rm mergers}$ additionally includes all resolved objects accreted along the main branch. The divergence between CDM and $\sigma/m = 70\,{\rm cm^2/g}$ is already present by $z \sim 12$ and progressively increases toward lower redshift, reflecting the cumulative effect of enhanced subhalo disruption over cosmic time. In contrast, $N_{\rm total}$ remains consistent across all models at all redshifts, confirming that the large-scale accretion history is insensitive to the self-interaction cross section.

post-infall disruption driven by self-interactions.

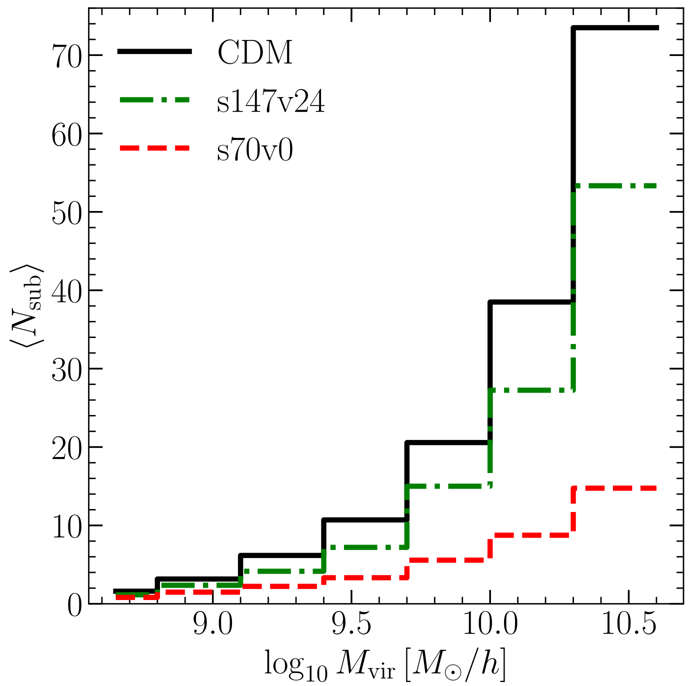

Figure 4.4: Average number of surviving subhalos per host at $z = 6$, $\langle N_{\rm sub}\rangle$, as a function of host mass for CDM (black, solid), $\sigma/m = 70\,{\rm cm}^2/{\rm g}$ (red, dashed), and $\sigma/m = 147\,{\rm cm}^2/{\rm g}$, $v_0 = 24\,{\rm km\,s}^{-1}$ (green, dot-dashed). The suppression of surviving subhalos in SIDM relative to CDM grows with host mass, reaching nearly a factor of 6 in the highest mass bin. While resolution effects may contribute to the apparent mass trend at the low-mass end, the physical expectation is consistent with this behavior: more massive hosts have deeper potential wells and richer merger histories, providing more opportunities for self-interactions to disrupt infalling substructure and making the effects of SIDM most pronounced in larger halos.

4.4.4 Gravothermal Collapse

The gravothermal collapse timescale in Eq. 4.2 depends on the NFW halo through its central density, ρ_s, and characteristic velocity dispersion, both of which are set entirely by a halo's concentration, $c_{\rm vir}$ and $M_{\rm vir}$. As a result, $t_{\rm coll}$ is itself a function of position in the $c_{\rm vir}$–$M_{\rm vir}$ plane. The predicted gravothermal collapse timescales calculated using the isolated halo approximation, defined in Section 4.3.1, to our simulated halos is given in Figure 4.5. By applying the analytic prescription to NFW profiles measured from our `LB` CDM halo catalog at $z = 6$, we find that of 10^5 halos above our resolution limit ($N_{\rm particles} > 200$), ~ 800 are predicted to collapse within the age of the Universe by $z \approx 6$. This number falls to ~ 150 when requiring a minimum of 10^3 particles per halo, and to 4 under a stricter 10^5-particle threshold motivated by the numerical convergence requirements identified by Palubski et al. (2024). [2] This strong sensitivity to the resolution cut reflects that the analytically-favored collapse candidates are concentrated among the lowest-mass, least-resolved halos in our catalog, so certifying a fully-collapsed core by the Palubski et al. (2024) standard is only possible for a handful of systems. Detecting the onset of collapse, however, does not require this level of resolution: a halo beginning to undergo core collapse should show a measurably enhanced central density relative to both CDM and its own earlier-time profile well before the innermost cusp is numerically converged. We therefore turn directly to the density profiles themselves, focusing on our 100 most massive host halos, each comprising at least 9.5×10^4 particles, comfortably above the 10^3-particle threshold and approaching the strictest 10^5 standard discussed above and the corresponding 4 collapse candidates at this resolution threshold.

We analyze the density profiles of our 100 most massive host halos for signs of gravothermal collapse in Figure 4.6. To highlight the temporal evolution of the halo cores we highlight two representative redshifts, $z = 10$ and 6, across our three models. We find that in the case of CDM the density profile is largely unchanged across the redshifts explored, as expected for purely gravitational interactions in a DMO simulation. For the SIDM models, as the effective cross section is increased we find that halos have lower central densities, indicating they are deeper in the core expansion regime by $z = 6$, with no clear signs of collapsed halos (e.g. densities far exceeding their previous densities at $z = 10$, or approaching densities comparable to CDM). Appendix 4A extends this analysis to the sample of s70v0 halos matched to the CDM collapse candidates (with $N_{\text{particles}} > 5 \times 10^4$) identified in Figure 4.5, tracking each halo's density profile continuously from formation to $z = 6$.

[2] These numbers increase to ~ 1200, ~ 300, and 10 at each respective resolution when a value of $C = 0.84$ is used.

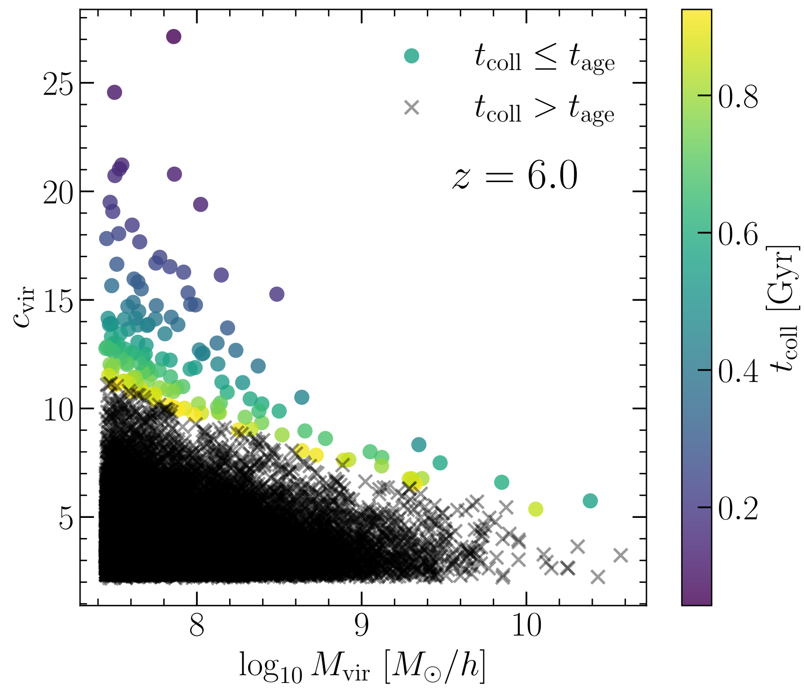

Figure 4.5: Concentration vs virial mass for halos in the `LB` CDM simulation at $z = 6$, colored by the predicted gravothermal collapse timescale $t_{\rm coll}$. Black points indicate halos with $t_{\rm coll} > t_{\rm H}$, i.e., those not expected to collapse by $z = 6$ under isolated NFW initial conditions. NFW profiles are fit to halo structural parameters measured directly from the CDM halo catalog, and collapse timescales are computed following the analytic prescription of Outmezguine et al. (2023). Of the $\sim 10^5$ halos in the full sample, ~ 800 are predicted to collapse, falling to ~ 150 when requiring a minimum of 10^3 particles per halo, and to 4 under a stricter 10^5 particle threshold motivated by the numerical convergence study of Palubski et al. (2024).

This picture, central densities decreasing but never approaching collapse by $z = 6$, and thus deeper in the coring regime, can be sharpened by tracking each halo's evolution continuously rather than at two fixed epochs. We show in Figure 4.7 the top 100 most massive halos in the $\rho_{\rm c}$–$r_{\rm c}$ plane. We exclude halos with a concentration below one standard deviation from our mean concentration ($C_{\rm vir} \sim 2$) and those for which an NFW halo is a poor fit as the analytic prescription requires input from NFW parameters, leaving 86 halos in the sample. The color bar indicates the concentration of the halos Halos trace a characteristic curve in this plane: they enter on the lower track as they approach coring, with decreasing $\rho_{\rm c}$ and increasing $r_{\rm c}/r_{\rm s,Kly}$ (where r_c is the core radius, $r_{\rm s,Kly}$ is the scale radius (Klypin et al., 2011;

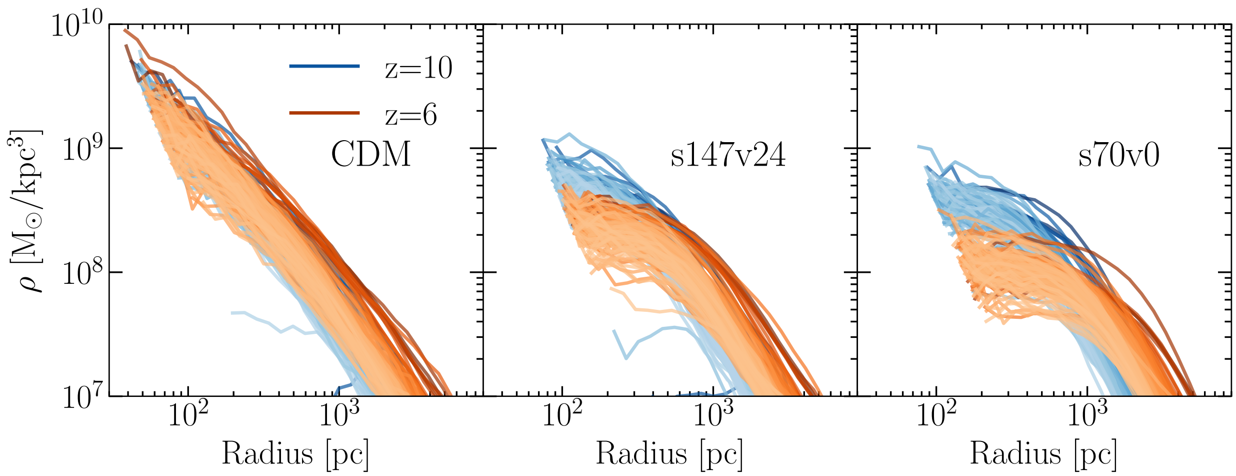

Figure 4.6: Radial density profiles for the top 100 host halos with assuming CDM, a velocity dependent cross section of $\sigma/m = 147\,\mathrm{cm}^2/\mathrm{g}, v_0 = 24\,\mathrm{km\,s}^{-1}$, and a constant cross section of $\sigma/m = 70\,\mathrm{cm}^2/\mathrm{g}$. The profiles are largely unchanged in CDM, however as we increase the effective cross section the profiles start to core as we approach later times, i.e. $z = 6$. There are no clear signs of gravothermal collapse in our sample.

Behroozi et al., 2012a)). Upon reaching minimum core density and maximum core radius, halos turn around and follow the upper track as they undergo gravothermal collapse.

We show the median track followed by our halo sample as the black dotted line, where each point on the line corresponds to a different snapshot/redshift. Our halos start, interestingly, from the lower left (likely due to resolution effects at early times), after which they follow the track and approach coring. This figure shows that by $z = 6$, most halos are deep in the coring phase, but that they have not quite turned around and started to undergo gravothermal collapse, in agreement from the findings in Figure 4.6.

Overall, we find no clear evidence of gravothermal core collapse in our simulations by $z = 6$. This is not necessarily evidence that collapse is absent in all halos at $z = 6$ as our sample is limited to a finite comoving volume. Gravothermal collapse may occur in halos residing in quieter, more isolated environments than those captured here, or at masses below our resolution limit. We discuss the physical origin of this non-detection, in particular the role of mergers and environmental heating in delaying collapse relative to isolated halo predictions,

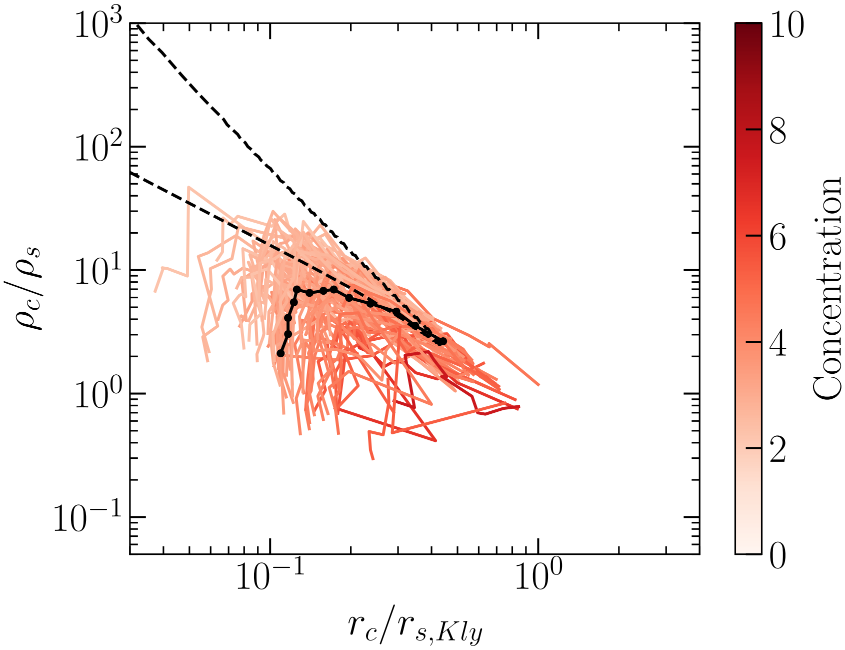

Figure 4.7: All massive `LB` s70v0 halos in the core-parameters plane, shaded by concentration. The median values at each timestep are given as the solid black line while the dashed black line represents the analytic prediction. We find that the median values tend to approach the track and then core, however we do not find evidence that they begin to turn around towards collapse.

in Section 4.5.

4.5 Discussion

Taken together, our results paint a consistent picture: the large-scale structure formation, halo abundance (Figure 4.1) and the total number of objects accreted over a halo's history (Figure 4.3, bottom panel), is unaffected by self-interactions until at least $z = 6$, while every diagnostic that is sensitive to a halo's internal dynamics: shape, subhalo survival, and core structure already show a clear SIDM imprint by $z \sim 12$–15, with the difference to CDM increasing until $z = 6$. The differences between CDM and SIDM are expected from late-time evolution (see Tulin & Yu (2018) and Adhikari et al. (2025); SIDM is built to leave ΛCDM's large-scale successes intact while modifying small-scale structure. What is new here is how early that imprint is established. The halos in our sample are still assembling at these redshifts, yet the divergence from CDM is already present and, in the case of subhalo survival, still growing. In the following subsections we discuss the physical origin of this early divergence (Section 4.5.1), its implications for gas accretion and star formation efficiency in the earliest galaxies (Section 4.5.2), and why, despite this ongoing small-scale evolution, we find no evidence for gravothermal core collapse in our sample (Section 4.5.3).

4.5.1 From large to small scales

The first departure from CDM appears not in the abundance of halos but in their internal structure. CDM halos span a wide range of axis ratios in Figure 4.2, reflecting the triaxial nature of collisionless structure formation (Dubinski & Carlberg, 1991). Increasing the effective cross section systematically pushes halos toward sphericity, with clear separation from CDM by $z \approx 15$. This is the same isotropization mechanism established in SIDM halos at $z \approx 0$ (Tulin & Yu, 2018) operating on the very halos assembling the first massive galaxies. Shape is therefore a probe of a halo's internal dynamical state and can be an early indication of self-interactions reshaping structure within the first billion years.

The same scattering that isotropizes host halos also governs the fate of their subhalos, and it is at the subhalo scale that the effect is largest. CDM's efficient hierarchical assembly produces an abundance of surviving subhalos (Kauffmann et al., 1993); SIDM reduces this population (Bhattacharyya et al., 2022), and we find the suppression is already substantial within the first billion years (Figures 4.3 and 4.4). Critically, this is a survival effect, not an accretion effect: $N_{\rm total}$ tracks CDM at all redshifts, while $N_{\rm sub}$ falls increasingly behind, down to $\sim$25% of the CDM count by $z = 6$ for $\sigma/m = 70\,{\rm cm}^2/{\rm g}$, marginalized over halo mass. Physically, this is the same momentum exchange responsible for the shape trend, now acting on infalling substructure: the higher the effective cross section the further cored our halos are and the easier it is to disrupt them (Dave et al., 2001).

This also explains the mass dependence in Figure 4.4: the suppression is modest at $z = 12$

and grows to nearly a factor of 6 in our most massive hosts by $z = 6$. More massive hosts have deeper potential wells and richer merger histories, so infalling subhalos both move faster and pass through the host more often, giving self-interactions more opportunities to act. A purely numerical origin cannot explain this trend on its own: our resolution cut is identical across all three models, so an incomplete sample would suppress subhalo counts by the same factor in each, not produce a widening, cross-section-dependent gap. The mass trend is therefore likely physical and it means the SIDM signature we find here should be most pronounced in precisely the massive host halos JWST is best equipped to identify at high redshifts.

4.5.2 Suppressed substructure and gas inflow

Section 4.5.1 established that subhalo suppression in SIDM is a survival effect, not an accretion effect, and that it grows with host mass, reaching a factor of ~ 6 relative to CDM in our most massive hosts by $z = 6$ (Figure 4.4). This raises a natural question: what happens to the gas these disrupted subhalos were carrying?

To estimate the gas budget available in these early (sub)halos, we turn to the TNG50-1 simulation suite. Figure 4.8 shows the ratio of satellite gas mass to host gas mass enclosed within a given radius, for the 1000 most massive host halos in TNG. This ratio declines as the Universe evolves, as host halos accrete an increasing share of their gas directly, through mergers and cosmic filaments, rather than via satellites.

A simple, well-defined estimate can bracket the plausible gas budget freed by this enhanced disruption using the excess number of disrupted subhalos together with the median gas fractions measured in TNG50-1. In our most massive host bin, we find $\langle N_{\rm sub} \rangle \approx 75$ surviving subhalos in CDM versus ≈ 15 in SIDM ($\sigma_0/m = 70\ {\rm cm^2\,g^{-1}}$) at $z = 6$, so roughly 60 subhalos that survive in CDM are instead disrupted in this model. Taking a conservative estimate of order 50 excess disrupted subhalos and assigning each the median satellite gas fraction

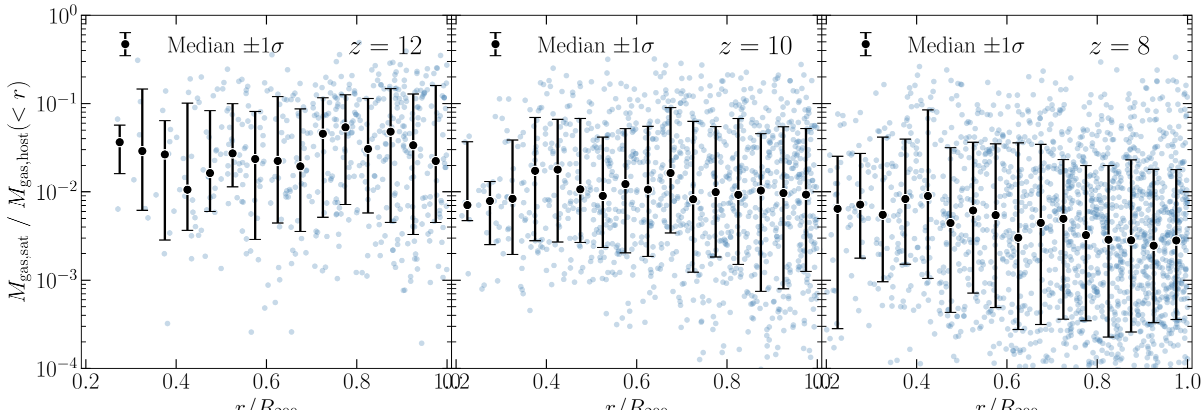

Figure 4.8: Fraction of gas in satellites compared to the host gas mass at a given radii. The blue scatter are individual satellites, while the black points are the median value in that bin with the corresponding standard deviation. As time progresses, the satellites contain less overall gas compared to the host.

from Figure 4.8 (∼2–4% of the host's gas mass), the summed gas content carried by these subhalos amounts to ∼100–200% of the host's gas mass. This is a substantial reservoir: even allowing for incomplete accretion of stripped gas onto the host, it suggests SIDM-enhanced subhalo disruption could meaningfully boost the gas supply available for star formation in precisely the massive, early hosts JWST is most sensitive to. We stress that this estimate assigns the population median gas fraction to the excess disrupted subhalos rather than tracking their individual gas content, so it does not capture whether these particular objects are gas-rich or gas-poor relative to that median; hydrodynamic simulations of SIDM that self-consistently track gas accretion alongside subhalo disruption would be needed to pin this down accurately.

4.5.3 Absence of Gravothermal Collapse

Section 4.4.4 showed that while the analytic prescription predicts a non-negligible population of collapse candidates, we find no realized collapse in our density profiles or in the ρ_c–r_c diagnostic (Figures 4.6 and 4.7): halos enter the coring phase but do not collapse. We now argue this is a direct consequence of the merger-dominated environment in which these halos assemble.

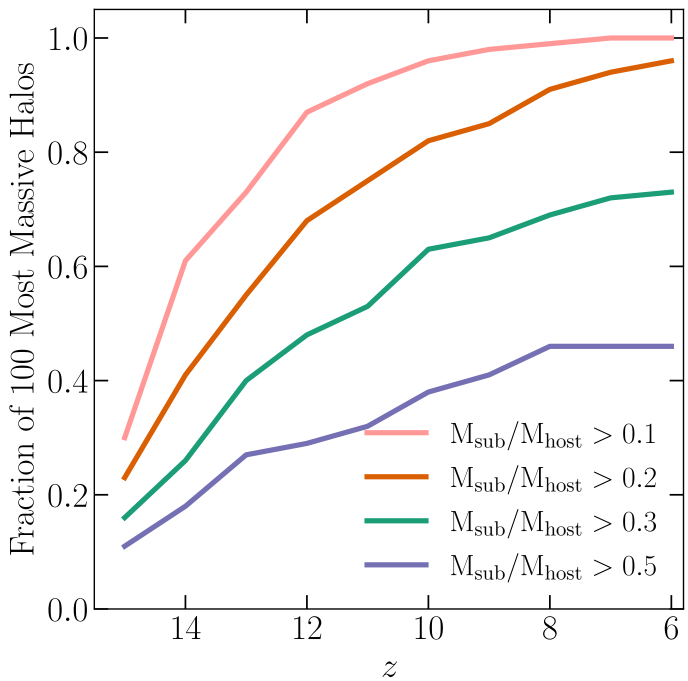

Figure 4.9: Cumulative fraction of our 100 most massive host halos at $z = 6$ that have experienced at least one merger above a given mass ratio, $M_{\rm sub}/M_{\rm host}$, by a given redshift, beginning at $z = 15$. Curves show thresholds of $M_{\rm sub}/M_{\rm host} > 0.1$, 0.2, 0.3, and 0.5. By $z = 6$, essentially all of our most massive halos have experienced at least one merger with $M_{\rm sub}/M_{\rm host} > 0.1$, a fraction that saturates by $z \sim 8$; more than 95% have undergone a merger with $M_{\rm sub}/M_{\rm host} > 0.2$; $\sim 73\%$ have experienced a merger with $M_{\rm sub}/M_{\rm host} > 0.3$; and $\sim 46\%$ have experienced a near-equal-mass merger with $M_{\rm sub}/M_{\rm host} > 0.5$.

Figure 4.9 quantifies how pervasive mergers are for our 100 most massive host halos at $z = 6$ ($M_{\rm vir}/h \in [2.6 \times 10^9, 3.5 \times 10^{10}]\, M_\odot$), the same sample used in Figures 4.6 and 4.7. It shows the cumulative fraction of these halos that have experienced at least one merger above a given mass ratio, $M_{\rm sub}/M_{\rm host}$, since $z = 15$. By $z = 6$, essentially all ($\sim 100\%$) of our most massive halos have experienced at least one merger with $M_{\rm sub}/M_{\rm host} > 0.1$, a fraction that saturates already by $z \sim 8$. More than 95% have undergone a merger with $M_{\rm sub}/M_{\rm host} > 0.2$, nearly three-quarters ($\sim 73\%$) with $M_{\rm sub}/M_{\rm host} > 0.3$, and close to half ($\sim 46\%$) have experienced a near-equal-mass merger with $M_{\rm sub}/M_{\rm host} > 0.5$. As a whole, the median latest time to encounter a merger of ratio $M_{\rm sub}/M_{\rm host} > 0.2$, 0.3, and 0.5 is $z = 8$, 9, and 10 respectively. By any reasonable accounting, then, a genuinely quiescent window, one free of major mergers, is the exception rather than the rule for halos in this mass range at these redshifts.

These findings align with those highlighted by Silverman et al. (2026); mergers can delay gravothermal collapse by injecting kinetic energy into the host halo. Silverman et al. (2026) find that while more massive mergers are more likely to impact collapse, other parameters, including merger radiality and infall time, also play a role. Although we do not have the time resolution to study radiality and infall time in detail, we present full density-profile histories of a handful of halos that do not undergo gravothermal collapse in Appendix 4A.

These halos highlight that the actual disruption to, and delay of, gravothermal collapse is more nuanced than a single fixed mass ratio can capture. A sample extending beyond the six halos considered in Silverman et al. (2026), together with longer simulation times and higher resolution than achieved here, would help quantify the relationship between merger ratio and collapse delay more systematically.

We stress that this does not preclude collapse at later times. Silverman et al. (2026) find that mergers occurring sufficiently early in a halo's gravothermal evolution, i.e. before maximum core expansion, do not permanently suppress collapse, just delay it. Given how massive and frequent mergers are for halos with $M_{\rm vir} > 2.6 \times 10^9 \ M_\odot/h$ at redshifts between 15 – 6 (Figure 4.9), it is perhaps unsurprising that our density profiles instead show halos still actively evolving through the core-expansion phase at $z = 6$ (Figure 4.7). Whether these halos ultimately collapse depends on their future merger history: if major mergers subside after $z = 6$, the conditions for gravothermal collapse could be met on timescales accessible to lower-redshift observations. This makes our high-redshift halo population an interesting target for follow-up zoom-in simulations evolved to $z = 0$.

This has direct bearing on a proposed explanation for the "little red dots" (LRDs) recently identified by JWST: several studies have suggested that gravothermal core collapse in SIDM can produce the supermassive black hole seeds needed to explain these systems (Jiang et al., 2025; Feng et al., 2025; Roberts et al., 2025), building on analytic collapse timescales that, for isolated halos, can fall well below the Hubble time (Balberg et al., 2002; Outmezguine et al., 2023). These studies rely predominantly on analytic and semi-analytic modeling of isolated halos. Our simulations provide a direct test of whether collapse is realized once halos are embedded in a full cosmological environment. We find that it is not, at least not by $z = 6$ for the cross sections tested here: the same merger-driven resetting that delays collapse in our sample would, by the same argument, delay or suppress the SMBH seed formation these models invoke.

Three caveats temper this conclusion. First, our simulated volume is finite, so a null detection does not rule out collapse occurring within our simulated redshift range; it is possible that halos in quieter, more isolated environments than those captured in our boxes do complete collapse, simply not among the halos we resolve. Second, we adopt $\sigma/m = 70\,\mathrm{cm}^2/\mathrm{g}$ as our most optimistic constant cross section, following the collisional constraints relevant at these halo velocities; a stronger effective cross section, whether from a different velocity dependence or a larger normalization, would shorten the gravothermal timescale and could plausibly overcome the merger-driven delays we identify here. Third, our simulations may lack the resolution and numerical treatment needed to reliably capture gravothermal collapse even in halos that are otherwise good candidates. Palubski et al. (2024) show that adaptive force softening, the standard approach in GIZMO, can introduce numerical heating that competes with the genuine self-interaction-driven heat transport, in some regimes artificially accelerating collapse and in others stalling or reversing it, and that consistent, converged collapse times require of order 10^6 particles within the halo, with $\sim 10\%$ scatter persisting even at that resolution. Our most massive `LB` host halos are resolved with of order 5×10^5 particles, and the bulk of our sample falls well below this benchmark. We therefore cannot cleanly separate the merger-driven delay we argue for above from a resolution-driven suppression of collapse: it is possible that some fraction of our sample would show collapse signatures in a numerical scheme and resolution regime built for that purpose. Disentangling these two effects will require dedicated higher-resolution, idealized-halo comparisons of the kind performed by Palubski et al. (2024) and Mace et al. (2024), ideally applied to halos drawn from, or embedded in, the same cosmological merger histories characterized here.

With these caveats in mind, our results still suggest that a realistic accounting of merger histories is a necessary ingredient for SIDM-driven LRD formation channels, and that the halos capable of hosting collapsed cores at a given epoch may be preferentially those that have recently experienced a sufficiently quiescent window, rather than the general host population assumed in isolated-halo treatments. LRDs are observed across $4 \lesssim z \lesssim 9$ (Labbé et al., 2023; Kokorev et al., 2024), with the bulk of the population at $z < 6$, outside the redshift range probed here; testing the merger-history picture directly against the smaller population at $z \gtrsim 6$, and extending our simulations to lower redshift to confront the bulk of the observed sample, are natural follow-ups.

4.6 Conclusions

We have investigated the effect of self-interactions on dark matter halos in the first billion years of the Universe using a suite of cosmological N-body simulations spanning multiple box sizes, resolutions, and cross sections. Our main findings are as follows:

1. **Large-scale structure formation is unaffected by self-interactions.** The halo mass function (Figure 4.1) is nearly identical across CDM and both SIDM models at all four redshifts probed ($z = 12, 10, 8, 6$), in good agreement with analytic predictions calibrated to high redshift. SIDM therefore does not alter the abundance of massive host halos capable of hosting the brightest JWST observed galaxies at these epochs, if self-interactions play a role in the observed population it must operate below the host halo scale.

2. **Halo shape is the earliest and most persistent SIDM imprint.** Halos in our SIDM models are systematically rounder than their CDM counterparts, with clear separation already present by $z \approx 15$ and maintained through $z \approx 6$ (Figure 4.2). This isotropization, driven by the same momentum exchange responsible for our other structural results, establishes that self-interactions are dynamically important within a few hundred Myr of halo assembly, far earlier than typically probed in SIDM studies.

3. **Subhalo suppression is a survival effect, not an accretion effect, and grows with host mass.** $N_{\rm total}$ remains consistent between CDM and SIDM at all redshifts, while $N_{\rm sub}$ in SIDM falls increasingly behind the CDM value, down to $\sim$25% of the CDM count by $z = 6$ for $\sigma/m = 70\,{\rm cm}^2/{\rm g}$ (Figure 4.3). The suppression is mass-dependent, growing to nearly a factor of 6 in our most massive hosts by $z = 6$ (Figure 4.4), consistent with deeper potential wells and richer merger histories giving self-interactions more opportunities to disrupt infalling substructure. This makes the SIDM signature identified here most pronounced in precisely the massive hosts JWST is best equipped to detect.

4. **Enhanced subhalo disruption plausibly boosts gas delivery in rare, massive hosts, though the population-median effect is modest.** The median satellite gas fraction measured in TNG50-1 is only $\sim$2–4% of the host's gas budget (Figure 4.8), weighted down by numerous low-mass satellites. However, individual gas-rich satellites reach well above 10% within $0.5\,R_{200}$, and applying this 10% benchmark to the $\sim 6\times$ SIDM disruption enhancement in our most massive hosts would deliver $\sim$60% additional gas relative to CDM. This suggests SIDM-enhanced disruption could meaningfully boost star formation specifically in the rare, massive galaxies JWST is most sensitive to at these redshifts, though translating delivered gas into realized star formation, and quantifying this effect directly, requires dedicated hydrodynamic simulations we leave to future work.

5. **We find no evidence of realized gravothermal collapse by $z = 6$.** While the analytic isolated-halo prescription predicts a non-negligible population of collapse candidates (Figure 4.5), our density profiles (Figure 4.6) and ρ_c–r_c diagnostic (Figure 4.7) instead show halos still actively deepening through the core-expansion phase, with no examples of turnaround. We attribute this to the merger-dominated environment of early halo assembly: each of our 100 most massive hosts has experienced a merger with $M_{\rm sub}/M_{\rm host} > 0.1$ by $z = 6$, and more than 95%, 73%, and 46% have exceeded thresholds of 0.2, 0.3, and 0.5, respectively (Figure 4.9). Each such event, whether a large merger or multiple small mergers, resets a halo's gravothermal clock, making the fully quiescent windows assumed in isolated-halo predictions the exception rather than the rule at these redshifts.

Taken together, these results paint a consistent picture of *where*, not *whether*, self-interactions matter in the early Universe: large-scale halo abundance and total accretion history are indistinguishable from CDM, while every diagnostic sensitive to internal dynamics — shape, subhalo survival, and core structure — carries a clear and, in the case of subhalo suppression, still-growing SIDM imprint by $z \sim 12$–15. This division is precisely what SIDM is designed to produce: it leaves ΛCDM's large-scale successes intact while reshaping small-scale structure. What we do *not* find is support for SIDM as a straightforward resolution to the JWST massive-galaxy tension via either channel we tested here: neither the modest gas budgets available from subhalo disruption nor the ongoing, merger-interrupted approach to gravothermal collapse obviously produces the efficient early star formation these systems appear to require. This should not be read as ruling out a role for self-interactions in early galaxy formation, but rather as narrowing the mechanisms by which that role could plausibly be realized, and underscoring that resolving it will require moving beyond the dark-matter-only approximation used here. This motivates two directions for future work: hydrodynamic simulations to directly test the star-formation response to subhalo-liberated gas, and how self-interactions modify that response relative to CDM, rather than bracketing it analytically; and simulations evolved to $z = 0$ to quantify, within a full cosmological framework, whether the ongoing core expansion we observe at $z = 6$ eventually completes and turns around to collapse as merger activity subsides at lower redshift, potentially via targeted zoom-ins to reach the resolution needed to track individual halos through collapse.

Appendix 4A: Collapse Candidates

Figure 4.10 shows the radial density profile evolution from each halo's earliest resolved snapshot (> 200 particles) to $z = 6$ pulled from s70v0 halos matched from CDM collapse candidates with $\geq 5 \times 10^4$ particles at $z = 6$ (Figure 4.11 is an analog for matched halos in s147v24) . The density profiles require at least 20 particles per radial bin. Each profile is colored by redshift, from yellow ($z \approx 20$) to dark blue ($z \approx 6$).

The inset panels show the core evolution, where the core is taken to be the density at 200 pc (a similar result is obtained if instead we use 200 particles as the core size). The star markers are colored by the largest merger mass ratio $M_{\mathrm{sub}}/M_{\mathrm{host}}$ the halo experienced within that snapshot interval and corresponds to the *viridis* colorbar in the lower right corner of the figure.

As with the broader 100 most massive halo sample in Figure 4.6, none of the matched candidates shown here exhibit a density profile indicative of completed gravothermal collapse. Instead, the sample reinforces the picture from Section 4.5.3: cores tend to deepen over cosmic time with brief increases in density. Subsequent mergers then reset the collapse timeline, with enough frequency that an actual collapse becomes too rare to appear in the simulated volume.

Halo 10187133 (middle row, first column) illustrates this behavior directly: its core density rises to a maximum at $z = 11$, immediately following a merger with $M_{\mathrm{sub}}/M_{\mathrm{host}} = 0.87$ (and a total incoming mass ratio from all progenitors/mergers of 2.35) completed between $z = 12$–11. The core density then drops steadily back into the coring phase over the following snapshots, and includes another 0.21 ratio merger in which the core density drops after. Its s147v24 counterpart (10256178) shows an almost identical merger history — $M_{\mathrm{sub}}/M_{\mathrm{host}} = 0.21$ at $z = 8$ and 0.88 at $z = 12$, matching the s70v0 values to within a few percent — consistent with merger timing itself being largely insensitive to cross section (Section 4.5.1).

The core density tracks the same shape in both models, but s147v24 sits higher than s70v0 by a factor of up to ~ 2.8 at $z = 6$, increasing monotonically from $z = 8$ (possibly beginning to collapse, if left alone), where s70v0 dips back downwards by $z = 6$. This hints that a similar merger occurring under a smaller effective cross section ($\sim 60 \mathrm{cm}^2/gatthesevelocities$) can lead to varying core densities. Finer snapshot resolution and tracking could ensure the difference doesn't arise from the stochasticity of the simulations themselves.

Halo 10796328 (middle row, third column) shows a similar trend: its core density increases before being interrupted by a near-equal-mass merger ($M_{\mathrm{sub}}/M_{\mathrm{host}} = 0.98$) at $z = 9$ and again by a $M_{\mathrm{sub}}/M_{\mathrm{host}} = 0.45$ merger at $z = 7$. Its s147v24 counterpart (10881478) shares several of the same mergers:$M_{\mathrm{sub}}/M_{\mathrm{host}} = 0.21$ at $z = 8$ and a near-equal-mass merger of ~ 0.99 at $z = 9$, and a 0.37 merger at $z = 7$ compared to the 0.45 merger in s70v0. While this merger event in s70v0 is sufficient to drop the density by the next snapshot, in s147v24 the density is found to still rise at $z = 6$, possibly on its way to collapsing. There is an additional resolved merger at $z = 11$ ($M_{\mathrm{sub}}/M_{\mathrm{host}} = 0.37$) which is not apparent in s70v0, despite comparable total mass growth at that step ($2.12\times$ vs. $2.01\times$). It is possible that the halo was disrupted in s70v0 before being registered as a merger, due to the larger cross section, though without finer timing resolution it is difficult to tell.

Halo 10487695 (bottom row, second column) experiences a particularly active history with mergers of $M_{\mathrm{sub}}/M_{\mathrm{host}} = 0.52$ at $z = 10$ and 0.34 at $z = 9$. Its profile becomes visibly diminished by $z = 7$, possibly due to a temporary centering issue or due to the consistent pummeling the halo experiences, it then recovers by $z = 6$ reflecting a typical cored profile shape. Unlike the other seven candidates shown here, Halo 10487695 has no s147v24 counterpart.

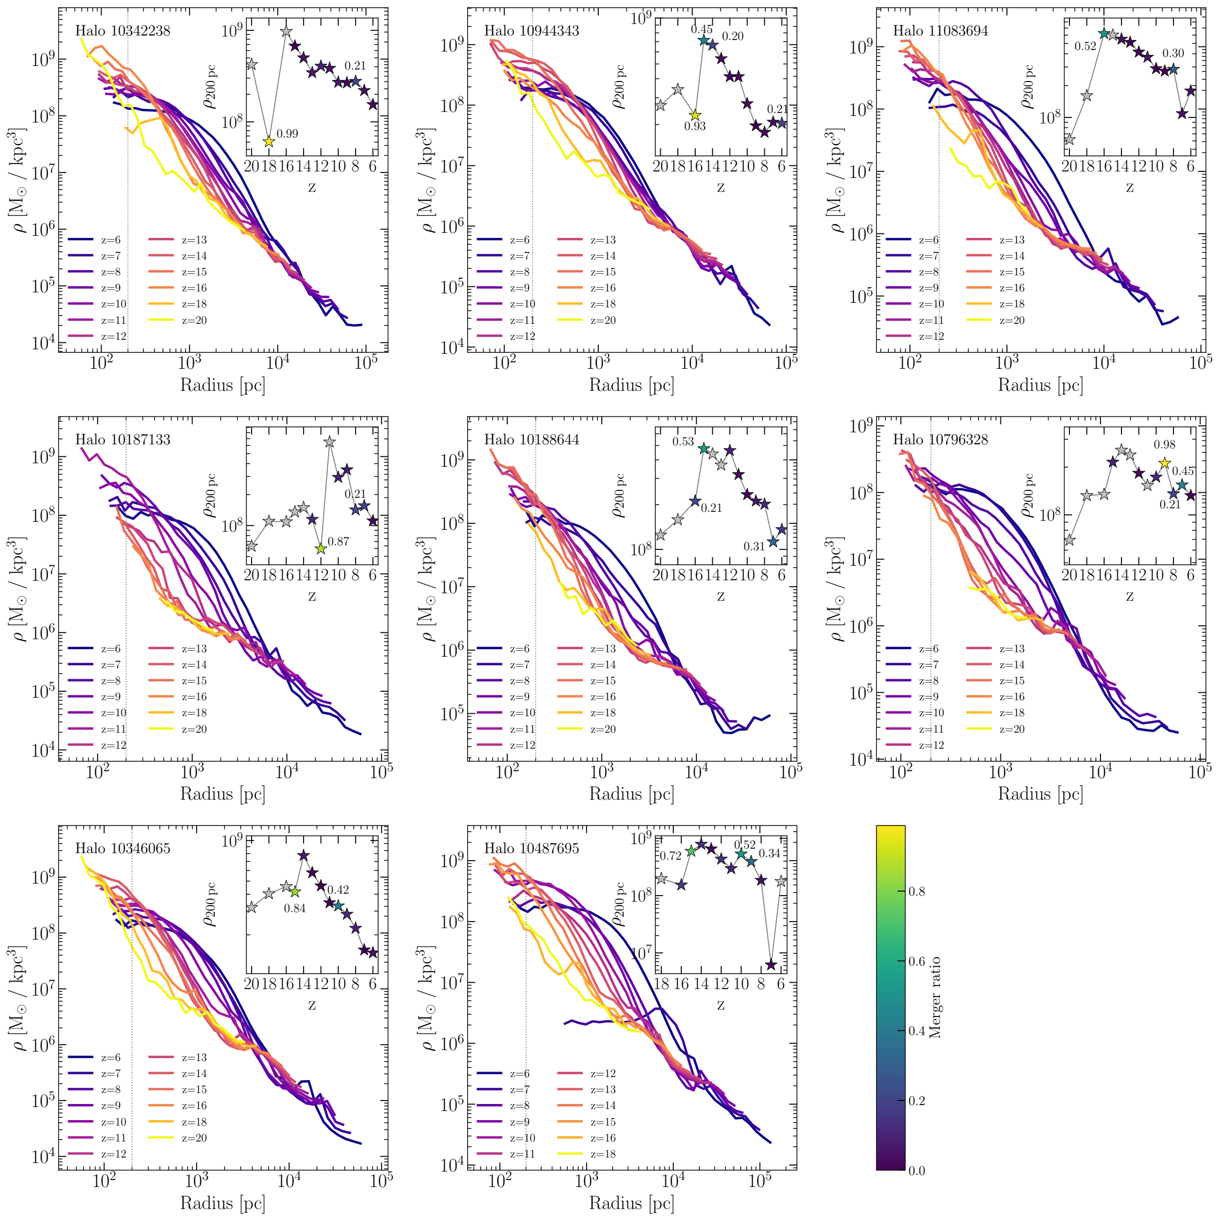

Figure 4.10: Radial density profile evolution for the s70v0 halos matched to CDM gravothermal collapse candidates (Figure 4.5), each comprising $> 5 \times 10^4$ particles at $z = 6$. Colors indicate redshift, from $z \approx 20$ (yellow) to $z = 6$ (dark blue). *Insets*: core density at 200 pc versus redshift; points are colored by the largest single merger mass ratio, $M_{\rm sub}/M_{\rm host}$, completed in that snapshot interval (viridis colorbar, bottom right). As in Figure 4.6, none of these halos show density profiles indicative of completed gravothermal collapse; cores repeatedly deepen and are then interrupted by mergers before turnaround.

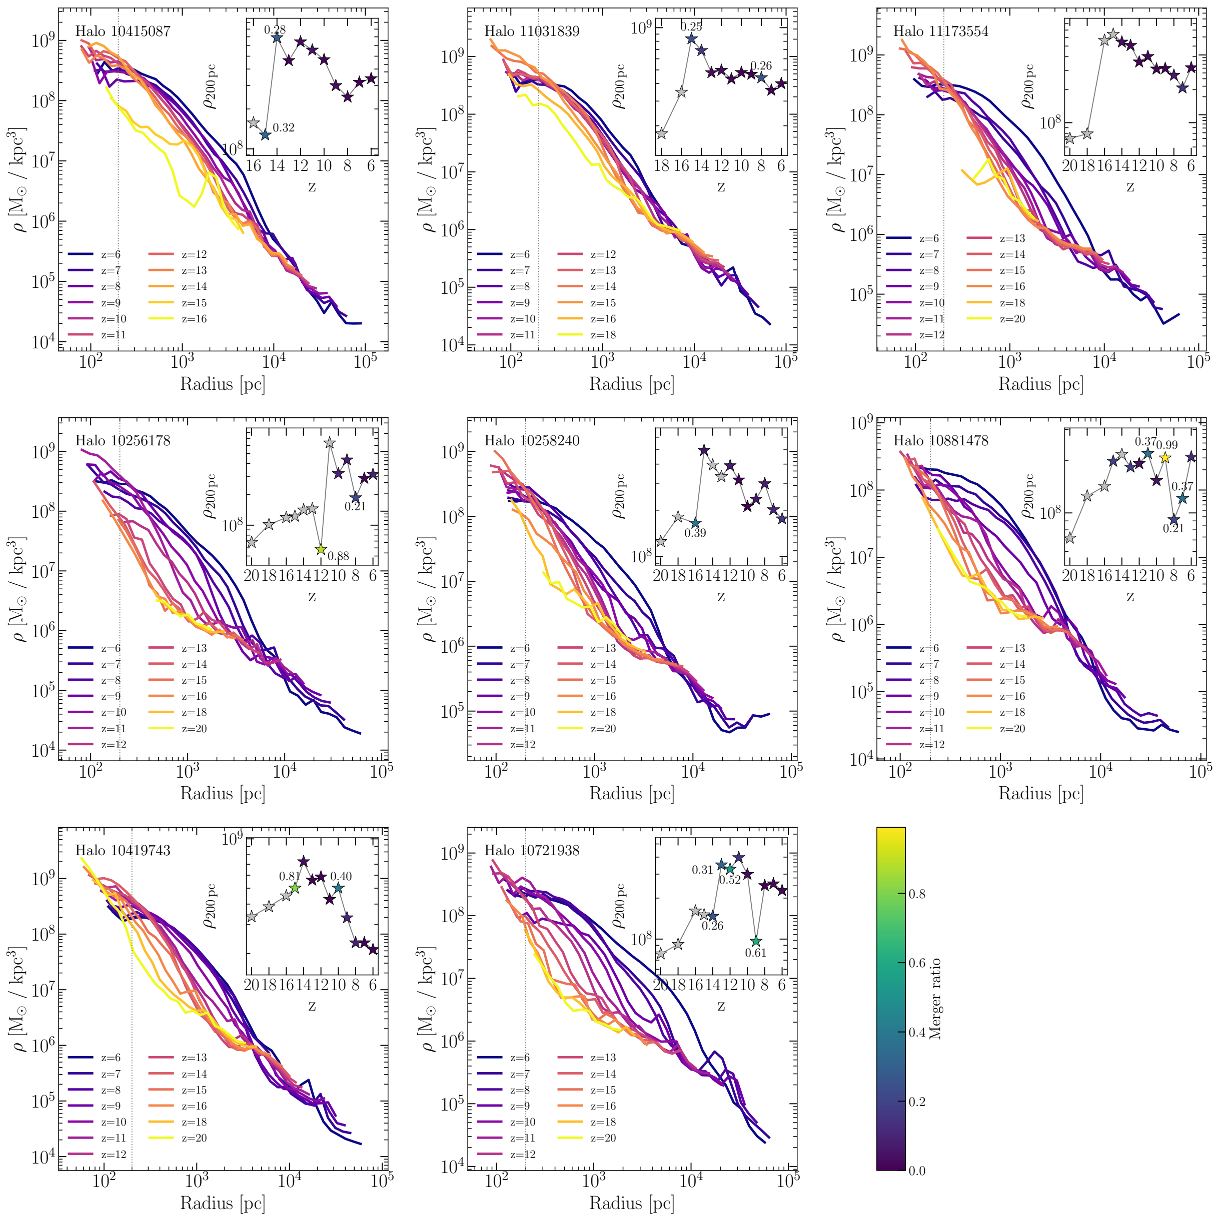

Figure 4.11: As Figure 4.10, but for the s147v24 halos matched to the same CDM gravothermal collapse candidates. Panel order matches Figure 4.10 for the seven candidates with a counterpart in both models; the eighth panel (Halo 10721938) is the one CDM candidate matched only in s147v24, not s70v0.

Chapter 5

Measuring Energy and Deleptonization of the Next Galactic Supernova without Spectral Assumptions

5.1 Abstract

The next Milky Way core-collapse supernova will deliver more than ten thousand neutrino events across a global network of detectors — providing the first opportunity to characterize the full-flavor emission and probe the newly formed proto-neutron star (PNS). In particular, the total emitted energy $E_{\rm tot}$ and the net deleptonization number $N_{\rm delep} \equiv N_{\nu_e} - N_{\bar{\nu}_e}$ are the two leading-order observables of the explosion. Measuring them precisely is the es-

This chapter is an adaptation of work to be submitted to *Physical Review D* later this year. The coauthors on this work are Shirley W. Li, John Beacom, Luke F. Roberts, and Francesco Capozzi.

sential first step toward a comprehensive characterization of the event. Here, we present a model-independent measurement of them from the flavor-resolved energy spectra at Super-Kamiokande, DUNE, and JUNO, using B-spline spectral unfolding without assuming a parametric spectral shape. *We show that $E_{\rm tot}$ scales robustly with PNS mass, and with sufficient precision to make it a direct observable for measuring the remnant mass.* We also show that $N_{\rm delep} > 0$ can be robustly established, directly confirming the neutronization of the stellar core. We further demonstrate that DUNE plays an irreplaceable role: without a dedicated ν_e anchor, $E_{\rm tot}$ reconstruction degrades by an order of magnitude, and a non-zero $N_{\rm delep}$ cannot be established. Our results demonstrate the power of model-independent, full-flavor spectral unfolding as a framework for extracting the physics of the next Galactic supernova.

5.2 Introduction

Core-collapse supernovae release $\simeq 3 \times 10^{53}$ erg of gravitational binding energy, nearly all of which is radiated in neutrinos over $\sim$1 min (Janka, 2012; Burrows & Vartanyan, 2021; Li et al., 2021). The next Milky Way core-collapse supernova will deliver more than 10^4 neutrino events across a global network of detectors — orders of magnitude beyond the 19 events from SN 1987A (Hirata et al., 1987, 1988; Bionta et al., 1987; Bratton et al., 1988) — providing an unprecedented opportunity to characterize the full-flavor neutrino emission (Scholberg, 2012; Mirizzi et al., 2016; Adams et al., 2013). Because nearby supernovae are rare (Diehl et al., 2006; Rozwadowska et al., 2021), we likely have just one chance over the next few decades to get this right.

The two most important quantities characterizing the burst are the total emitted energy $E_{\rm tot}$ and the net deleptonization number $N_{\rm delep} \equiv N_{\nu_e} - N_{\bar{\nu}_e}$. These are the leading-order observables of the explosion: $E_{\rm tot}$ quantifies the total energy release of the collapse, and $N_{\rm delep}$ measures the net neutronization of the collapsing core. Both quantities also connect directly to the properties of the newly formed proto-neutron star (PNS): $E_{\rm tot}$ is set by the

PNS gravitational binding energy, and N_{delep} reflects the degree of neutronization required to produce the remnant baryonic mass.

The standard approach to inferring these quantities from detector data is to fit observed event spectra to a parametric spectral template, typically a two- or three-parameter pinched Fermi-Dirac shape. While tractable, this imposes a rigid spectral form that is not guaranteed by realistic supernova simulations (Fiorillo et al., 2023), and the resulting model bias can propagate into systematic errors on E_{tot} and N_{delep} at the level of tens of percent.

In this chapter, we instead adopt a model-independent approach based on B-spline spectral unfolding. The detector response matrix is inverted from the observed event counts to recover the true neutrino energy spectrum without assuming any parametric shape; integrating the recovered spectrum directly yields E_{tot} and N_{delep}. We apply this to three complementary channels that together span the full neutrino flavor space: Super-Kamiokande detects $\bar{\nu}_e$ via inverse beta decay ($\simeq$6500 events at 10 kpc); DUNE detects ν_e via charged-current absorption on argon ($\simeq$4000 events); and JUNO constrains ν_x via elastic scattering (Abusleme et al., 2025b).

DUNE is indispensable to this decomposition. Without a dedicated ν_e anchor, the ν_e and ν_x contributions to the JUNO signal cannot be independently resolved, E_{tot} reconstruction degrades by an order of magnitude, and a non-zero N_{delep} cannot be established.

This chapter is organized as follows. Section 5.3 reviews the physics of supernova neutrino emission and describes the simulation inputs used in this work. Section 5.4 describes the detection channels and detector response models. Section 5.5 demonstrates that parametric spectral fitting fails systematically and motivates the model-independent approach. Section 5.6 discusses the unfolding problem and describes the B-spline framework. Section 5.7 presents the reconstruction results. Section 5.8 connects these to PNS physics. We conclude in Section 5.9.

5.3 Supernova Neutrino Emission

Core-collapse supernovae are among the most remarkable events in the Universe. They mark the deaths of massive stars and the births of compact remnants — neutron stars, or in some cases black holes. At their core, these events are powered by a simple physical process: the gravitational collapse of the stellar iron core releases an enormous amount of gravitational binding energy, nearly all of which is carried away by neutrinos and antineutrinos of all flavors over a $\simeq$1-min burst (Li et al., 2021). Simultaneously, the collapse converts a large fraction of the stellar protons into neutrons, releasing a net excess of ν_e over $\bar{\nu}_e$ that encodes the total deleptonization of the core. The total emitted energy $E_{\rm tot}$ and the net deleptonization number $N_{\rm delep} \equiv N_{\nu_e} - N_{\bar{\nu}_e}$ are therefore the two leading-order observables of the explosion and the primary targets of this work.

The collapse begins when the iron core of a massive star ($\mathrm{M} \gtrsim 8\,\mathrm{M}_\odot$) grows massive enough that electron degeneracy pressure can no longer support it against gravity (Janka, 2012; Burrows & Vartanyan, 2021). Once this maximum supportable mass (the effective Chandrasekhar mass) is exceeded, the core collapses nearly in free fall, reaching nuclear saturation density ($\rho_0 \simeq 2.7 \times 10^{14}$ g cm^{-3}, the density inside an atomic nucleus) within $\simeq$100 ms. At that point the nuclear equation of state stiffens dramatically and abruptly halts the collapse — the core bounces — sending a shock wave outward into the still-infalling outer layers of the star. As the shock propagates outward and breaks through the region where neutrinos are trapped (the neutrinosphere), it heats the overlying material and drives rapid neutronization, $e^- + p \to \nu_e + n$, producing a brief intense burst of ν_e. Despite its name, this neutronization burst carries only $\simeq$1–3% of the total emitted energy, but it already accounts for $\simeq$40% of the total lepton-number release (Li et al., 2021).

Over the following $\simeq$0.1–1 s, matter from the outer stellar layers continues to fall inward and accrete onto the compact object, depositing gravitational energy and sustaining high neutrino luminosities in all flavors. Meanwhile, the outgoing shock loses energy dissociating the infalling nuclei and stalls before it can reach the outer envelope. It is the neutrinos themselves — streaming outward from the hot compact core — that deposit enough energy behind the stalled shock to eventually revive it and drive the explosion (Bethe & Wilson, 1985). This accretion phase contributes a substantial fraction of the total energy $E_{\rm tot}$ and lepton number $N_{\rm delep}$; and ν_e and $\bar{\nu}_e$ are somewhat enhanced relative to ν_x because charged-current reactions couple electron-flavor neutrinos more strongly to the dense, neutron-rich matter.

Once neutrino heating revives the stalled shock and a successful explosion ensues, the compact remnant—now called a proto-neutron star (PNS)—settles into hydrostatic equilibrium with an initial radius of $\simeq$50 km. Over the following tens of seconds, it radiates its residual thermal and lepton-number content as it contracts toward its final radius of $\simeq$12 km, eventually becoming transparent to neutrinos and cooling into a cold neutron star (Li et al., 2021; Mirizzi et al., 2016).

Across all three phases, the three neutrino flavors are not emitted identically—their spectra and total numbers differ, reflecting their different couplings to the dense PNS matter. Broadly, though, all six flavors carry comparable luminosities and average energies; the flavor differences are a modest $\sim$10–20% effect. Within this, a clear hierarchy in average energy emerges from the coupling strengths. The ν_e interact via charged current primarily on the abundant neutrons in the neutron-rich PNS, making them the most tightly coupled; they therefore decouple at the lowest temperature and carry the lowest average energy. Their total number is also boosted above the thermal baseline by the net deleptonization excess. The $\bar{\nu}_e$ couple via charged current to the less abundant protons, decouple at a somewhat higher temperature, and carry a somewhat higher average energy. The heavy-lepton flavors ν_x interact only via neutral current, are the least coupled, and tend to carry the highest average energies — though the precise ordering relative to $\bar{\nu}_e$ can vary across models. Finally, all four ν_x species are nearly identical to one another: PNS temperatures are below the muon and tau mass thresholds, so all heavy-lepton flavors interact through the same neutral-current channels.

The time-integrated emission from all three phases together encodes these two observables quantitatively. The total emitted energy, $E_{\rm tot} \simeq (3/5)GM_{\rm PNS}^2/R_{\rm PNS} \simeq 3 \times 10^{53}$ erg, equals the gravitational binding energy of the PNS and hence scales directly with its mass and compactness. The net deleptonization number reflects the total neutronization of the core: since each electron captured to form a neutron produces one ν_e, we expect $N_{\rm delep} \simeq M_{\rm PNS}/2m_p \simeq 8 \times 10^{56}$ for a typical remnant, growing with $M_{\rm PNS}$. Measuring both precisely requires integrating the full neutrino signal across all flavors over the entire burst duration—which is why a multi-detector, full-flavor strategy is essential.

5.3.1 Supernova Simulations

Computing the neutrino emission from a core-collapse supernova requires large-scale numerical simulations that solve neutrino radiation transport coupled self-consistently to hydrodynamics and gravity. These are among the most computationally demanding simulations in astrophysics; state-of-the-art three-dimensional models can require months of runtime on modern supercomputers (Janka, 2012; Burrows & Vartanyan, 2021). Different simulation groups working with independent codes and methods largely agree on the broad picture — total energetics, overall emission timescales, and the rough flavor hierarchy — while differing in details such as the explosion timing, the luminosity evolution during accretion, and the precise spectral shapes. These differences reflect genuine uncertainties in the input physics, including the nuclear equation of state and neutrino opacities in dense matter, which we account for by testing our method across a range of models.

Since our unfolding method makes no assumption about the spectral shape of the neutrino flux, our results should in principle be robust against the choice of input model. Nonetheless, we need to adopt a nominal flux to demonstrate the method concretely; we then verify robustness by testing across a range of models. For our nominal input, we use the one-dimensional (1D) supernova simulations of Fiorillo et al. (2023). Although multidimensional simulations capture important effects such as convection and asymmetric explosion geometries, they are too computationally expensive to evolve for the full tens of seconds needed to capture the complete neutrino burst; 1D simulations can be run through the entire cooling phase, making them the appropriate choice for computing time-integrated quantities such as $E_{\rm tot}$ and $N_{\rm delep}$. As our fiducial model we take the 1.93 $M_\odot$ remnant with the SFHx equation of state (Steiner et al., 2013; Hempel & Schaffner-Bielich, 2010).

Importantly, the nominal model we use, and also most simulation codes, do not directly output neutrino energy spectra; rather, they provide time-resolved luminosities $L_\nu(t)$, mean energies $\langle E\rangle(t)$, and mean-squared energies $\langle E^2\rangle(t)$ for each flavor, from which a spectral shape must be inferred. To construct the flux at each time step, we fit a pinched Fermi-Dirac distribution to the provided $\langle E\rangle(t)$ and $\langle E^2\rangle(t)$, with the pinch parameter $\alpha(t)$ derived from the ratio $\langle E^2\rangle/\langle E\rangle^2$. We then time-integrate these instantaneous spectra over the full burst to obtain the total flavor-decomposed flux dN/dE_ν. Crucially, because $\langle E\rangle(t)$ and $\alpha(t)$ both evolve throughout the burst, the resulting time-integrated spectrum is *not* a single pinched Fermi-Dirac — it has a more complex shape that no simple parametric form captures accurately, which is precisely why a model-independent unfolding approach is needed.

One further point concerns the treatment of the heavy-lepton neutrino flavors. The four flavors: ν_μ, $\bar{\nu}_\mu$, ν_τ, and $\bar{\nu}_\tau$ are collectively denoted ν_x. In detail, their emission spectra can differ slightly from one another; for example, ν_τ and $\bar{\nu}_\tau$ experience a slightly modified effective potential inside the star due to the larger tau lepton mass. However, no terrestrial detector can distinguish among these four flavors: all ν_x interact with detector material only via neutral-current reactions, and their cross sections with nucleons and electrons are nearly identical at supernova energies. We therefore treat all four as identical and work with their per-flavor average spectrum; the total heavy-lepton contribution is then 4 times this per-flavor flux.

Figure 5.1 shows the time-integrated flux for the fiducial model. The three neutrino flavors have average energies $\langle E_{\nu_e} \rangle = 11.15$ MeV, $\langle E_{\bar{\nu}_e} \rangle = 13.81$ MeV, and $\langle E_{\nu_x} \rangle = 13.06$ MeV.

The fiducial model shown in Fig. 5.1 is of course just one realization; to assess how sensitive our results are to this choice, we apply our unfolding procedure to all 20 models in the same simulation suite: four equations of state — DD2, LS220 (Lattimer & Swesty, 1991), SFHo, and SFHx (Steiner et al., 2013; Hempel & Schaffner-Bielich, 2010) —crossed with five final baryonic PNS masses of 1.36, 1.44, 1.62, 1.77, and 1.93 $M_\odot$.

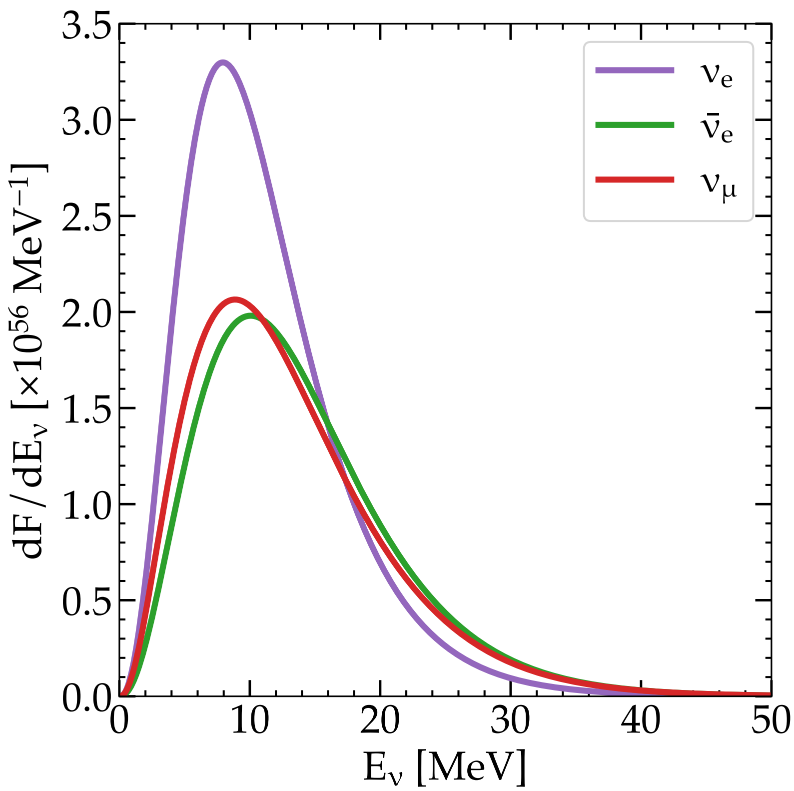

Figure 5.1: Time-integrated neutrino flux for the fiducial model (1.93 $M_\odot$ PNS remnant, SFHx equation of state). Shown are the electron neutrino (purple), electron antineutrino (green), and per-flavor ν_x (red).

5.3.2 Neutrino Oscillations

A further complication is that supernova neutrinos can change flavor as they escape the star. In a supernova, the oscillation problem is far more complex than in vacuum. Dense stellar matter modifies the effective neutrino masses and can drive resonant flavor conversion (the Mikheyev-Smirnov-Wolfenstein, or MSW, effect). At the extreme neutrino densities near the PNS, neutrino-neutrino interactions can drive collective oscillations, and recent work has highlighted the possibility of fast flavor conversion that could occur on timescales of nanoseconds very close to the neutrinosphere. The net effect of all these processes on the time-integrated flavor spectra reaching Earth remains genuinely uncertain — this is one of the central unsolved problems in supernova neutrino physics (Mirizzi et al., 2016).

In this work, we make the simplifying assumption that no oscillations occur: the neutrino fluxes from simulation are used directly as inputs to the detector response, without any flavor transformation. This is the standard convention in the field and provides a clean, well-defined benchmark. In practice, different oscillation scenarios would modify the flavor decomposition of the detected signal; interpreting future measurements in terms of specific oscillation scenarios is an important but separate problem, which we leave to future work.

5.4 Supernova Neutrino Detection

In this section, we describe the four detection channels used in this analysis and compute the detected spectrum for our nominal flux model for each. Together, SuperK via inverse beta decay (IBD), DUNE via charged-current absorption on argon, and JUNO via both neutrino–electron and neutrino–proton elastic scattering cover all three neutrino flavors. We adopt a fiducial source distance of 10 kpc throughout, a representative value for a Galactic core-collapse supernova (Mirizzi et al., 2016)

5.4.1 Super-Kamiokande - Inverse Beta Decay

The best detection channel for a single flavor is inverse beta decay (IBD) interactions for $\bar{\nu}_e$. The interaction is through free protons in the detector,

$$\bar{\nu}_e + p \to e^+ + n\,. \tag{5.1}$$

Among the experiments we consider, Super-K has been the longest running with this channel available. In addition, the newly operating JUNO and the future Hyper-Kamiokande detector both have this channel available. In this chapter, we use Super-K as the fiducial case, demonstrating that it already achieves excellent reconstruction performance, and then add brief discussions on how results differ for JUNO and Hyper-K.

Super-K is a water Cherenkov detector that has been operating for over two decades. Because the CCSN burst lasts at most $\sim$1 minute (Li et al., 2021), non-CCSN backgrounds are negligible, and in principle a larger detection volume of $\sim$32 kt could be adopted; here, to be conservative, we use the fiducial mass of 22.5 kt (Suzuki, 2019), corresponding to $N_p \simeq 1.5 \times 10^{33}$ free proton targets. The recently loaded gadolinium has a high neutron capture rate and releases a multi-MeV gamma cascade, providing a delayed coincidence signal that distinguishes IBD from other CCSN channels — such as neutrino–electron scattering, which produces only a recoil electron with no neutron tag. In this work, we assume IBD in Super-K can be cleanly separated from other channels with no background contamination.

The IBD cross section is well known with little theoretical uncertainty. We use the analytic prescription from Strumia & Vissani (2003):

$$\begin{aligned}\sigma &\simeq 10^{-43}\ \mathrm{cm}^2\, p_e E_e \\ &\times\ E_\nu^{-0.07056+0.02018\ln E_\nu - 0.001953\ln^3 E_\nu}\,,\end{aligned} \tag{5.2}$$

where E_ν is the incident antineutrino energy in MeV. $E_e = E_\nu - \Delta$ is the outgoing positron total energy, where $\Delta = m_n - m_p \simeq 1.293$ MeV, and $p_e = \sqrt{E_e^2 - m_e^2}$ is the positron momentum. This parameterization reproduces the exact cross section to sub-percent accuracy for neutrino energies $E_\nu \lesssim 300$ MeV, well above the supernova neutrino energy range considered here.

Once the positron is produced, it loses energy by ionization while emitting Cherenkov light, comes to rest, and annihilates into two 0.511 MeV gamma rays. The total detected energy is therefore $E_{\rm det} = E_{e^+} + 0.511$ MeV. Detector energy resolution is modeled by Gaussian smearing, using the parametrization from the Super-K IV solar neutrino analysis, $\sigma_E(E) = -0.05525 + 0.3162\sqrt{E} + 0.04572E$ (Abe et al., 2024), with a detection threshold of 3.5 MeV.

Figure 5.2 shows the time-integrated IBD spectrum for our fiducial SN model. The total number of detected events above threshold is 6.9×10^3 with an average detected energy of 20.8 MeV.

The calculation from true neutrino spectrum to detected spectrum is naturally framed as a matrix multiplication. The observed event counts $n_j^{\rm obs}$ in detected-energy bin j are related to the true neutrino flux $\phi_i \equiv \phi(E_i)$ in true-energy bin i by

$$n_j^{\rm exp} = \sum_i M_{ji}\,\phi_i\,\Delta E_i\,, \tag{5.3}$$

where M_{ji}, the response matrix, gives the expected counts in detected bin j per unit flux in true bin i, encoding both the interaction cross section and detector energy resolution. This relation is the starting point for the spectral unfolding methods described in Sec. 5.6. The left panel of Fig. 5.3 shows this matrix for Super-K IBD. The matrix takes the form of a narrow diagonal band: IBD kinematics are nearly one-to-one, $E_{\rm det} \simeq E_\nu - 0.782$ MeV, so without smearing the matrix would reduce to a single line whose brightness grows with energy (reflecting the rising cross section); Gaussian smearing broadens this into the band visible in the figure.

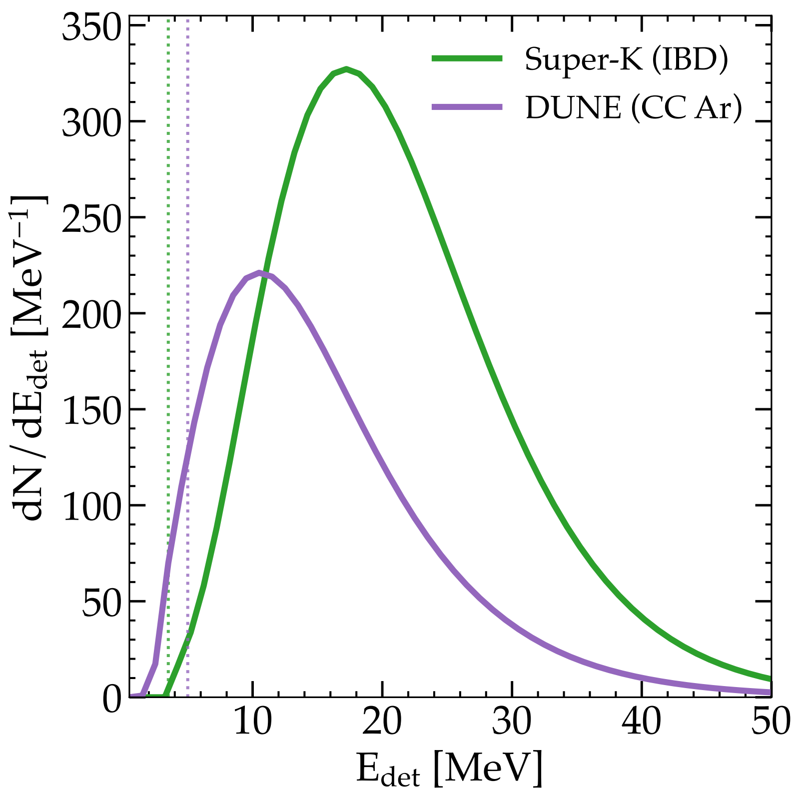

Figure 5.2: Time-integrated detected neutrino spectra at SuperK (inverse beta decay; green) and DUNE (charged-current ν_e on argon; purple) as a function of detected energy E_{det} for the fiducial core-collapse supernova model. Vertical dotted lines indicate the corresponding detector energy thresholds. Source distance 10 kpc.

As shown in Fig. 5.4, the 3.5 MeV detector threshold corresponds to a source threshold of $E_\nu^{\text{min}} \simeq 4.3$ MeV. Because the $\bar{\nu}_e$ spectrum peaks well above this value, the threshold penalty is negligible: only 0.6% of $N_{\bar{\nu}_e}$ and 0.4% of $E_{\bar{\nu}_e}^{\text{tot}}$ fall below threshold for our fiducial model. IBD at Super-K is therefore sensitive to essentially the full $\bar{\nu}_e$ emission.

The IBD channel is also available at JUNO and Hyper-Kamiokande. JUNO has a comparable fiducial mass, so the IBD event count is similar; JUNO's advantage is better energy resolution and a lower threshold, which in principle allows a more precise unfolding. Hyper-K will have a much larger fiducial volume — roughly 8 times that of Super-K — providing a substantial boost in statistics. Its low-energy performance is still being characterized, with

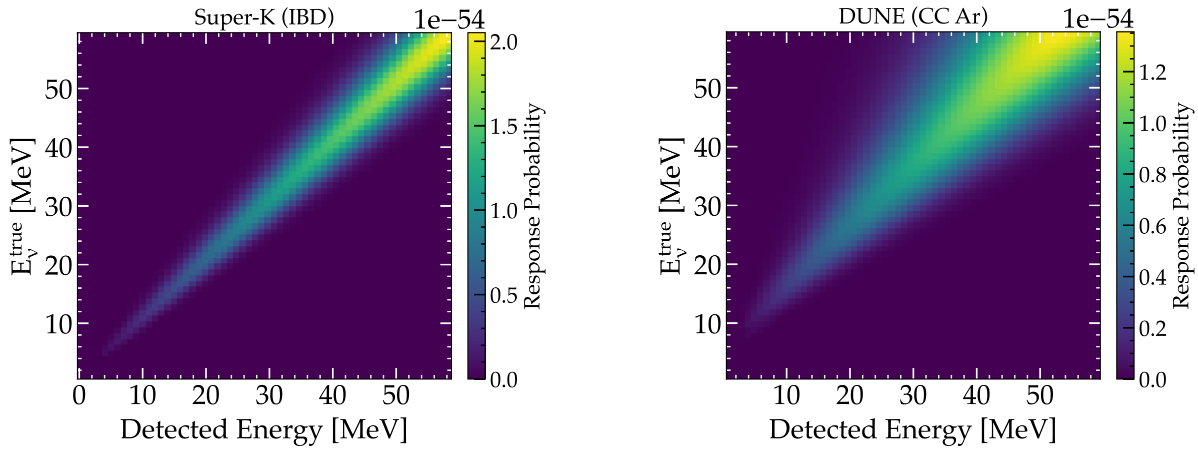

Figure 5.3: Left: Response matrix for the Super-Kamiokande inverse beta decay channel. Each column shows the probability distribution of reconstructed positron energies for a given true $\bar{\nu}_e$ energy. The narrow diagonal band reflects the excellent energy resolution of the water Cherenkov detector; the matrix is empty below the 3.5 MeV detection threshold. Right: Response matrix for the DUNE charged-current channel ($\nu_e + {}^{40}\mathrm{Ar} \to e^- + {}^{40}\mathrm{K}^*$). The broader diagonal band reflects the 20% energy resolution; the matrix is empty below the 3 MeV detection threshold.

projections suggesting energy resolution comparable to Super-K, possibly with a somewhat higher threshold. When the next galactic supernova occurs, combining IBD data from all three detectors will give the best possible constraint on the $\bar{\nu}_e$ spectrum. In practice, we find that Super-K alone can already achieve excellent $\bar{\nu}_e$ reconstruction, as we show in Sec. 5.7. .

5.4.2 DUNE – ν_e Charged-Current Absorption on Argon

The best detection channel for the ν_e flavor is charged-current (CC) absorption on argon nuclei,

$$\nu_e + {}^{40}\mathrm{Ar} \to e^- + {}^{40}\mathrm{K}^* \,. \tag{5.4}$$

This channel is uniquely available at the Deep Underground Neutrino Experiment (DUNE) (Abi et al., 2021, 2020).

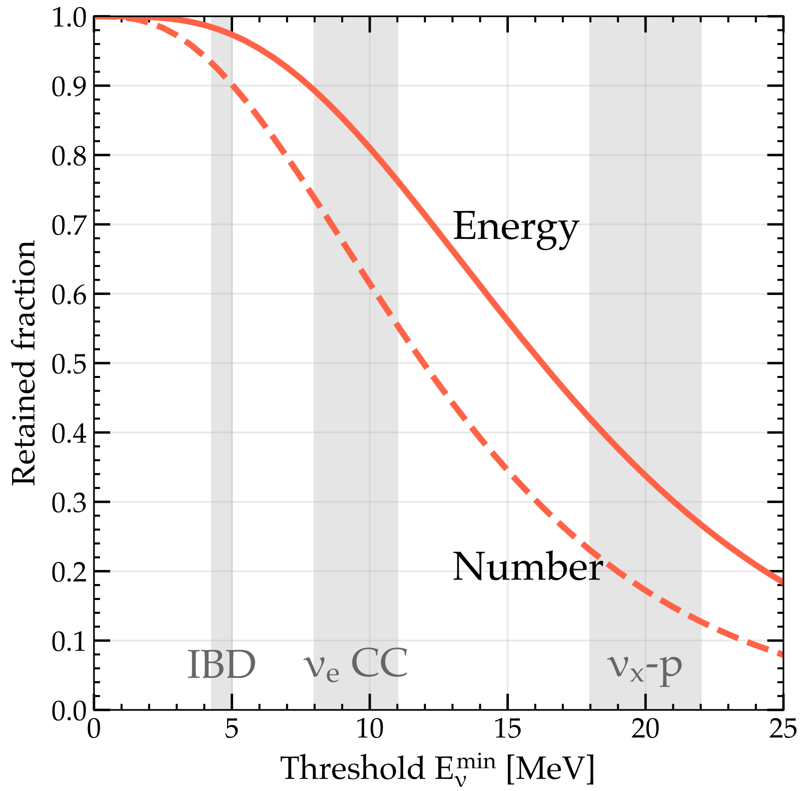

Figure 5.4: Fraction of the ν_x total energy (solid) and total number (dashed) retained above a given threshold $E_\nu^{\rm min}$, for the fiducial supernova model. Shaded bands mark the effective threshold ranges of each detection channel: IBD at Super-K (4.3–5 MeV), ν_e charged-current at DUNE (7–11 MeV), and ν_x–p elastic scattering at JUNO (18–22 MeV). The threshold penalty is negligible for IBD, moderate for ν_e CC, and severe for ν_x–p.

DUNE is a liquid argon time-projection chamber (LArTPC) currently under construction at the Sanford Underground Research Facility (SURF). Its far detector will consist of four 10-kt modules, for a total of 40 kt of liquid argon, which we adopt as the fiducial mass. Charged particles traversing the liquid argon produce ionization charge that drifts under an applied electric field and is collected at wire readout planes, from which particles are reconstructed.

The ν_e CC cross section on ^{40}Ar has sizable uncertainties. Roughly speaking, one can think of the interaction as a sum over transitions from the ^{40}Ar ground state to 15 discrete excited states of ^{40}K*. We adopt the formula from Bhattacharya et al. (2009),

$$\sigma(E_\nu) = \sum_i \frac{G_F^2 |V_{ud}|^2}{\pi} B_i \, E_e^i \, p_e^i \, F(Z, E_e^i) \,, \tag{5.5}$$

where G_F is the Fermi constant, V_{ud} is the relevant CKM matrix element, and $B_i = B_i(F) + B_i(\mathrm{GT})$ is the summed Fermi and Gamow–Teller reduced transition probability for transition i (Konopinski & Uhlenbeck, 1935). For each transition i, E_e^i and p_e^i are the electron total energy and momentum, and $F(Z, E_e^i)$ is the Fermi function accounting for Coulomb distortion of the outgoing electron wavefunction in the field of the daughter nucleus ($Z = 19$ for ^{40}K).

For each nuclear transition i, the outgoing electron total energy is $E_e^i = E_\nu - Q - \Delta E_i$, where $Q = 1.504$ MeV is the ground-state Q-value and ΔE_i is the excitation energy of the i-th state in ^{40}K*, ranging from 2.3 to 8.0 MeV across the 15 transitions. We model the detector energy resolution as Gaussian smearing with $\sigma(E)/E = 20\%$ and apply a detection threshold of $E_{\mathrm{det}} \geq 3$ MeV. Figure 5.2 shows the time-integrated detected spectrum for our fiducial supernova model. The total detected event count above threshold is 4.0×10^3 with an average detected energy of 15.7 MeV.

The right panel of Fig. 5.3 shows the response matrix for this channel. As with the IBD matrix, the matrix takes the form of a diagonal band, but broader than the IBD case, reflecting the 20% Gaussian energy resolution. The discrete transition structure from the 15 nuclear levels is smeared away. The matrix is empty below the 3 MeV detection threshold.

The detection threshold has a more noticeable impact for this channel than for IBD. As shown in Fig. 5.4, the 3 MeV detection threshold on the electron corresponds to an effective source ν_e threshold of approximately 7 MeV, set by the lowest-energy nuclear transition ($\Delta E_0 = 2.3$ MeV combined with the ground-state Q-value gives $E_\nu^{\min} = Q + \Delta E_0 + E_{\mathrm{det}}^{\min} \simeq$ 6.8 MeV). The threshold penalty is moderate but non-negligible: roughly 27% of N_{ν_e} and

12% of $E_{\nu_e}^{\rm tot}$ lie below this threshold for our fiducial model. This asymmetry between the count and energy fractions reflects the energy weighting of the two integrals: the count integral $\int (dN/dE_\nu)\, dE_\nu$ treats all energies equally, so the sub-threshold region contributes its full 27% to N_{ν_e}; the energy integral $\int E_\nu\, (dN/dE_\nu)\, dE_\nu$ downweights low energies by the factor E_ν, reducing the sub-threshold contribution to 12%. As a result, N_{ν_e} is intrinsically more sensitive to sub-threshold extrapolation errors than E_{ν_e}, and reconstruction of the total count is correspondingly harder. This contrasts with the IBD case, where losses are sub-percent, and motivates careful treatment of the low-energy portion of the ν_e spectrum in the unfolding (Sec. 5.7).

The ν_e CC channel on argon is unique to DUNE among the detectors we consider. This makes DUNE irreplaceable: without it, the ν_e and ν_x contributions to the JUNO and Hyper-K signals cannot be independently disentangled, and neither $E_{\rm tot}$ nor a non-zero $N_{\rm delep}$ can be established. We demonstrate this explicitly through an ablation study in Sec. 5.7.

5.4.3 JUNO – Neutrino-Electron Elastic Scattering

In this and the following subsection, we describe the channels used to constrain ν_x, the hardest flavor to detect. Neutrino-electron elastic scattering,

$$\nu + e^- \to \nu + e^- \,, \tag{5.6}$$

proceeds for all neutrino flavors. This makes it a powerful channel for constraining ν_x, but requires careful treatment of the ν_e and $\bar{\nu}_e$ contributions, which are non-negligible. Since electron targets are present in all detectors we consider, this channel is available at every detector in our analysis. We take JUNO to be our nominal case because of its superb energy resolution and low threshold. In the joint fit, we also include the Hyper-K channel for its large statistics; we compare its contribution in Sec. 5.7.

JUNO is a 20 kt liquid scintillator detector currently operating in Jiangmen, Guangdong, China (An et al., 2016; Abusleme et al., 2025a). The scintillator is based on linear alkylbenzene (LAB), providing $N_e \simeq 6.9 \times 10^{33}$ electron targets and $N_p \simeq 1.5 \times 10^{33}$ free proton targets. JUNO achieves an exceptional energy resolution of $\sigma_E/E = 3\%/\sqrt{E/\mathrm{MeV}}$ and a detection threshold of $E_{\rm det} \geq 0.2$ MeV.

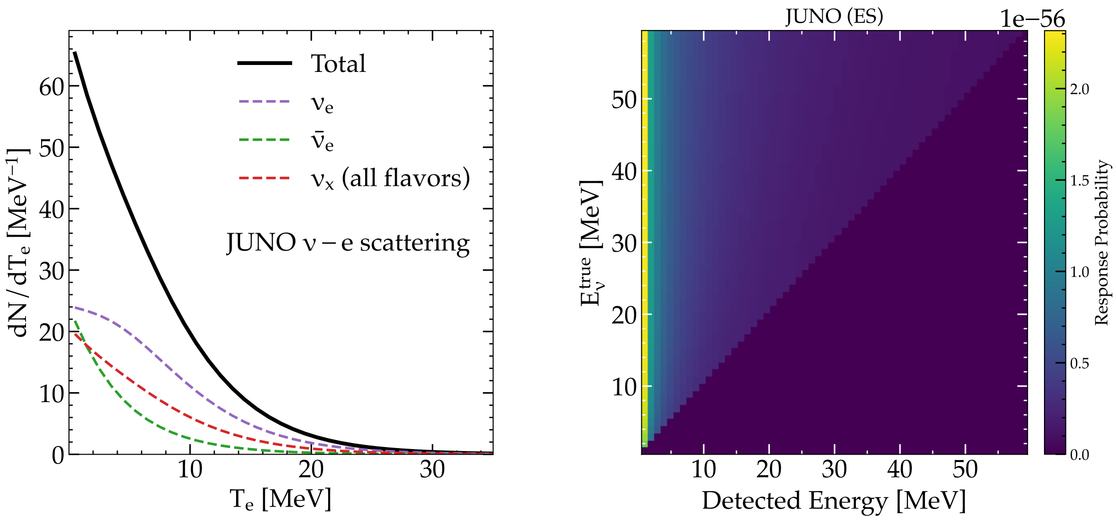

Figure 5.5: Left: Time-integrated JUNO neutrino–electron elastic scattering detected spectrum, broken down by flavor contribution, for the fiducial supernova model at 10 kpc. Because ν_e scatters via both charged- and neutral-current interactions while ν_x scatters via neutral current only, the flavor contributions are unequal and must be disentangled using the dedicated ν_e and $\bar{\nu}_e$ measurements from DUNE and SuperK. Right: Response matrix for the JUNO neutrino–electron elastic scattering channel. The triangular structure reflects the kinematics of elastic scattering: for a given E_ν, the recoil electron energy ranges from zero to a maximum $T_e^{\max}(E_\nu)$, producing a broadening toward lower detected energies. The very low threshold of 0.2 MeV allows this channel to probe the full low-energy portion of the neutrino spectrum.

The differential cross section is well known. We use the standard formula (Marciano & Parsa, 2003),

$$\frac{d\sigma}{dT_e} = \frac{2G_F^2 m_e}{\pi}\left[g_1^2 + g_2^2\left(1-\frac{T_e}{E_\nu}\right)^2 - g_1 g_2 \frac{m_e T_e}{E_\nu^2}\right], \tag{5.7}$$

where T_e is the electron recoil kinetic energy, ranging from zero to a kinematic maximum $T_e^{\max} = 2E_\nu^2/(m_e + 2E_\nu)$. The couplings g_1 and g_2 depend on neutrino flavor. For ν_e, the charged-current amplitude adds to the neutral-current, giving $g_1^{\nu_e} = \frac{1}{2} + \sin^2\theta_W \simeq 0.73$ and $g_2^{\nu_e} = \sin^2\theta_W \simeq 0.23$. For ν_x (neutral-current only), $g_1^{\nu_x} = -\frac{1}{2} + \sin^2\theta_W \simeq -0.27$ and $g_2^{\nu_x} = \sin^2\theta_W \simeq 0.23$. For antineutrinos, g_1 and g_2 are swapped relative to the corresponding neutrino.

Figure 5.5 left panel shows the event rate in this channel. For our fiducial supernova model, the total detected events above threshold is 5.0×10^2, with an average detected energy of 6.3 MeV. The flavor breakdown is unequal: ν_e contributes $\simeq 244$ events (50%), ν_x (all four heavy-lepton flavors combined) contributes $\simeq 143$ events (30%), and $\bar{\nu}_e$ contributes $\simeq 97$ events (20%). The dominance of ν_e reflects its enhanced cross section from the CC amplitude: $|g_1^{\nu_e}| \simeq 0.73$ versus $|g_1^{\nu_x}| \simeq 0.27$.

Figure 5.5 right panel shows the response matrix for this channel. In contrast to the narrow diagonal band of IBD or DUNE CC, the ν-e matrix is lower-triangular: for a given E_ν, the recoil energy can range from zero to $T_e^{\max}(E_\nu)$, spreading the signal over a broad range of detected energies. This partial energy transfer makes the matrix less diagonal and the unfolding more challenging than for the quasi-one-to-one IBD kinematics.

With a 0.2 MeV threshold, the threshold penalty for this channel is negligible: essentially the full ν_x source spectrum is accessible, comparable to the IBD case for $\bar{\nu}_e$.

In summary, this channel in JUNO is statistically limited ($\simeq$140 ν_x events), has degraded neutrino energy resolution from the differential scattering kinematics, and carries substantial contamination from ν_e and $\bar{\nu}_e$ — though the very low threshold is an asset. The Hyper-K ν-e sample provides complementary coverage: its statistics are about 8 times larger, but it has a threshold of a few MeV. The ν-p channel described below also provides complementary information that partially compensates these limitations.

5.4.4 JUNO – Neutrino-Proton Elastic Scattering

In JUNO, all six neutrino flavors can scatter elastically off the free proton targets in the detector via the neutral-current interaction,

$$\nu + p \to \nu + p\,. \tag{5.8}$$

The outgoing protons can be detected via their scintillation emission. Being a purely neutral-current process, this channel is flavor-universal: all six neutrino species contribute with the same cross section, making it a complementary probe to the flavor-specific channels above.

The differential cross section in terms of the proton recoil kinetic energy T is (Dasgupta & Beacom, 2011)

$$\begin{aligned}\frac{d\sigma}{dT} = \frac{G_F^2 M_p}{2\pi}\Bigg[&(g_V+g_A)^2 + (g_V-g_A)^2\left(1-\frac{T}{E_\nu}\right)^2 \\ &- (g_V^2 - g_A^2)\frac{M_p T}{E_\nu^2}\Bigg], \quad 0 \le T \le T_{\max},\end{aligned} \tag{5.9}$$

where $g_V = \frac{1}{2} - 2\sin^2\theta_W \approx 0.02$ and $g_A \approx 0.63$ are the proton vector and axial-vector weak couplings, and $T_{\max} = 2E_\nu^2/(2E_\nu + M_p)$ is the kinematic maximum recoil energy . The vector coupling nearly vanishes due to cancellations between up- and down-quark contributions, so the cross section is dominated by the axial-vector response. The kinematic maximum is small: at $E_\nu = 30$ MeV, $T_{\max} \simeq 1.8$ MeV.

A key complication of this channel is proton quenching. Nonrelativistic recoil protons have a high ionization density, which suppresses their scintillation light output relative to an electron of the same kinetic energy — an effect described by Birks' law,

$$T'(T) = \int_0^T \frac{dE}{1 + k_B\,\langle dE/dx\rangle}\,, \tag{5.10}$$

where T' is the electron-equivalent detected energy and k_B is Birks' constant (An et al., 2016).

In the JUNO liquid scintillator, the quenching is severe: the JUNO TDR parametrizes the relation as a power law $T' = A\,T^{\beta}$ with $A = 0.209$ and $\beta = 1.434$ (An et al., 2016), so that a proton of true kinetic energy T produces only about 20% of the equivalent electron light at $T \sim 1$ MeV, and less at lower energies. The analysis threshold of $T_p^{\rm det} \geq 0.2$ MeV on the quenched energy translates to a source neutrino energy threshold of $E_\nu \gtrsim 22$ MeV, since only neutrinos of that energy or greater can impart sufficient recoil to produce a detectable signal. After quenching, the detected proton energy $T_p^{\rm det}$ is further smeared by JUNO's energy resolution of $\sigma_E/E = 3\%/\sqrt{T_p^{\rm det}/\text{MeV}}$, and events below 0.2 MeV are excluded.

Figure 5.6 left panel shows the resulting detected spectra. Because the cross section is flavor-universal, the relative contribution of each species is set entirely by the fraction of its flux above the ∼22 MeV threshold. The ν_x component (four flavors summed) yields $\simeq 1.1 \times 10^3$ events and overwhelmingly dominates the total of $\simeq 1.5 \times 10^3$ events. Among the charged-lepton flavors, $\bar{\nu}_e$ contributes $\simeq 270$ events while ν_e contributes only $\simeq 130$, because the softer ν_e spectrum ($\langle E\rangle = 11.1$ MeV) has a smaller fraction of its flux above 22 MeV than the harder $\bar{\nu}_e$ spectrum ($\langle E\rangle = 13.8$ MeV).

Figure 5.6 right panel shows the response matrix for an averaged ν_x channel, i.e., averaging between ν_x and $\bar{\nu}_x$. The most striking feature is the extreme compression of the detected energy axis: the kinematic maximum recoil is $T_{\rm max} \lesssim 2$ MeV for $E_\nu \lesssim 30$ MeV, and after quenching the detected maximum $T_p^{\rm max}$ falls below 0.5 MeV at $E_\nu = 30$ MeV. As a result, the entire accessible neutrino energy range ($E_\nu \gtrsim 22$ MeV) maps into a narrow strip of $T_p^{\rm det} \sim 0.2$–1 MeV.

The threshold penalty for this channel is the most severe among the four, as shown by the ν_x-p band in Fig. 5.4. At the effective threshold of ∼22 MeV, only ≃13% of the total ν_x number and ≃27% of the total ν_x energy are accessible. This contrasts sharply with IBD and ν-e scattering, for which the penalty is sub-percent and negligible, respectively.

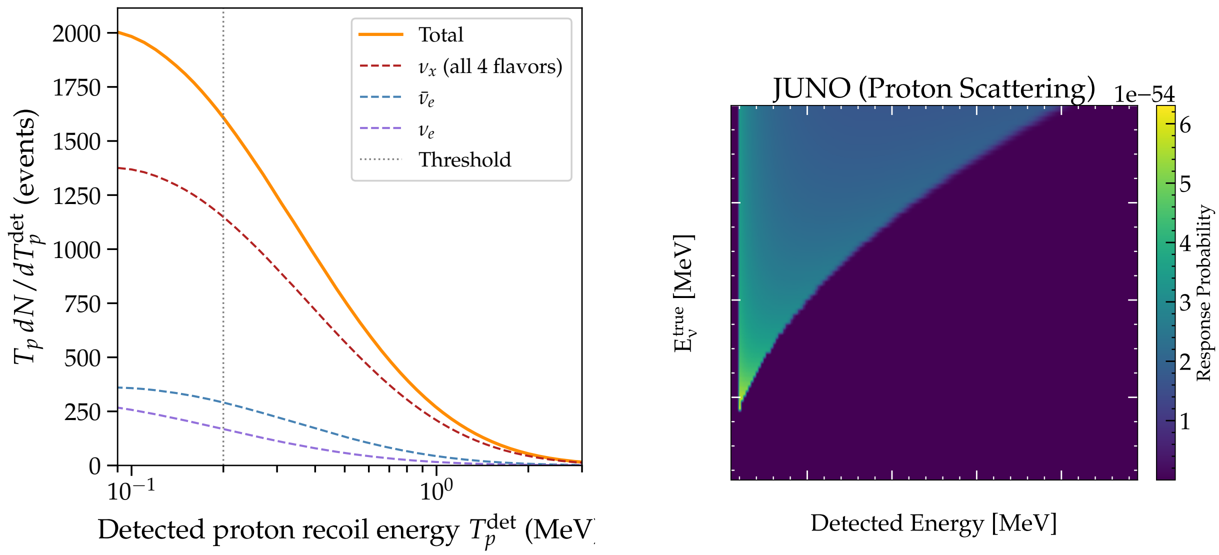

Figure 5.6: Left: Time-integrated JUNO neutrino–proton elastic scattering detected spectrum as a function of detected proton recoil kinetic energy T_p^{det}, broken down by flavor contribution, for the fiducial supernova model at 10 kpc. The vertical dotted line marks the detection threshold at 0.2 MeV. The narrow energy range results from strong quenching of the proton recoil in liquid scintillator. Because this channel is purely neutral-current, all six neutrino flavors contribute with equal cross sections. Right: Response matrix for the JUNO neutrino–proton elastic scattering channel. The heavy quenching compresses the entire accessible neutrino energy range ($E_\nu \gtrsim 22$ MeV) into a narrow strip of detected energies $T_p^{\mathrm{det}} \sim 0.2$–$1$ MeV.

In summary, the JUNO ν-p channel is uniquely sensitive to the high-energy tail of the ν_x spectrum above $\simeq$22 MeV. The ν_x contribution dominates for two reasons: all four heavy-lepton flavors contribute with equal cross sections, and the high spectral threshold preferentially selects hotter components of the source flux. Despite the severe threshold and narrow detected energy range, this channel provides an important spectral anchor at the high-energy tail of the ν_x spectrum that no other channel in our analysis reaches.

5.5 Limitations of Parametric Spectral Fitting

The standard approach to measuring E_{tot} and N_{delep} from supernova neutrino data is to fit the observed event spectrum to a parametric spectral template. The most common choice is the three-parameter *pinched* (α-spectrum) form (Keil et al., 2003; Tamborra et al., 2012),

$$\frac{dN}{dE} = \frac{N_{\rm tot}}{E_{\rm ave}} \frac{(\alpha+1)^{\alpha+1}}{\Gamma(\alpha+1)} \left(\frac{E}{E_{\rm ave}}\right)^{\alpha} e^{-(\alpha+1)E/E_{\rm ave}}, \tag{5.11}$$

with free parameters $N_{\rm tot}$ (total neutrino number), $E_{\rm ave}$ (mean energy), and α (pinching parameter). Fitting Eq. (5.11) per flavor to the detected event spectrum yields $E_{\rm tot} = N_{\rm tot}\, E_{\rm ave}$ summed over all flavors, from which $E_{\rm tot}$ and $N_{\rm delep}$ are extracted.

The pinched form is physically motivated for *instantaneous* neutrino emission, which is well approximated as a thermally distributed flux from a surface at a roughly uniform temperature (Janka & Hillebrandt, 1989). The question is whether it remains adequate for the *time-integrated* spectrum that a detector actually records. The time-integrated flux is the superposition of emission from all phases of the burst — neutronization spike, accretion, and Kelvin-Helmholtz cooling — during which the temperature, luminosity, and spectral shape all evolve substantially. The integral of a sequence of pinched spectra need not itself be a pinched spectrum.

Figure 5.7 addresses this directly. It shows the time-integrated ν_e, $\bar{\nu}_e$, and per-flavor ν_x spectra from our reference Garching model ($1.93\,\mathrm{M}_\odot$ + SFHx EOS), together with their best-fit pinched spectra. The fits are visually quite good: the pinched curves closely track the true spectra, and no obvious systematic shape discrepancy is apparent. One might conclude that the time-integrated spectrum is, for practical purposes, well described by the pinched form.

This conclusion is misleading. Figure 5.8 shows what happens when a nine-parameter pinched-spectrum fit — Eq. (5.11) per flavor, fitted jointly to all five channels via the same Poisson NLL minimization used throughout this chapter — is used to reconstruct $E_{\rm tot}$ and $N_{\rm delep}$ from 10^3 Poisson toy realizations. The joint distribution of $E_{\rm tot}/E_{\rm tot}^{\rm true}$ and $N_{\rm delep}/N_{\rm delep}^{\rm true}$ is centered near (0.97, 1.96): $E_{\rm tot}$ is modestly underestimated while $N_{\rm delep}$ is overestimated

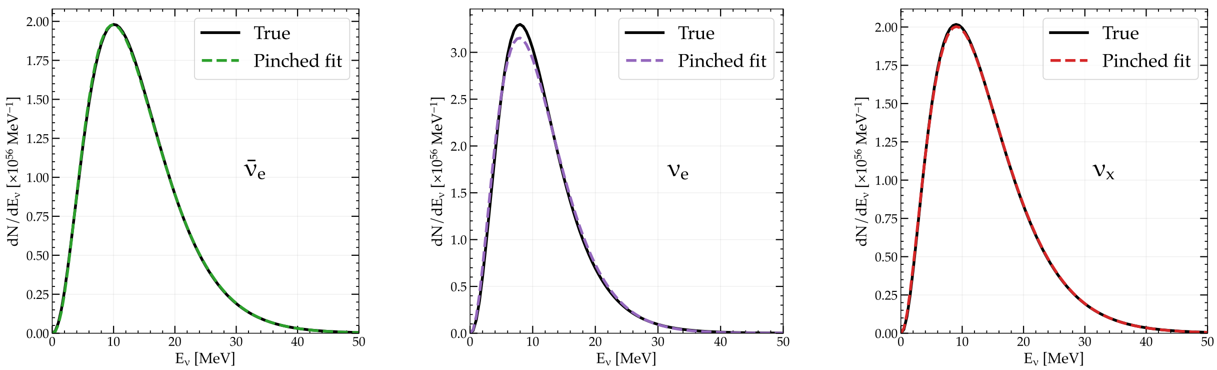

Figure 5.7: Time-integrated $\bar{\nu}_e$ (left), ν_e (center), and per-flavor ν_x (right) spectra from the reference Garching model ($1.93\,\mathrm{M}_\odot$ + SFHx EOS), compared to their best-fit pinched (α-spectrum) shapes (colored dashed). The visual agreement is good, which might suggest that parametric fitting is adequate. Figure 5.8 shows that it is not.

by nearly a factor of two. Quantitatively, the E_{tot} bias is -2.6% and the 68% width is 5.4%; the N_{delep} bias is $+96\%$ with a 68% width of 44%. The truth point at $(1,1)$ lies at the 2σ contour.

Crucially, *this bias is already present in the Asimov (no-noise) fit*, before any Poisson fluctuations are applied. Fitting the parametric model to the noiseless truth spectrum returns $\langle E_{\nu_e} \rangle = 9.82$ MeV and $\alpha_{\nu_e} = 1.52$ for the ν_e component, with N_{ν_e} inflated correspondingly. The fitter has traded the mean energy against the normalization to minimize the NLL under a mis-specified model: since no single pinched spectrum can exactly reproduce the time-integrated shape, the optimizer finds a compensating combination of parameters that fits the detected counts well while getting the underlying physics wrong. The inflated N_{ν_e} propagates directly into $N_{\mathrm{delep}} = N_{\nu_e} - N_{\bar{\nu}_e}$, since $N_{\bar{\nu}_e}$ is independently anchored by SK IBD. The result is a systematic overestimate that does not diminish with more data: *adding more events or more channels makes the wrong answer more precisely wrong.*

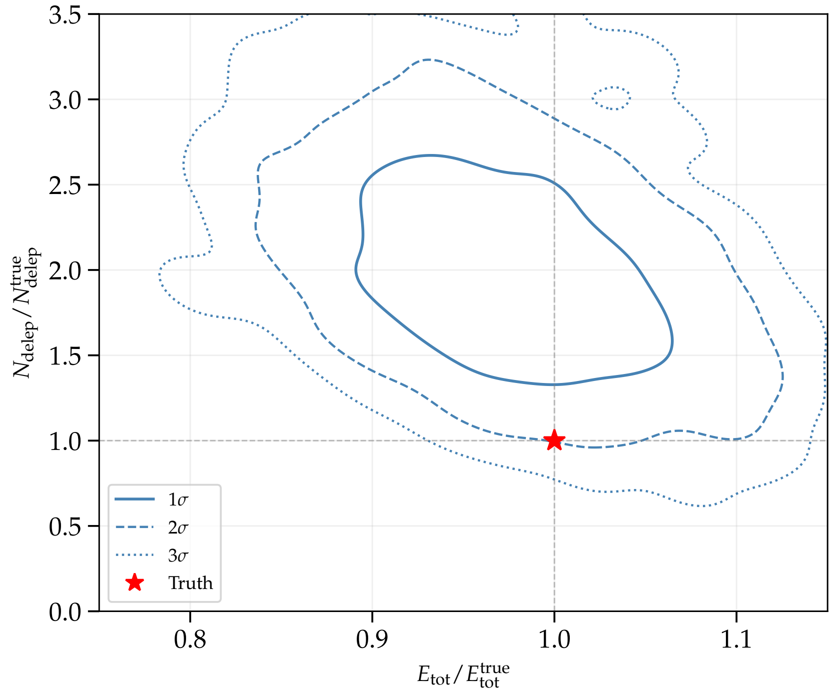

Figure 5.8: Joint reconstruction of E_{tot} and N_{delep} using a nine-parameter pinched-spectrum fit to all five channels, from 10^3 Poisson toys. Contours enclose 1σ, 2σ, and 3σ of the toy distribution; the red star marks truth at $(1, 1)$. Despite the visual adequacy of the spectral fits in Fig. 5.7, the reconstructed N_{delep} is biased by +96% — a model bias already present in the Asimov fit that does not improve with more data.

5.6 Unfolding Considerations

When the next Galactic supernova neutrino burst occurs, detectors around the world will record event spectra across multiple interaction channels. The question we ask is how well one can infer the true neutrino emission from these detections—and in particular, how precisely one can recover the total energy and deleptonization of the explosion. Inferring the true flux from detected spectra is the spectral unfolding problem.

At its core, unfolding is an ill-conditioned matrix inversion problem, and there is no single best method: the right choice depends on the specific regularization one is willing to impose. Spectral unfolding seeks to invert Eq. 5.3 to recover $\phi(E_\nu)$. The response matrix is generally ill-conditioned: kinematic smearing (e.g., the broad recoil energy distribution in elastic scattering) and detector energy resolution spread the signal across detected bins in ways that are difficult to invert without amplifying noise. Different unfolding methods handle this regularization differently.

Throughout this work, we treat the response matrices as known exactly: uncertainties on interaction cross sections, detector geometry, and energy resolution are not propagated. The only uncertainty we consider is the statistical (Poisson) fluctuation of the detected event counts.

5.6.1 Iterative Bayesian Unfolding

Iterative Bayesian Unfolding (IBU) (D'Agostini, 1995, 2010) is a natural first approach to spectral unfolding. It represents the true spectrum as a histogram with no imposed functional form and makes no shape assumptions whatsoever. Starting from a flat prior π_i over true bins, the algorithm applies Bayes' theorem to compute the probability that a detected event in bin j originated from true bin i,

$$P(C_i \mid E_j) = \frac{M_{ji}\,\pi_i}{\sum_k M_{jk}\,\pi_k}\,, \tag{5.12}$$

and uses this to update the unfolded spectrum,

$$n_i^{(t+1)} = n_i^{(t)} \sum_j \frac{M_{ji}\,n_j^{\text{obs}}}{\sum_k M_{jk}\,n_k^{(t)}}\,. \tag{5.13}$$

The prior is updated to $\pi_i \propto n_i^{(t)}$ at each step. Above the detector threshold, IBU is genuinely model-independent: no spectral shape is assumed, and the reconstruction reflects only the data and the response matrix.

Unfolding always involves a tradeoff between bias and noise amplification. For IBU, the iteration count acts as the implicit regularization parameter: more iterations recover finer spectral detail at the cost of amplifying noise. The stopping point must therefore be chosen to balance these competing effects.

For all channels, convergence is assessed via the relative change in the spectral estimate between successive iterations,

$$\delta^{(k)} = \frac{\|n^{(k)} - n^{(k-1)}\|}{\|n^{(k-1)}\|}, \tag{5.14}$$

and the iteration is terminated when $\delta^{(k)} < 6 \times 10^{-3}$, or after a maximum of 100 iterations, whichever occurs first. For the well-measured channels (SK IBD and DUNE CC), the choice of tolerance has negligible impact on the result. The JUNO channel is more sensitive: reducing the tolerance below 10^{-3} narrows the toy-to-toy spread but degrades spectral shape fidelity, shifting the reconstructed mean away from the true spectrum particularly at low energies (see Appendix 5A). This reflects the underdetermined nature of the JUNO inversion, where additional iterations amplify noise rather than refine the solution. The chosen tolerance of 6×10^{-3} balances variance reduction against reconstruction bias.

Figure 5.9 shows the IBU reconstruction for each channel. For Super-K (left panel, green band) the unfolded spectrum closely tracks the true $\bar{\nu}_e$ distribution across the full energy range. The IBD threshold at 4.3 MeV falls well below the peak of the $\bar{\nu}_e$ spectrum and costs a negligible fraction of the signal. For DUNE (center panel, purple band), the 3 MeV electron detection threshold propagates to a cutoff in reconstructed neutrino energy near ∼10 MeV, below which IBU returns nothing; the reconstruction works well above this threshold, though uncertainties grow toward it.

For JUNO (right panel, red band), we consider only the ν_x-electron scattering sample, with no contamination from ν_e or $\bar{\nu}_e$. Even so, the reconstructed spectrum has the largest uncertainty of the three channels and does not follow the true shape well. This reflects the lower-triangular response matrix of elastic scattering: a single neutrino energy maps to a broad

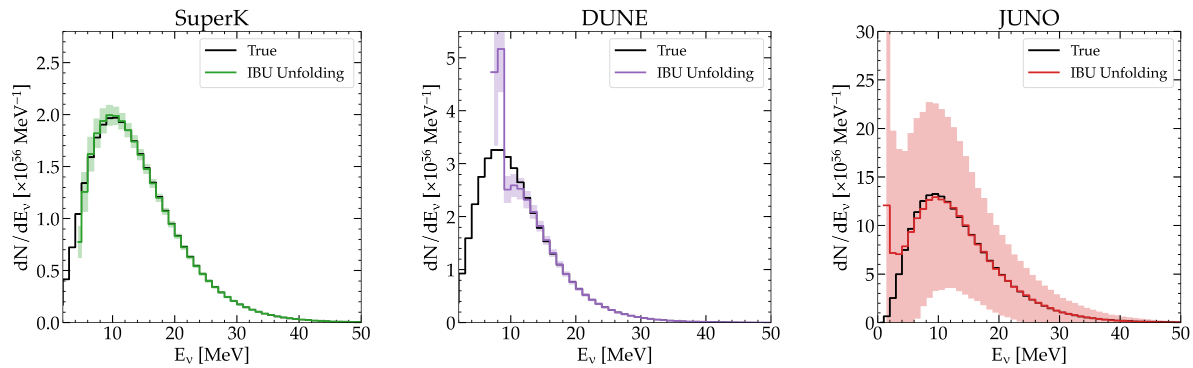

Figure 5.9: IBU reconstruction of the neutrino spectra for the three detectors considered in this work. The black curves show the true spectra and the colored curves show the unfolded results obtained from iterative Bayesian unfolding. Shaded bands represent the spread from repeated Poisson realizations of the detected spectra. Differences in the reconstruction performance reflect the detector response and interaction channels: inverse beta decay in Super-K, charged-current ν_e absorption in DUNE, and elastic scattering in JUNO.

range of recoil electron energies (see Figure 5.5), making the inversion poorly constrained.

The fundamental limitation of IBU, and many other unfolding methods, is that it cannot recover flux below the detector threshold — there is nothing to invert where no signal exists. This is most severe for the JUNO ν-p channel, but also serious for the DUNE ν_e CC channel. In addition, a totally unconstrained unfolding of the $\simeq 100$ ν_x-electron sample in JUNO is also challenging.

These limitations motivate the B-spline approach described in the following section.

5.7 B-spline Spectral Unfolding

Both limitations of IBU point toward the same solution: parametrize the flux with a smooth basis that can extrapolate below threshold under a minimal prior, and allow a joint simultaneous fit across all channels and flavors.

5.7.1 Fitting Procedure

We represent each flavor's flux as a cubic B-spline,

$$\phi(E_\nu) = \sum_{k=1}^{N_c} c_k \, B_k(E_\nu) \,, \tag{5.15}$$

where $B_k(E_\nu)$ are cubic B-spline basis functions (piecewise polynomials, each nonzero over a localized energy interval) spanning 0–80 MeV, and $c_k \geq 0$ are the coefficients to be fitted. Because adjacent basis functions overlap, adjacent coefficients c_k and c_{k+1} control neighboring energy intervals, so the vector $\mathbf{c}$ encodes the spectral shape in a spatially ordered way. We use $N_c = 14$ coefficients, providing sufficient spectral flexibility while remaining regularizable. Two physical constraints are imposed: the flux vanishes at $E_\nu = 0$ (boundary condition, $c_1 = 0$) and is non-negative everywhere (positivity, $c_k \geq 0$). The positivity constraint is essential for meaningful sub-threshold extrapolation.

We "unfold" the neutrino spectrum by fitting the spline coefficients: we take the spline-evaluated flux and forward fold it to predict a detected spectrum, and we find the coefficients by minimizing the Poisson negative log-likelihood in detector space,

$$\mathcal{L}(\{c_k\}) = \sum_j \left[n_j^{\text{pred}} - n_j^{\text{obs}} \ln n_j^{\text{pred}} \right] . \tag{5.16}$$

To regularize the solution, we include an additional smoothness penalty,

$$\mathcal{L}_{\text{reg}} = \mathcal{L} + \lambda \, \|\mathbf{D}_2 \, \mathbf{c}\|^2 \,, \tag{5.17}$$

where $\mathbf{D}_2\mathbf{c}$ computes second differences of adjacent coefficients (the discrete analog of $d^2\phi/dE_\nu^2$), penalizing curvature in the unfolded spectrum, and λ controls the regularization strength. This penalty suppresses high-frequency oscillations in the unfolded spectrum without imposing any parametric shape; it is the minimal prior needed to make the sub-threshold extrapolation well-posed. Crucially, it is strictly weaker than the pinched Fermi-Dirac assumption of the standard approach: it rules out pathological oscillations but imposes no global spectral shape. We take $\lambda = 10^{-4}$ for Super-Kamiokande IBD and $\lambda = 10^{-3}$ for all other channels, chosen empirically by minimizing reconstruction bias on the reference model; the lower value for Super-K reflects its well-conditioned, nearly diagonal response matrix and high statistics, which require less regularization. The minimization is performed with the L-BFGS-B algorithm (Byrd et al., 1995; Zhu et al., 1997; Morales & Nocedal, 2011), a bounded quasi-Newton method that natively enforces the $c_k \geq 0$ constraints.

In the full analysis, we simultaneously fit three independent B-splines (ϕ_{ν_e}, $\phi_{\bar{\nu}_e}$, ϕ_{ν_x}) across all five channels, for 42 parameters in total. Each channel's predicted counts depend on whichever flavor(s) it observes through its response matrix, so the joint likelihood automatically encodes the flavor mixing in JUNO and HK. For ν_x specifically, all three ν_x-sensitive channels — JUNO ν-e, JUNO ν-p, and HK ν-e — share a single B-spline with their Poisson NLLs summed, directly recovering ϕ_{ν_x} from their combined spectral information.

We assess reconstruction performance via a toy Monte Carlo procedure. For each toy, we Poisson-resample the nominal predicted counts in each detected-energy bin and refit the spline. From 10^4 toys we compute the bias (median reconstructed integral divided by truth, minus one) and the statistical width (half the 16th-to-84th percentile range divided by truth) of the integrated energy E_{tot} and deleptonization number N_{delep}, integrated over the full energy range.

5.7.2 Three-Flavor Spectral Reconstruction

Figure 5.10 shows the reconstructed spectrum for each flavor from the five-channel joint fit; we discuss the per-flavor performance in turn.

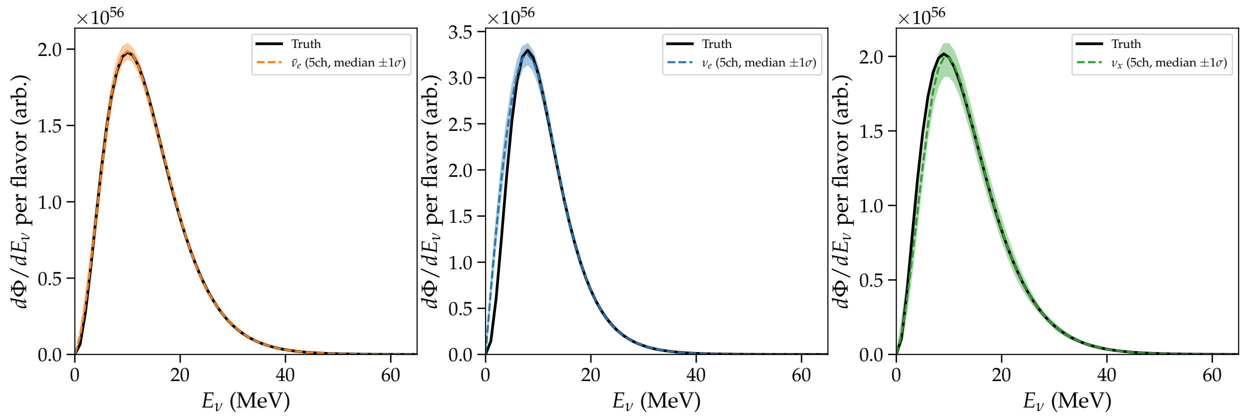

Figure 5.10: B-spline reconstruction of the $\bar{\nu}_e$ (left), ν_e (center), and per-flavor ν_x (right) spectra from the five-channel joint fit at 10 kpc. In each panel, the black curve shows the true spectrum; the colored curve and shaded band show the median and 68% interval from 10^4 Poisson toys.

$\bar{\nu}_e$ reconstruction. In the joint fit, the $\bar{\nu}_e$ reconstruction is anchored by Super-Kamiokande IBD and is essentially self-sufficient: the other channels shift the $\bar{\nu}_e$ result negligibly, since IBD is the only purely $\bar{\nu}_e$-dominated channel and SK's $\simeq$6500 events dwarf the $\bar{\nu}_e$ contribution to any other channel. The IBD response matrix is also nearly one-to-one (as shown in Fig. 5.3), making this the best-conditioned unfolding problem of the five channels, with no need for aggressive sub-threshold extrapolation. The recovered total energy has a bias of +0.2% and a width of 1.3%; the recovered total count has a bias of +0.8% and a width of 2.0%. The single-channel performance is detailed in Appendix 5C.

ν_e reconstruction. In the joint fit, ν_e is anchored by DUNE CC. The sub-threshold extrapolation below $\sim$7 MeV benefits from the other channels in the joint fit. JUNO ν-e, with its 0.2 MeV threshold, is sensitive to ν_e at energies where DUNE is blind; this provides an indirect anchor on the low-energy spectral shape. However, JUNO ν-e is not a pure ν_e channel — ν_x also scatters there — so there is an intrinsic degeneracy between ϕ_{ν_e} and ϕ_{ν_x} in the JUNO data. This degeneracy is only partial: ϕ_{ν_x} is simultaneously constrained at higher energies by HK ν-e and JUNO ν-p, and the spline smoothness propagates those constraints back to low energies, preventing a free trade-off between ϕ_{ν_e} and ϕ_{ν_x} below threshold (see Appendix 5C). The recovered total energy has a bias of +1.5% and a width of 1.9%; the recovered total count has a bias of +6.4% and a width of 2.8%, consistent with the expectation from Fig. 5.4 that N_{ν_e} is disproportionately sensitive to the sub-threshold region.

ν_x **reconstruction.** The ν_x spectrum is the hardest to reconstruct: it enters JUNO and HK only through flavor-mixed channels, with no dedicated single-flavor anchor. Unlike $\bar{\nu}_e$, which is isolated by IBD at Super-K, or ν_e, which is isolated by CC absorption at DUNE, ν_x shares every available channel with ν_e and $\bar{\nu}_e$; separating the three contributions requires solving for all flavors simultaneously from the same data.

Figure 5.11 shows the result of a joint fit to the three ν_x-sensitive channels alone — JUNO ν-e, JUNO ν-p, and HK ν-e — without contamination from ν_e and $\bar{\nu}_e$ contributions to these samples. We can see that above the 22 MeV threshold, the ν-p channel in JUNO can pin down the spectrum very well, despite the highly compressed response matrix in this channel. Between about 5 MeV to 22 MeV, the HK sample contributes dominantly thanks to its large statistics. Below the ∼5 MeV threshold in HK, the JUNO ν-e sample constrains the spline. Despite the low event count in the JUNO ν-e channel, it is the most valuable for recovering the total integral: its near-zero threshold gives sensitivity to the bulk of the ν_x flux that the other channels miss. Quantitatively, adding HK ν-e to the two JUNO channels roughly halves the energy width from 8.3% to 4.8% and reduces the count width from 18% to 15%; the triple-channel fit recovers the total ν_x energy with a bias of +1.4% and a width of 4.8%, and the total count with a bias of +7.2% and a width of 15% (integrated over all E_ν).

In the full five-channel joint fit, simultaneously solving for all three flavors further improves the ν_x reconstruction beyond the three-channel result, as shown in the right panel of Fig. 5.10. At first this may seem counterintuitive: the three-channel ν_x-only fit seems like it should give a better fit because of the lack of contamination. But, actually, the improvement with

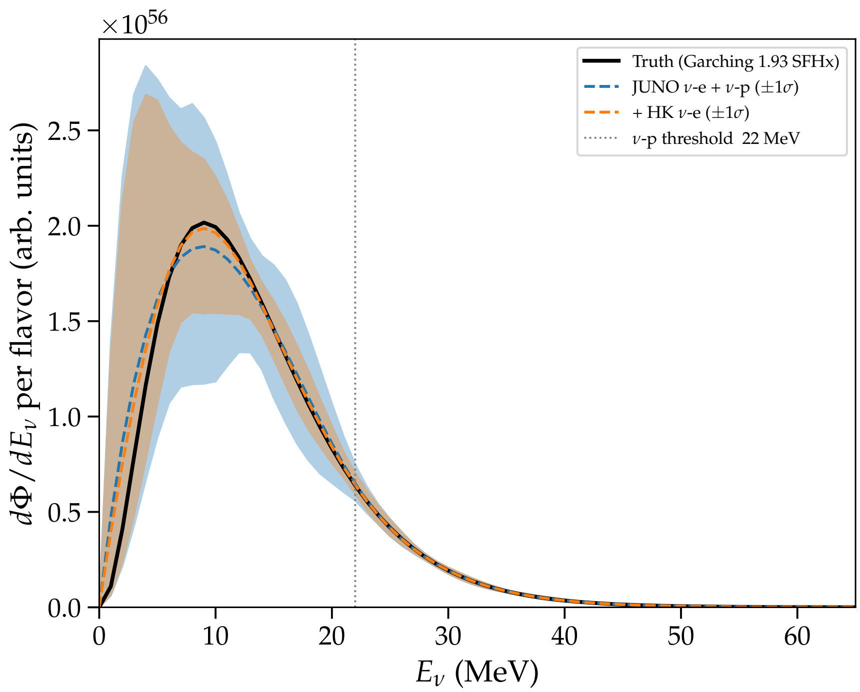

Figure 5.11: B-spline reconstruction of the per-flavor ν_x spectrum from a joint fit to the three ν_x-sensitive channels at 10 kpc. The blue band shows the result using JUNO ν-e and ν-p alone; the orange band shows the improvement from adding HK ν-e. The dotted vertical line marks the ν-p detection threshold at 22 MeV.

the joint fit comes from a shared event budget. Note that both the standalone DUNE ν_e fit and the three-channel ν_x fit suffer from a common pathology: at low energies where detected events are sparse or absent, the B-spline tends to overshoot the true flux, driven by smoothness and positivity alone with no data to anchor it. In the five-channel joint fit, ϕ_{ν_e} and ϕ_{ν_x} cannot both overshoot simultaneously, because they share the JUNO ν-e event budget: the total predicted counts at JUNO—from ν_e, $\bar{\nu}_e$, and ν_x combined—must match the observed total. If the ν_e spline rises too high at low energies, the predicted JUNO counts exceed what is observed, and the NLL pushes back; the ϕ_{ν_e} and ϕ_{ν_x} contributions are anti-correlated through this shared constraint. Moreover, DUNE's $\simeq$4000 events constrain ϕ_{ν_e} well above threshold, and spline smoothness propagates that constraint downward to energies where DUNE is blind; this anchors the low-energy ν_e contribution at JUNO, which in turn stabilizes the ν_x reconstruction there. The two effects together — mutual competition for the JUNO event budget, and DUNE's indirect low-energy anchor — suppress the low-energy overshoot in both flavors and bring the five-channel ν_x reconstruction substantially closer to truth. The recovered total energy has a bias of -1.3% and a width of 3.8%, and the recovered total count has a bias of -3.3% and a width of 3.8%.

5.7.3 Combined Three-Flavor Results

The key combined observables from the five-channel fit are the total emitted energy,

$$E_{\text{tot}} = E_{\nu_e} + E_{\bar{\nu}_e} + 4\,E_{\nu_x}\,, \tag{5.18}$$

where the factor of four accounts for the four degenerate heavy-lepton flavors, and the net deleptonization number,

$$N_{\text{delep}} = N_{\nu_e} - N_{\bar{\nu}_e}\,. \tag{5.19}$$

Figure 5.12 shows the joint distribution of E_{tot} and N_{delep} from 10^4 toys.

The total energy is recovered with a bias of -0.7% and a width of 2.5%, considerably better than any individual flavor in isolation. This improvement arises from a partial cancellation of per-flavor biases in the sum: the slight overestimate of E_{ν_e} (+1.5%) and the slight underestimate of E_{ν_x} (-1.3%) push in opposite directions and largely cancel, while the $\bar{\nu}_e$ term, anchored by Super-K IBD, contributes with nearly zero bias (+0.2%). Because ν_x enters with four per-flavor contributions it dominates E_{tot}, and its small negative bias pulls the sum slightly below truth—but only by 0.7%. The 2.5% width reflects the combined statistical uncertainty of all five channels propagated through the sum.

The deleptonization number presents a sharper challenge. Despite N_{ν_e} being recovered to +6.4% and $N_{\bar{\nu}_e}$ to +0.8%—both modest biases in isolation—the reconstructed N_{delep} carries a bias of +26% and a width of 15%. This amplification is a generic consequence of measuring a

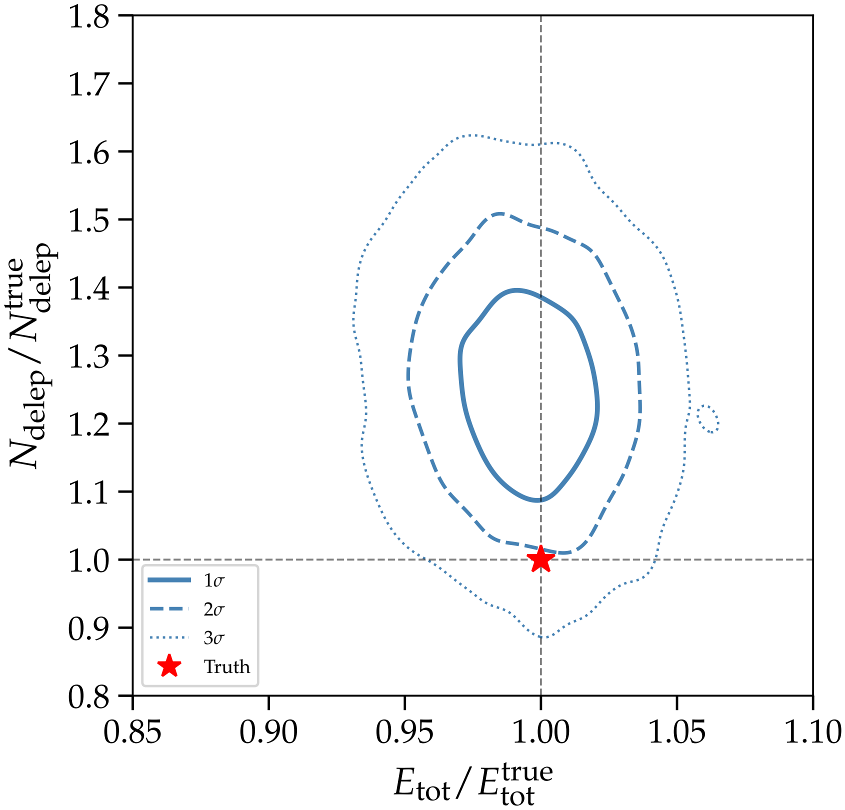

Figure 5.12: Joint reconstruction of E_{tot} and N_{delep} from the five-channel joint fit, based on 10^4 Poisson toys. Both axes are normalized to their true values. Contours enclose 1σ, 2σ, and 3σ of the toy distribution; the red star marks truth at $(1, 1)$. The tight E_{tot} distribution and the separation of N_{delep} from zero confirm robust recovery of both observables.

small difference between two large numbers: the deleptonization excess $N_{\mathrm{delep}} = N_{\nu_e} - N_{\bar{\nu}_e}$ is a small fraction of the total electron-neutrino count N_{ν_e}, so the same absolute error on N_{ν_e} that appears modest relative to N_{ν_e} itself becomes large relative to N_{delep}. The dominant source of bias is the $+6.4\%$ overestimate of N_{ν_e} driven by the DUNE sub-threshold extrapolation, as discussed in Sec. 5.7.2. Despite the large fractional bias, the distribution of N_{delep} is cleanly separated from zero: as seen in Fig. 5.12, the entire 1σ contour lies well above $N_{\mathrm{delep}} = 0$, robustly confirming a net positive neutronization of the stellar core and directly establishing that the collapse produced more ν_e than $\bar{\nu}_e$.

5.7.4 Results without DUNE

Among the five channels in our analysis, DUNE plays a qualitatively distinct role: it is the only channel with a dedicated ν_e signature, via charged-current absorption on argon. The remaining channels either isolate $\bar{\nu}_e$ or observe a flavor-mixed signal in which ν_e, $\bar{\nu}_e$, and ν_x all contribute. Removing any other single detector degrades reconstruction performance but leaves the flavor decomposition scheme qualitatively intact; removing DUNE eliminates the only clean ν_e anchor entirely.

In this subsection, we quantify this through an ablation study: we repeat the full joint fit with DUNE removed, leaving only SK IBD, JUNO ν-e, JUNO ν-p, and HK ν-e. The reconstructed spectra in this case are shown in Appendix 5D (Fig. 5.21). Without a dedicated ν_e anchor, the optimizer cannot independently resolve ϕ_{ν_e} and ϕ_{ν_x} from the JUNO and HK data alone, because both flavors contribute to those channels with partially degenerate spectral signatures. In the absence of DUNE, the optimizer tends to over-attribute JUNO events to ν_e, since ν_e carries the largest cross section in the ν-e channel. This artificially inflates ϕ_{ν_e}, which in turn starves ϕ_{ν_x}: once the excess JUNO events are assigned to ν_e, fewer remain to be explained by ν_x, and the optimizer suppresses ϕ_{ν_x} accordingly. Because ν_x enters E_{tot} with a factor of four, even a moderate suppression of ϕ_{ν_x} drags E_{tot} far below truth. Quantitatively, the ν_e energy width degrades from 1.9% to 6.8%, reflecting the loss of the DUNE spectral anchor. The ν_x total count bias worsens from -3.3% to -28%, a direct consequence of the event mis-attribution described above.

Figure 5.13 illustrates the full degradation in total reconstructed energy and deleptonization. The total energy bias worsens from -0.7% to -11%, driven by the cascading ν_x suppression amplified by the factor of four. Critically, the deleptonization number is recovered with a bias of $+84\%$ and a width of 39%: the same amplification mechanism discussed in Sec. 5.7.3—a small difference between two large numbers—now operates on a catastrophically inflated N_{ν_e} rather than a modestly overestimated one, and the 39% width is so large that the measurement can no longer robustly establish $N_{\rm delep} > 0$: the orange contours are broadly displaced from truth and the lower tail of the distribution approaches zero, erasing the confirmation of core neutronization that the five-channel fit achieves cleanly. Without DUNE, neither $E_{\rm tot}$ nor $N_{\rm delep}$ can be reliably extracted, and the full-flavor characterization program fails.

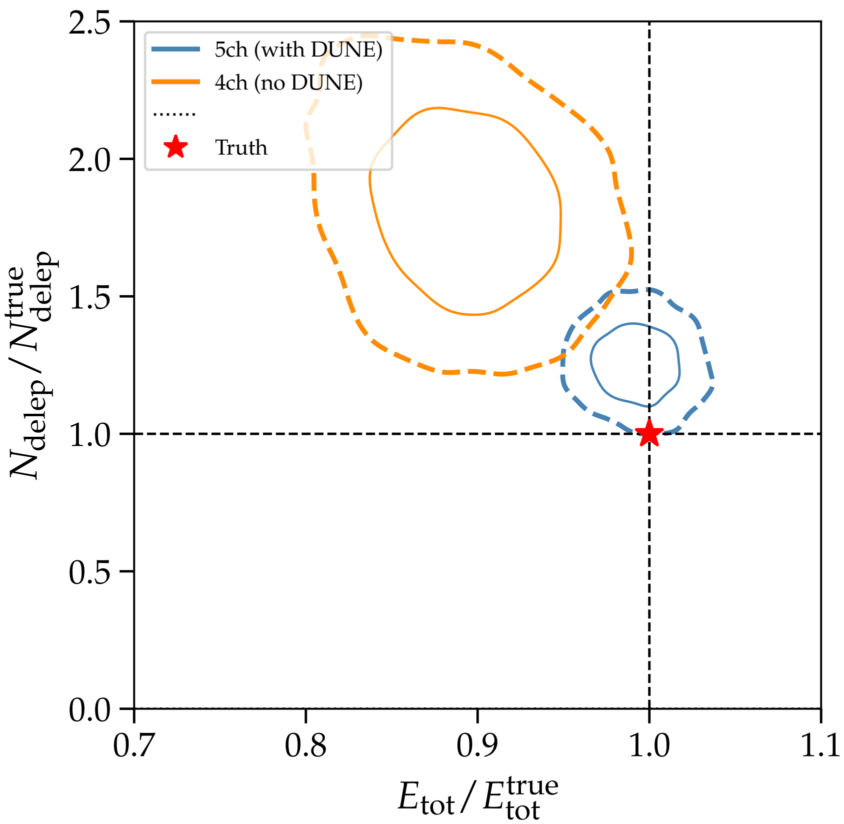

Figure 5.13: Joint reconstruction of $E_{\rm tot}$ and $N_{\rm delep}$ with DUNE (blue) and without DUNE (orange), from 10^4 Poisson toys each. Contours enclose 1σ, 2σ, and 3σ of each toy distribution; the red star marks truth. Without DUNE, both $E_{\rm tot}$ and $N_{\rm delep}$ are severely biased and the constraint on $N_{\rm delep} > 0$ is lost.

5.7.5 Reconstruction performance across progenitor models

All results in Secs. 5.7.2–5.7.4 use the reference model ($1.93\,M_\odot$ + SFHx EOS). To verify that the B-spline approach is not tuned to this particular progenitor, we extend the analysis to all 20 Garching models (five baryonic masses $\times$ four EOS), running 10^4 Poisson toys per model with the full five-channel fit.

We find that $E_{\rm tot}$ is recovered without significant bias across the entire grid: the absolute bias is $\lesssim 2\%$ for all 20 models, with no systematic trend in either mass or EOS. The 68% reconstruction width decreases monotonically with $M_{\rm bar}$, from $\approx 5.5\%$ at $1.36\,M_\odot$ to $\approx 2.5\%$ at $1.93\,M_\odot$. This improvement reflects the growing total event count: higher-mass progenitors emit roughly 2.4 times more energy in neutrinos, providing a tighter spectral constraint across all channels. At fixed mass, the four EOS variants give consistent reconstruction widths, confirming that the B-spline is not sensitive to the moderate spectral shape differences between EOS. Full numerical results for all 20 models are given in Tables 5.1 and 5.2.

While $E_{\rm tot}$ can be recovered with bias $\lesssim 2\%$ and width ≈ 2.5–5.5%, $N_{\rm delep}$ is substantially harder to reconstruct. The root cause is that $N_{\rm delep} = N_{\nu_e} - N_{\bar{\nu}_e}$ is a difference of two large numbers: a small fractional error in either flavor propagates into a large relative error on their difference. More concretely, the ν_e reconstruction carries a positive bias from the low-energy extrapolation below the DUNE threshold (discussed in Sec. 5.7.3), which inflates N_{ν_e} and thereby $N_{\rm delep}$.

Across the 20 Garching models, the 10^4-toy sweep gives a bias of +22% to +33% and a 68% width of ≈ 13–30%, both consistently larger than the corresponding $E_{\rm tot}$ figures by a factor of $\sim$5–10. The width decreases with $M_{\rm bar}$, from ≈ 26–30% at $1.36\,M_\odot$ to ≈ 13–16% at $1.93\,M_\odot$, mirroring the $E_{\rm tot}$ trend but at much larger absolute values. Even at the most favorable (highest-mass) models, the reconstruction width alone is comparable to the $\sim$14% theoretical uncertainty from $Y_{e,i}$ discussed in Sec. 5.8.3.

5.8 What can we learn about supernovae?

The reconstruction of Sec. 5.7 enables two concrete inferences about the newly formed PNS. First, $E_{\rm tot}$ tracks the gravitational binding energy of the remnant and directly constrains its mass; we quantify the precision of this mass inference and compare it to existing neutron-star mass measurements from X-ray and gravitational-wave observations. Second, $N_{\mathrm{delep}} > 0$ can be robustly established across the same model space, directly confirming neutronization of the stellar core.

5.8.1 E_{tot} as a PNS mass observable

Because E_{tot} traces the gravitational binding energy of the newly formed PNS, it scales directly with its mass and compactness. Figure 5.14 quantifies this mass–energy relation using both theoretical considerations and supernova simulations. The dashed curves show the cold-NS binding energy $E_{\mathrm{bind}} = (M_{\mathrm{bar}} - M_{\mathrm{grav}})c^2$ computed from the Tolman–Oppenheimer–Volkoff (TOV) equation for four EOS commonly used in supernova simulations. The TOV equation describes the hydrostatic equilibrium of a neutron star under general relativity; given an equation of state, it predicts the stellar mass and radius as a function of central density, and hence the gravitational binding energy released when diffuse matter assembles into a compact remnant.

We see that there is clear scaling between E_{tot} and M_{bar}. The variation is large — E_{tot} rises monotonically with M_{bar} from $\approx 1.8 \times 10^{53}$ erg at $1.36\,M_\odot$ to $\approx 4.3 \times 10^{53}$ erg at $1.93\,M_\odot$, a factor of 2.4 across the models in the figure — while the EOS spread at any fixed mass is only ∼10–15%. A measurement of E_{tot} therefore constrains M_{bar} robustly, with the EOS as the dominant residual systematic.

The TOV curves also reveal the detailed EOS dependence. At fixed M_{bar}, DD2, the stiffest EOS considered here (maximum gravitational mass $\approx 2.42\,M_\odot$, radius ≈ 13.5 km at $1.4\,M_\odot$), has the lowest binding energy. Gravitational binding energy scales roughly as GM^2/R: at fixed mass, a star with larger R has less binding energy because gravity has done less work assembling matter into a compact remnant. A stiffer EOS resists compression more strongly, producing a larger neutron star for the same baryonic mass and therefore a lower binding

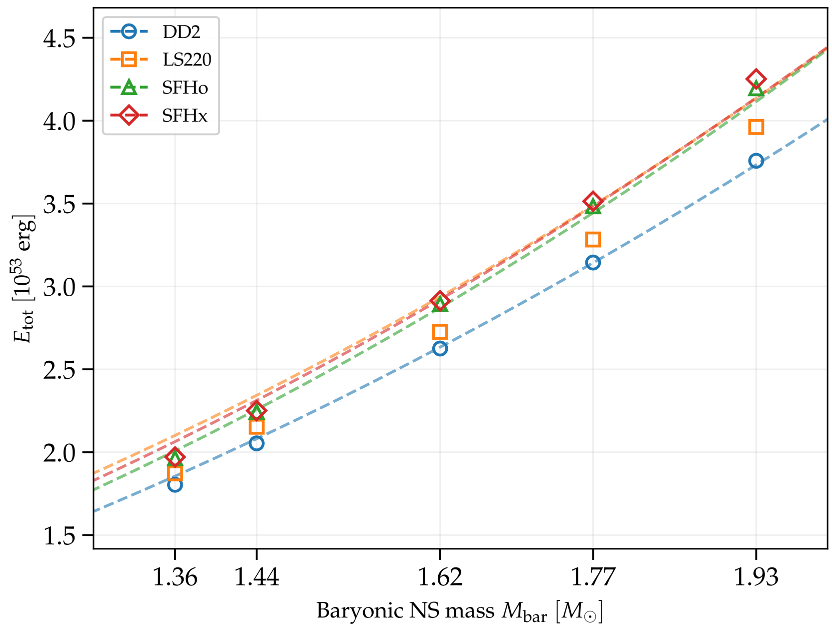

Figure 5.14: Total emitted energy $E_{\rm tot}$ vs. baryonic NS mass $M_{\rm bar}$ for the 20 Garching models (five masses $\times$ four EOS). Dashed curves show the cold-NS binding energy $(M_{\rm bar} - M_{\rm grav})c^2$ from the TOV equation for each EOS. Open symbols show the simulated $E_{\rm tot}$ from each model. The gap between the symbols and their TOV curve reflects residual thermal energy in the hot PNS at simulation end.

energy. The three softer EOS — LS220, SFHo, and SFHx — produce nearly identical curves lying $\sim$15% above DD2, reflecting their smaller, more compact remnants.

The open symbols show the simulated $E_{\rm tot}$ from the 20 Garching models. Because $\gtrsim 99\%$ of the binding energy is released through neutrinos (Janka, 2012; Burrows & Vartanyan, 2021; Li et al., 2021), we might naively expect the simulations to fall on the TOV curves. In fact, we see that all the simulations fall at or below their corresponding TOV curves. This is physically expected: the TOV binding energy is the asymptotic cold-NS limit, the total energy that would be emitted only after complete Kelvin–Helmholtz cooling. The Garching simulations run for finite time, and the hot PNS retains residual thermal energy at simulation end; the gap between each symbol and its TOV curve is precisely this remaining thermal reservoir. DD2 shows the smallest gap: its stiffer EOS produces a larger, less dense remnant that stores less thermal energy in the hot newly-born phase, so the simulations capture nearly the full emission before they end. The softer EOS — LS220, SFHo, and SFHx — produce more compact, hotter remnants that store more thermal energy, and their simulated $E_{\rm tot}$ therefore sits further below the TOV prediction. LS220 shows the largest gap despite having a similar cold radius to SFHo and SFHx, likely reflecting differences in the hot dense-matter thermodynamics. Since the mass inference in Sec. 5.8.2 equates the reconstructed $E_{\rm tot}$ with the cold-NS binding energy via TOV inversion, this residual thermal gap introduces a systematic underestimate of $M_{\rm bar}$; with the Garching simulations capturing $\gtrsim 98\%$ of the total emission, the effect is $\lesssim 2\%$ and is absorbed into the reconstruction biases tabulated in Table 5.1.

A further argument for the robustness of $E_{\rm tot}$ as a mass probe is its insensitivity to neutrino flavor transformation. Our analysis assumes no oscillations between the source and detector (Sec. 5.3), but in reality collective effects, MSW resonances, or shock-induced flavor conversion may redistribute energy among flavors during propagation. Because oscillations are unitary transformations on the flavor space, they conserve the total neutrino energy summed over all species: any energy transferred from ν_e to ν_x, for example, cancels exactly in the combination $E_{\rm tot} = E_{\nu_e} + E_{\bar{\nu}_e} + 4E_{\nu_x}$. The value of $E_{\rm tot}$ reconstructed from a joint multi-detector fit is therefore equal to the total energy emitted at the source regardless of the oscillation history — it is a statement about the PNS binding energy, not about flavor composition at Earth.

The following subsection shows that our B-spline reconstruction recovers $E_{\rm tot}$ with sufficient precision across all 20 models to exploit this scaling directly.

5.8.2 Precision of PNS mass measurement

Section 5.7.5 established that the five-channel B-spline recovers $E_{\rm tot}$ with a 68% width of ≈ 2.5–5.5% and bias $\lesssim 2\%$ across all 20 Garching models. To convert this into a statement about $M_{\rm bar}$, we invert the TOV relation: for each toy-reconstructed $E_{\rm tot}$ sample, we solve

$E_{\rm bind}(M_{\rm bar}) = E_{\rm tot}$ to obtain the corresponding $M_{\rm bar}$ estimate. The 68% half-width of the resulting $M_{\rm bar}$ distribution, normalized to the true $M_{\rm bar}$, is the mass measurement precision.

Figure 5.15 shows this precision as a function of $M_{\rm bar}$ for three scenarios. The green curve (fixed EOS, SFHo) draws samples from the SFHo models and inverts using the SFHo TOV curve; it represents the floor set by reconstruction noise for a known EOS, and falls from $\approx$ 2.6% at $1.36\,M_\odot$ to $\approx 1.2\%$ at $1.93\,M_\odot$. This improvement with mass reflects two reinforcing effects: higher-mass PNS emit more total energy, tightening the $E_{\rm tot}$ reconstruction, and they sit on a steeper part of the $E_{\rm bind}(M_{\rm bar})$ curve, so the same $\Delta E_{\rm tot}$ maps to a smaller $\Delta M_{\rm bar}$. The gray dashed curve marginalizes over the four EOS with equal weight; the EOS uncertainty adds in quadrature with the reconstruction noise and raises the total to $\approx$ 3–4%, roughly flat across the mass range because the EOS spread in $M_{\rm bar}$ at fixed $E_{\rm tot}$ is nearly independent of mass. The black dotted curve further folds in a 2% distance uncertainty ($\Delta E_{\rm tot}/E_{\rm tot} \approx 2\,\Delta d/d \approx 4\%$); the additional degradation is modest, confirming that EOS uncertainty is the dominant systematic for this measurement.

For context, the shaded blue band in Fig. 5.15 shows the typical precision of existing NS mass measurements from X-ray pulse-profile modeling with NICER (Miller et al., 2019; Riley et al., 2019) ($\approx$ 11% for PSR J0030+0451) and from gravitational-wave observations of binary NS mergers (Abbott et al., 2017a,b) (individual component mass measurements typically 8–15% at 1σ). All three of our curves sit well below this band, indicating that a full flavor neutrino measurement has the best constraining power on $\mathrm{M}_{\rm bar}$. The five-channel B-spline fit achieves $\approx$ 1–4% precision on $M_{\rm bar}$ from a single supernova detection, competitive with or better than existing methods, without requiring a binary companion (though most massive stars (i.e. CCSN progenitors) are found in binaries $\sim$ 70% of the time (Sana et al., 2012; Moe & Di Stefano, 2017)), and probing the neutron star at the moment of its formation. We note that NICER and gravitational-wave analyses measure gravitational mass $M_{\rm grav}$, while the TOV inversion returns baryonic mass $M_{\rm bar} \approx 1.1$–$1.2\,M_{\rm grav}$ for typical neutron stars;

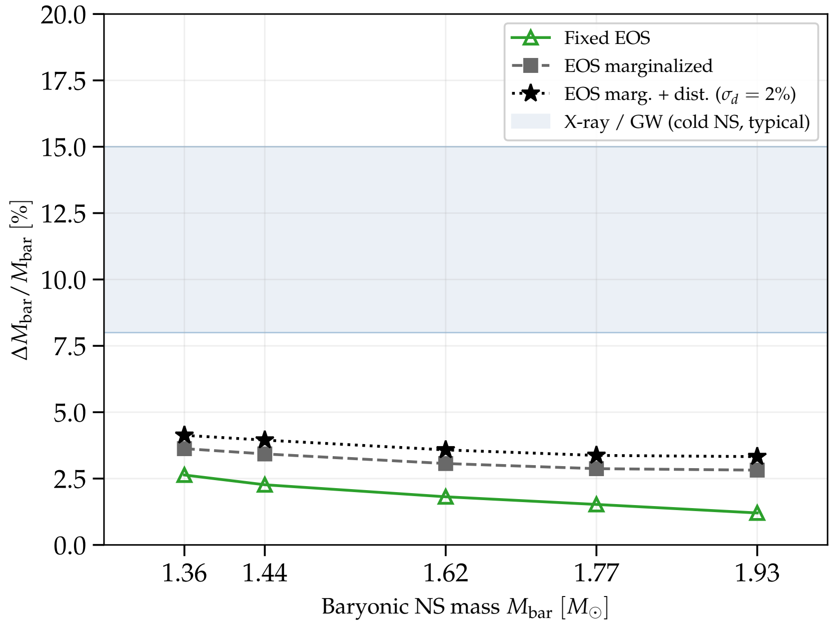

Figure 5.15: Precision $\Delta M_{\rm bar}/M_{\rm bar}$ with which the PNS baryonic mass can be inferred from the reconstructed $E_{\rm tot}$, as a function of $M_{\rm bar}$, for the 20 Garching models. Green solid (fixed EOS, SFHo): reconstruction floor assuming the EOS is known. Gray dashed (EOS marginalized): equal weight over four EOS in the inversion, folding in EOS uncertainty. Black dotted (EOS marg. + dist.): additionally includes a 2% distance uncertainty, which propagates as a 4% uncertainty on $E_{\rm tot}$. The blue shaded band shows the typical precision of cold-NS mass measurements from X-ray pulse-profile modeling (NICER) and gravitational-wave observations of binary mergers. All three curves lie well below the band.

the fractional precisions quoted here and in the existing literature are nonetheless directly comparable.

5.8.3 $N_{\rm delep}$ as a PNS deleptonization observable

Because $N_{\rm delep} = N_{\nu_e} - N_{\bar{\nu}_e}$ measures the net electron lepton number carried away by neutrinos, it directly traces the deleptonization of the newly formed PNS. As the core collapses, electrons are captured onto protons ($e^- + p \rightarrow n + \nu_e$), converting the pre-collapse core into a neutron-rich remnant. The total deleptonization after complete cooling is $N_{\rm delep}^{\infty} = N_{\rm bar}(Y_{e,i} - Y_{e,f})$, where $Y_{e,f} \sim 0.05$–0.09 is the cold-NS electron fraction in beta equilibrium (calculable from the TOV equation), and $Y_{e,i}$ is the initial electron fraction of the PNS-forming material. Figure 5.16 shows the resulting $N_{\rm delep}$ versus $M_{\rm bar}$ for the same four EOS as Fig. 5.14.

Unlike $E_{\rm tot}$, where the initial gravitational energy of the diffuse pre-collapse core is negligible and the TOV binding energy is purely a statement about the final cold NS, $N_{\rm delep}$ is sensitive to the initial conditions of the progenitor. The iron core at the onset of collapse has $Y_{e,i} \approx$ 0.42–0.46, but the PNS also accretes material from the surrounding silicon shell with $Y_{e,i}$ closer to 0.50, so the effective baryon-weighted initial electron fraction of the full PNS is uncertain over a range $\Delta Y_{e,i} \sim 0.06$–0.08. This propagates to $N_{\rm delep}$ as

$$\frac{\Delta N_{\rm delep}}{N_{\rm delep}} \approx \frac{\Delta Y_{e,i}}{Y_{e,i} - Y_{e,f}} \sim 14\%, \tag{5.20}$$

roughly twice the EOS spread at fixed mass. The shaded bands in Fig. 5.16 reflect this range, $Y_{e,i} \in [0.44, 0.50]$, for each EOS; all 20 Garching simulation points fall within the bands.

We see that $N_{\rm delep}$ still scales monotonically with $M_{\rm bar}$, rising from ≈ 0.6–0.7×10^{57} at $1.36\,M_\odot$ to ≈ 0.9–1.0×10^{57} at $1.93\,M_\odot$. Within each $Y_{e,i}$ slice, the EOS spread is ~ 7–8%, with LS220 sitting lowest because it predicts the highest cold-NS $Y_{e,f} \approx 0.085$.

Because the $Y_{e,i}$ uncertainty band is wider than the EOS spread at fixed mass, the two effects are entangled, and the consequence for physical inference is fundamental. A measurement of $N_{\rm delep}$ directly constrains the combination $M_{\rm bar}(Y_{e,i} - Y_{e,f})$, not $M_{\rm bar}$ or $Y_{e,f}$ individually. If $Y_{e,i}$ can be estimated from progenitor modeling, one can extract $Y_{e,f}$ and probe the nuclear symmetry energy. Without it, one must marginalize over the full range of $Y_{e,i}$, and the EOS constraint is significantly diluted compared to what $E_{\rm tot}$ achieves. This is the key asymmetry between the two observables: $E_{\rm tot}$ is a clean mass probe with EOS as the dominant residual systematic, while $N_{\rm delep}$ is sensitive to both initial and final conditions.

We see that the gap between the simulations and the upper edge of the band is considerably

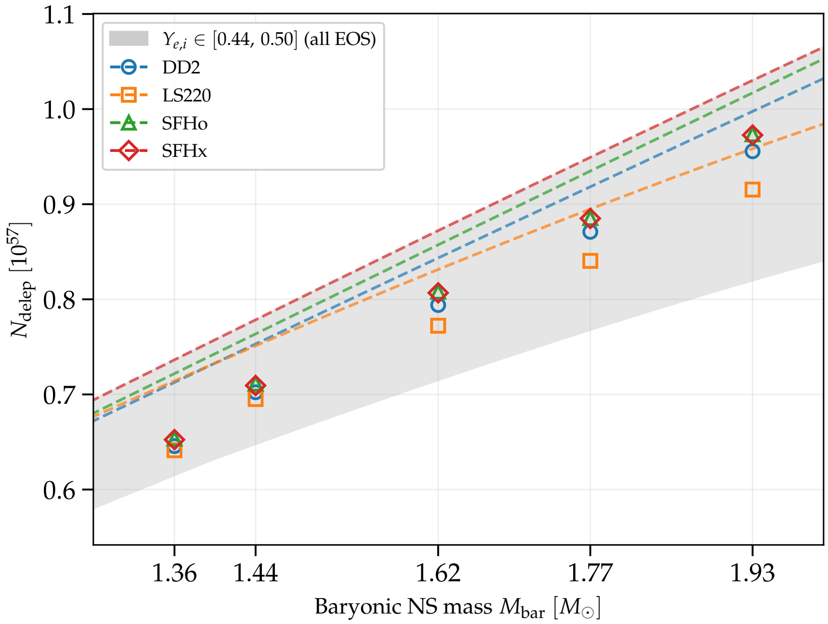

Figure 5.16: Net deleptonization number $N_{\text{delep}} = N_{\nu_e} - N_{\bar{\nu}_e}$ vs. baryonic NS mass M_{bar} for the 20 Garching models (five masses $\times$ four EOS). The gray band shows the asymptotic cold-NS value $N_{\text{bar}}(Y_{e,i} - Y_{e,f})$ from the TOV equation, spanning $Y_{e,i} \in [0.44, 0.50]$ and all four EOS, reflecting the uncertainty in the initial electron fraction of the PNS-forming material. Colored dashed curves show the upper edge ($Y_{e,i} = 0.50$) for each EOS individually. Open symbols show the simulated N_{delep} from each model, all falling within the band. The gap between the simulations and the upper edge is larger than in Fig. 5.14, reflecting the longer deleptonization timescale compared to thermal cooling.

larger than the corresponding gap in Fig. 5.14: the simulations lie near the bottom of the uncertainty band rather than close to the upper edge. This reflects a second effect independent of the $Y_{e,i}$ uncertainty: deleptonization takes longer than thermal cooling. The trapped ν_e excess must diffuse outward against a steep electron chemical potential gradient through the dense, neutron-rich core, and the deleptonization timescale substantially exceeds the Kelvin–Helmholtz cooling timescale. The Garching simulations capture the bulk of the thermal emission but only a fraction of the total deleptonization; the remaining gap is the lepton number still trapped in the PNS at simulation end. For an actual galactic core-collapse supernova, detectors would record neutrino emission over a longer duration, capturing a larger fraction of the deleptonization tail; the reconstructed N_{delep} would therefore exceed the simulation values in Fig. 5.16 and approach the asymptotic TOV prediction more closely. The simulation points are lower bounds on the total emitted N_{delep}, and should not be compared directly with $N_{\mathrm{bar}}(Y_{e,i} - Y_{e,f})$ without accounting for the incomplete deleptonization at simulation end.

Unlike E_{tot}, which is invariant under neutrino flavor transformation (Sec. 5.8.1), $N_{\mathrm{delep}} = N_{\nu_e} - N_{\bar{\nu}_e}$ is sensitive to oscillation physics. Flavor oscillations conserve the total neutrino energy but not the individual flavor lepton numbers: if ν_e converts to ν_x en route to Earth, N_{ν_e} detected at Earth is reduced while N_{ν_x} increases, shifting the difference $N_{\nu_e} - N_{\bar{\nu}_e}$ measured by our detectors. The magnitude of the shift depends on the oscillation scenario — MSW matter effect, collective oscillations, shock-induced transformations — and on the neutrino mass hierarchy, introducing an additional theoretical uncertainty on top of the $Y_{e,i}$ degeneracy already discussed.

The combination of large reconstruction uncertainty and irreducible $Y_{e,i}$ degeneracy means that N_{delep} cannot currently serve as a quantitative probe of PNS mass or EOS in the way that E_{tot} can. Its role is confirmatory: a robust detection of $N_{\mathrm{delep}} > 0$ directly confirms that the stellar core underwent neutronization, as demonstrated in Sec. 5.7.3. Quantitative mass inference from N_{delep} would require both a substantial reduction in the ν_e reconstruction bias and an independent estimate of $Y_{e,i}$ from progenitor modeling.

5.9 Conclusions

When the next Milky Way core-collapse supernova occurs, detectors around the world will record more than 10^4 neutrino events — orders of magnitude beyond the 19 events from SN 1987A. Making the most of this singular opportunity requires measuring the full-flavor neutrino emission without imposing assumptions about its spectral shape. Here we have shown that B-spline spectral unfolding, applied jointly across five complementary detection channels, provides a powerful and model-independent framework for recovering the two leading-order observables of the explosion: the total emitted energy $E_{\rm tot}$ and the net deleptonization number $N_{\rm delep} \equiv N_{\nu_e} - N_{\bar{\nu}_e}$.

Across all 20 Garching models spanning five baryonic masses and four nuclear equations of state, we find that $E_{\rm tot}$ is recovered with bias $\lesssim 2\%$ and a 68% width of ≈ 2.5–5.5%, with no systematic trend in progenitor mass or EOS. $N_{\rm delep} > 0$ is robustly established in all models, directly confirming core neutronization, though its quantitative precision is limited by the small-difference amplification discussed in Sec. 5.7.3 and will improve with better treatment of the DUNE sub-threshold ν_e reconstruction. The tight $E_{\rm tot}$ reconstruction is sufficient to make it a direct mass observable: by inverting the TOV relation, we infer the PNS baryonic mass at ≈ 1–4% precision from a single detection, competitive with or better than existing neutron-star mass measurements from X-ray pulse-profile modeling (Miller et al., 2019; Riley et al., 2019) and gravitational-wave observations (Abbott et al., 2017a,b), while probing the remnant at the moment of its birth.

DUNE is indispensable. Without a dedicated ν_e anchor, the ν_e and ν_x contributions to the flavor-mixed JUNO and HK signals cannot be independently resolved, the $E_{\rm tot}$ bias worsens from -0.7% to -11%, and it becomes impossible to robustly establish $N_{\rm delep} > 0$. DUNE's MeV capabilities — energy resolution, detection threshold, and background rejection — are a direct input to the science return of the next galactic supernova, and investments in improving them will pay dividends.

The most important extensions of this work are incorporating realistic neutrino oscillation effects — which affect $N_{\rm delep}$ but not $E_{\rm tot}$, for reasons discussed in Sec. 5.8 — and extending the framework to time-resolved reconstruction, which would allow the cooling and deleptonization evolution of the PNS to be tracked directly. The next Milky Way core collapse will likely occur within the next few decades. *We may have only one chance to extract the full physics from it; the analysis framework should be in place and validated before that* *happens.*

Appendix 5A: Obtaining ν_x at JUNO via IBU

The ν_x spectrum at JUNO is obtained in two steps: a low-energy extrapolation of the DUNE ν_e unfolding, followed by a flavor subtraction from the total JUNO signal.

Low-energy extrapolation of the DUNE ν_e spectrum

The finite detection threshold at DUNE ($E_{\text{det}} \geq 3$ MeV) translates to an effective source threshold of ~10 MeV, below which the ν_e spectrum cannot be reliably unfolded. To recover the missing low-energy contribution, we linearly extrapolate each Poisson realization of the unfolded spectrum from its last reconstructed bin down to zero flux at zero energy. The result is shown in Figure 5.17. This extrapolation improves the integrated recovery from $\sim 72\%$ to $\sim 107\%$ of the true neutrino count and from $\sim 88\%$ to $\sim 103\%$ of the true emitted energy—both within 10% of the truth.

Flavor subtraction to isolate ν_x

Although JUNO alone achieves good integrated reconstruction, its signal contains contributions from all neutrino flavors. To isolate ν_x, we subtract the unfolded $\bar{\nu}_e$ spectrum from Super-K and the extrapolated ν_e spectrum from DUNE from the total JUNO unfolded signal. The resulting ν_x reconstruction is shown in Figure 5.18.

At energies below ~10 MeV, the subtraction becomes unreliable because the DUNE ν_e contribution in this region is itself extrapolated rather than directly unfolded; the ν_x uncertainties grow accordingly. At higher energies, the reconstructed spectrum tracks the true distribution well, though with sizable uncertainties inherited from the individual detector unfoldings. Integrated over the full spectrum, the subtraction recovers $\sim 110\%$ of the true ν_x event count and $\sim 103\%$ of the true emitted energy—within $\sim 10\%$ and $\sim 3\%$ of truth, respectively.

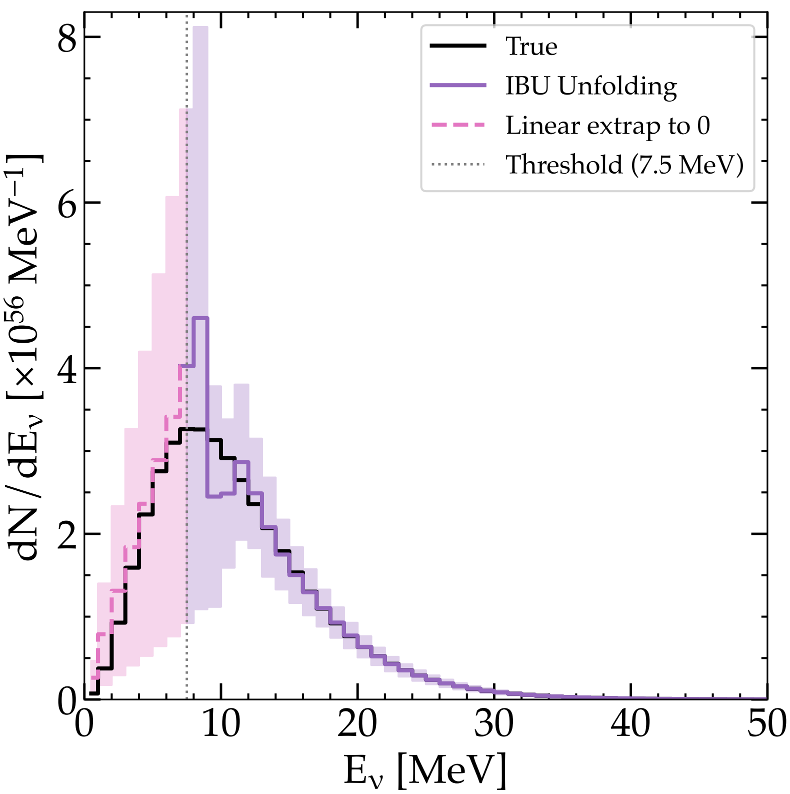

Figure 5.17: Unfolded ν_e spectrum from DUNE (purple; same as Figure 5.9) and the linearly extrapolated result (pink), obtained by extending each Poisson realization from the last reconstructed bin to zero flux at zero energy. The extrapolation recovers the true integrated count and energy to within 10% and 3%, respectively.

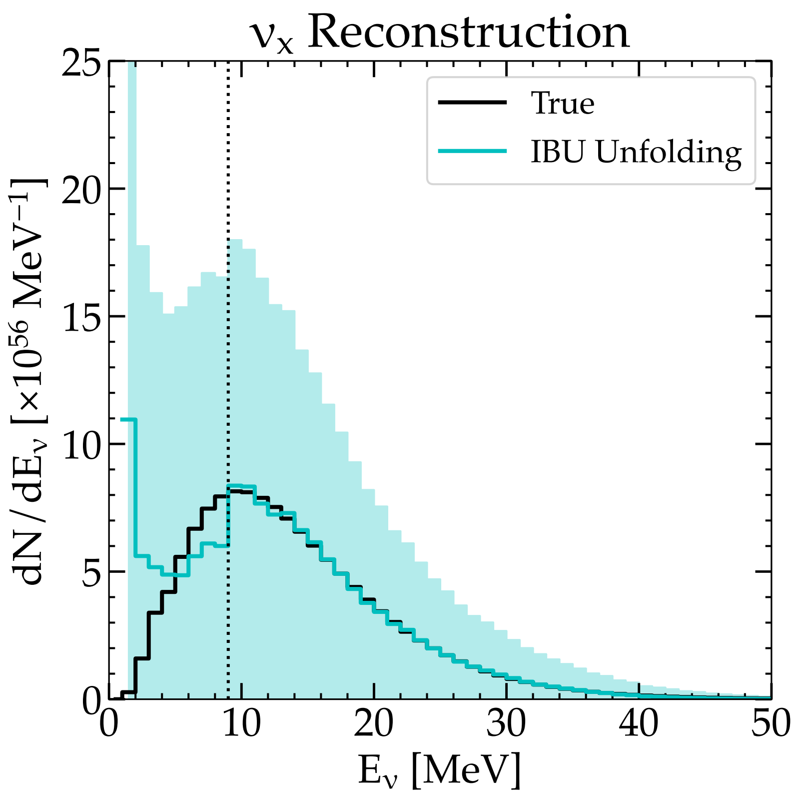

Figure 5.18: Reconstructed ν_x spectrum at JUNO, obtained by subtracting the unfolded $\bar{\nu}_e$ (Super-K) and ν_e (DUNE) contributions from the total JUNO signal. The black curve is the true spectrum; the colored curve and shaded band show the IBU mean and 1σ spread from repeated Poisson realizations. The dotted vertical line marks the DUNE unfolding threshold below which ν_x uncertainties grow due to reliance on the low-energy extrapolation.

Convergence Tolerance

As noted in Section 5.6, the IBU stopping criterion is based on the relative change $\delta^{(k)} = \|n^{(k)} - n^{(k-1)}\|/\|n^{(k-1)}\|$, with a tolerance of 6×10^{-3}. Figure 5.19 illustrates the effect of tightening this threshold to $\delta^{(k)} < 10^{-3}$ for the JUNO channel. While the tighter criterion reduces the toy-to-toy variance — visible as a narrower shaded band — it does so at the cost of spectral shape fidelity: the reconstructed mean shifts away from the true spectrum, particularly at low energies. This behavior is a manifestation of the underdetermined nature of the JUNO inversion; additional iterations beyond the chosen stopping point amplify noise rather than refine the solution. The tolerance of 6×10^{-3} is therefore preferred as it better preserves the reconstructed spectral shape while keeping statistical uncertainties controlled.

Appendix 5B: Sweep validation: E_{tot} reconstruction across all models

Table 5.1 lists the simulated E_{tot}, reconstruction bias, and 68% half-width for all 20 Garching models from the validation sweep of Sec. 5.7.5. Table 5.2 gives the corresponding N_{delep} results. Bias and width are expressed as a percentage of the true value.

Appendix 5C: Single-Channel B-spline Reconstruction

Figure 5.20 shows the standalone B-spline reconstruction for the two single-flavor dedicated channels: Super-Kamiokande IBD ($\bar{\nu}_e$, left) and DUNE CC (ν_e, right).

The SK IBD case is the cleanest application of the method: a single flavor, a near-diagonal response matrix, and $\simeq$6500 detected events. The system is highly overconstrained, requiring little regularization and no sub-threshold extrapolation. The B-spline median tracks the true spectrum closely across the full energy range, with a narrow 68% band. The recovered total

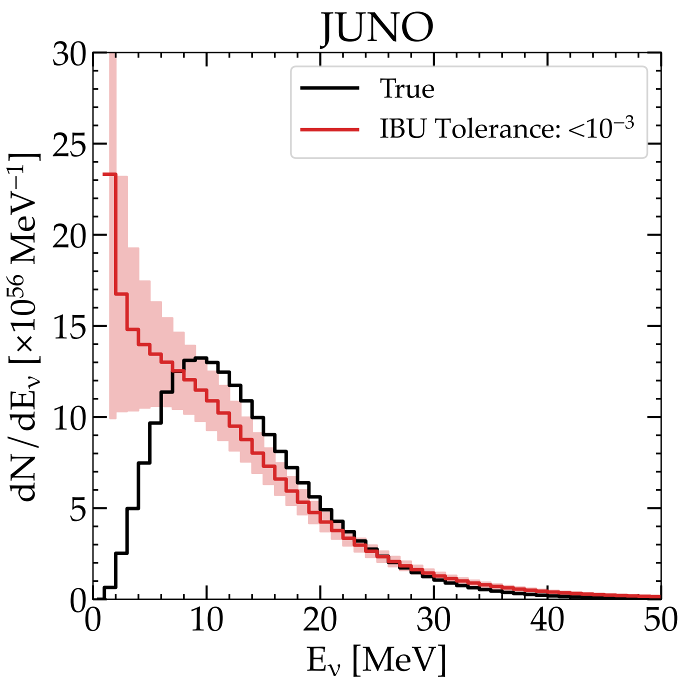

Figure 5.19: Reconstructed JUNO neutrino energy spectrum under a tightened IBU stopping tolerance, $\delta^{(k)} < 10^{-3}$, for the 1.93 $M_{\odot}$, SFHx progenitor model. The black curve shows the true spectrum; the red curve and shaded band show the mean and spread of reconstructed spectra across 10,000 Poisson realizations. Relative to the fiducial tolerance of 6×10^{-3}, the tighter criterion narrows the reconstructed uncertainty band but degrades spectral shape fidelity, producing a pronounced excess at low energies ($E_\nu \lesssim 10$ MeV) where additional IBU iterations amplify statistical noise rather than refine the solution.

Table 5.1: E_{tot} reconstruction performance for all 20 Garching models (five baryonic masses $\times$ four EOS), from 10^4 Poisson toys per model using the five-channel B-spline fit. Bias $= (\tilde{E} - E_{\text{true}})/E_{\text{true}}$ where $\tilde{E}$ is the median toy value. Width = 68% half-width/E_{true}.

M_{bar} [$M_\odot$]	EOS	E_{tot} [10^{53} erg]	Bias [%]	Width [%]
1.36	DD2	1.803	−0.4	5.5
	LS220	1.872	−0.9	5.4
	SFHo	1.963	−0.7	5.1
	SFHx	1.971	−0.4	5.0
1.44	DD2	2.053	0.0	4.8
	LS220	2.154	−0.8	4.7
	SFHo	2.244	−0.4	4.6
	SFHx	2.251	−0.8	4.6
1.62	DD2	2.625	−0.4	3.8
	LS220	2.726	−0.9	3.6
	SFHo	2.893	−0.5	3.3
	SFHx	2.913	−0.6	3.4
1.77	DD2	3.144	−0.5	3.2
	LS220	3.283	−1.1	3.1
	SFHo	3.483	−0.7	2.9
	SFHx	3.514	−0.5	2.9
1.93	DD2	3.758	−0.5	2.6
	LS220	3.962	−1.2	2.7
	SFHo	4.195	−0.6	2.3
	SFHx	4.252	−0.8	2.3

energy has a bias of +0.3% and a width of 1.5%; the recovered total count has a bias of +0.9% and a width of 3.8%.

The DUNE CC case is more challenging. Above the effective source threshold ($\sim$7 MeV), the median closely tracks the true spectrum with a 68% band comparable in width to the SK IBD case. Below threshold, the spline extrapolates using only the smoothness prior and positivity constraint, with no detected-event leverage. The recovered total energy has a bias of +3.9% and a width of 3.4%. The total count bias is larger (+14%), reflecting the difficulty of extrapolating the $\sim$27% of ν_e flux that lies below threshold. In the five-channel joint fit, the other channels help constrain the spectral shape and reduce this bias substantially.

Table 5.2: N_{delep} reconstruction performance for all 20 Garching models, from 10^4 Poisson toys per model using the five-channel B-spline fit. Bias $= (\tilde{N} - N_{\text{true}})/N_{\text{true}}$ where $\tilde{N}$ is the median toy value. Width = 68% half-width/N_{true}.

M_{bar} [$M_\odot$]	EOS	N_{delep} [10^{57}]	Bias [%]	Width [%]
1.36	DD2	0.646	+29.8	27.6
	LS220	0.641	+32.7	29.7
	SFHo	0.653	+28.8	27.3
	SFHx	0.652	+28.3	26.2
1.44	DD2	0.702	+27.7	24.0
	LS220	0.695	+31.0	27.6
	SFHo	0.711	+27.6	23.9
	SFHx	0.709	+28.4	25.9
1.62	DD2	0.794	+27.0	18.8
	LS220	0.772	+28.6	19.5
	SFHo	0.808	+23.3	17.5
	SFHx	0.807	+21.5	17.5
1.77	DD2	0.871	+22.8	16.1
	LS220	0.840	+29.5	15.4
	SFHo	0.885	+26.6	14.8
	SFHx	0.885	+24.4	15.3
1.93	DD2	0.956	+25.5	13.0
	LS220	0.915	+28.6	15.9
	SFHo	0.974	+25.7	14.2
	SFHx	0.973	+27.6	13.6

Appendix 5D: Spectral Reconstruction without DUNE

Figure 5.21 shows the per-flavor reconstructed spectra from the four-channel fit with DUNE removed, providing the spectral view of the degradation quantified in Sec. 5.7.4. The three panels tell a clear story. The $\bar{\nu}_e$ reconstruction (left) is completely unaffected: Super-K IBD is self-sufficient, and removing DUNE leaves it unchanged. The ν_e reconstruction (center) is catastrophically inflated, with the median spectrum roughly twice the true flux across the full energy range and a very wide uncertainty band—the direct spectral manifestation of the event mis-attribution described in Sec. 5.7.4. The ν_x reconstruction (right) is correspondingly suppressed, the median lying well below truth with a large uncertainty band especially at low energies. The contrast between the left panel and the center and right panels makes vivid

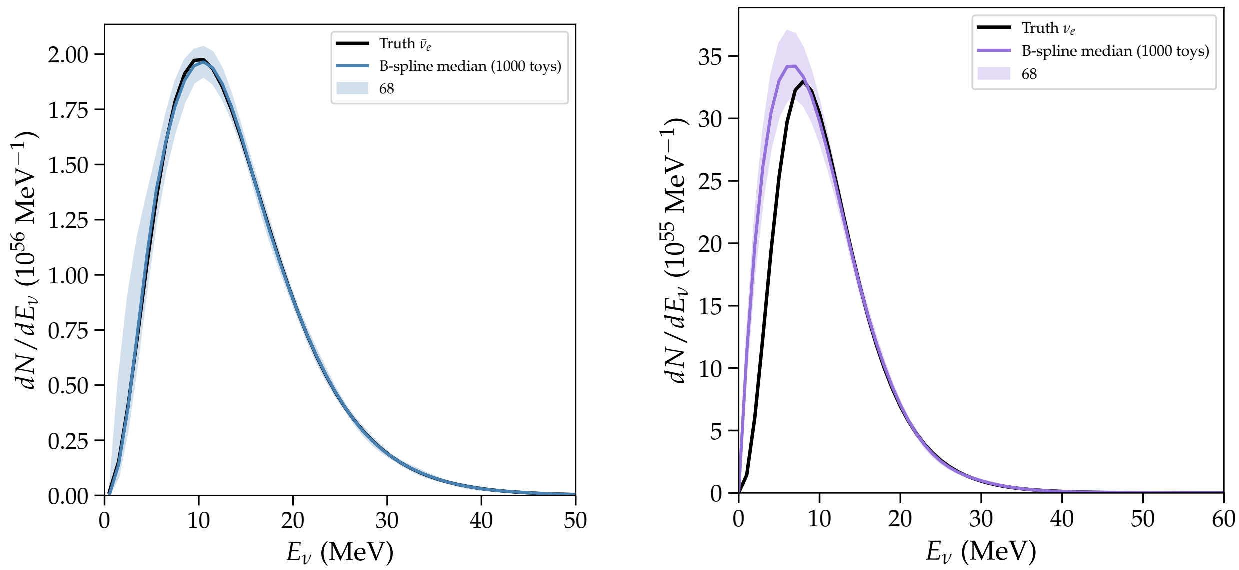

Figure 5.20: B-spline reconstruction of the $\bar{\nu}_e$ spectrum from Super-Kamiokande IBD alone (left) and the ν_e spectrum from DUNE CC alone (right), both at 10 kpc. In each panel, the black curve shows the true spectrum; the colored curve and shaded band show the median and 68% interval from 10^4 Poisson toys. The SK IBD reconstruction requires no sub-threshold extrapolation; the DUNE spline extrapolates below the effective source threshold ($\sim$7 MeV) using only the smoothness prior.

the irreplaceable role of DUNE: $\bar{\nu}_e$ needs only Super-K, but ν_e and ν_x cannot be disentangled without it.

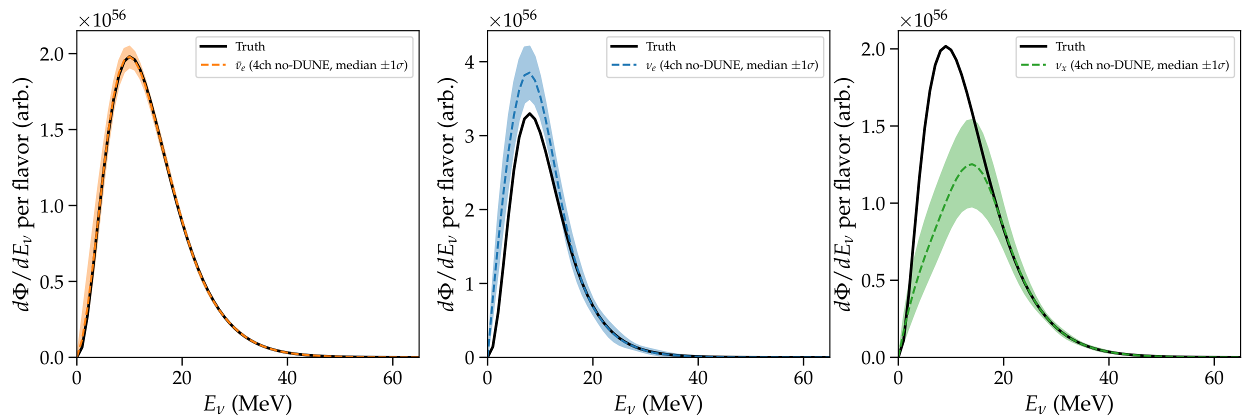

Figure 5.21: B-spline reconstruction of the $\bar{\nu}_e$ (left), ν_e (center), and per-flavor ν_x (right) spectra from the four-channel fit with DUNE removed, based on 10^4 Poisson toys. The black curve shows the true spectrum; colored dashed curves and shaded bands show the median and 68% interval. The $\bar{\nu}_e$ reconstruction is unaffected by the removal of DUNE. The ν_e spectrum is catastrophically inflated—roughly twice the true flux—because the optimizer over-attributes JUNO events to ν_e with no dedicated anchor to prevent it. The ν_x spectrum is correspondingly suppressed, its events absorbed by the unconstrained ν_e spline. This spectral breakdown directly causes the degraded E_{tot} and N_{delep} reconstruction shown in Fig. 5.13.

Bibliography

Abac, A. G., et al. 2024. https://arxiv.org/abs/2404.04248

Abac, A. G., Abouelfettouh, I., Acernese, F., et al. 2025, arXiv e-prints, arXiv:2508.18083, doi:10.48550/arXiv.2508.18083

Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016, Physical Review Letters, 116, 1, doi:10.1103/PhysRevLett.116.061102

Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017a, Phys. Rev. Lett., 119, 161101, doi:10.1103/PhysRevLett.119.161101

—. 2017b, ApJL, 850, L39, doi:10.3847/2041-8213/aa9478

—. 2019a, Physical Review X, 9, 031040, doi:10.1103/PhysRevX.9.031040

—. 2019b, ApJL, 882, L24, doi:10.3847/2041-8213/ab3800

Abbott, B. P., et al. 2020, Living Rev. Rel., 23, 3, doi:10.1007/s41114-020-00026-9

Abbott, D. C. 1982, ApJ, 259, 282, doi:10.1086/160166

Abbott, R., Abbott, T. D., Abraham, S., et al. 2020, ApJL, 896, L44, doi:10.3847/2041-8213/ab960f

Abbott, R., Abbott, T. D., Abraham, S., et al. 2020a, Phys. Rev. Lett., 125, 101102, doi:10.1103/PhysRevLett.125.101102

—. 2020b, The Astrophysical Journal, 900, L13, doi:10.3847/2041-8213/aba493

Abbott, R., Abbott, T. D., Acernese, F., et al. 2021a, arXiv e-prints, arXiv:2111.03606. https://arxiv.org/abs/2111.03606

Abbott, R., Abbott, T. D., Abraham, S., et al. 2021b, ApJL, 915, L5, doi:10.3847/2041-8213/ac082e

Abbott, R., Abbott, T. D., Abraham, S., et al. 2021, Physical Review X, 11, 021053, doi:10.1103/PhysRevX.11.021053

Abbott, R., Abbott, T. D., Acernese, F., et al. 2023, Physical Review X, 13, 011048, doi:10.1103/PhysRevX.13.011048

Abe, K., et al. 2024, Phys. Rev. D, 109, 092001, doi:10.1103/PhysRevD.109.092001

Abi, B., et al. 2020. `https://arxiv.org/abs/2002.03005`

—. 2021, Eur. Phys. J. C, 81, 423, doi:10.1140/epjc/s10052-021-09166-w

Abusleme, A., et al. 2025a. `https://arxiv.org/abs/2511.14590`

—. 2025b. `https://arxiv.org/abs/2511.14593`

Adame, A. G., Aguilar, J., Ahlen, S., et al. 2025, JCAP, 2025, 021, doi:10.1088/1475-7516/2025/02/021

Adams, S. M., Kochanek, C. S., Beacom, J. F., Vagins, M. R., & Stanek, K. Z. 2013, ApJ, 778, 164, doi:10.1088/0004-637X/778/2/164

Adams, S. M., Kochanek, C. S., Beacom, J. F., Vagins, M. R., & Stanek, K. Z. 2013, Astrophys. J., 778, 164, doi:10.1088/0004-637X/778/2/164

Adhikari, S., Banerjee, A., Boddy, K. K., et al. 2025, Reviews of Modern Physics, 97, 045004, doi:10.1103/m2vm-59y3

Afroz, S., & Mukherjee, S. 2024. `https://arxiv.org/abs/2411.07304`

Aghanim, N., et al. 2020, Astron. Astrophys., 641, A6, doi:10.1051/0004-6361/201833910

Agrawal, P., Hurley, J., Stevenson, S., et al. 2023, MNRAS, 525, 933, doi:10.1093/mnras/stad2334

Agrawal, P., Hurley, J., Stevenson, S., Szécsi, D., & Flynn, C. 2020, MNRAS, 497, 4549, doi:10.1093/mnras/staa2264

Ahmad, Q. R., Allen, R. C., Andersen, T. C., et al. 2002, Phys. Rev. Lett., 89, 011301, doi:10.1103/PhysRevLett.89.011301

Akins, H. B., Casey, C. M., Lambrides, E., et al. 2025, ApJ, 991, 37, doi:10.3847/1538-4357/ade984

Alekseev, E. N., Alekseeva, L. N., Krivosheina, I. V., & Volchenko, V. I. 1987, in European Southern Observatory Conference and Workshop Proceedings, Vol. 26, European Southern Observatory Conference and Workshop Proceedings, ed. I. J. Danziger, 237

An, F., et al. 2016, J. Phys. G, 43, 030401, doi:10.1088/0954-3899/43/3/030401

Antonini, F., Gieles, M., Dosopoulou, F., & Chattopadhyay, D. 2022. `https://arxiv.org/abs/2208.01081`

Arnett, W. D. 1966, Canadian Journal of Physics, 44, 2553, doi:10.1139/p66-210

Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, A&A, 558, A33, doi:10.1051/0004-6361/201322068

Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, AJ, 156, 123, doi:10.3847/1538-3881/aabc4f

Astropy Collaboration, Price-Whelan, A. M., Lim, P. L., et al. 2022, ApJ, 935, 167, doi:10.3847/1538-4357/ac7c74

Bahcall, N. A., Ostriker, J. P., Perlmutter, S., & Steinhardt, P. J. 1999, Science, 284, 1481, doi:10.1126/science.284.5419.1481

Baibhav, V., Berti, E., Gerosa, D., et al. 2019, Phys. Rev. D, 100, 064060, doi:10.1103/PhysRevD.100.064060

Balberg, S., Shapiro, S. L., & Inagaki, S. 2002, Astrophys. J., 568, 475, doi:10.1086/339038

Baldry, I. K., Driver, S. P., Loveday, J., et al. 2012, Monthly Notices of the Royal Astronomical Society, 421, 621, doi:10.1111/j.1365-2966.2012.20340.x

Banerjee, S. 2022, Astron. Astrophys., 665, A20, doi:10.1051/0004-6361/202142331

Banerjee, S., Baumgardt, H., & Kroupa, P. 2010, MNRAS, 402, 371, doi:10.1111/j.1365-2966.2009.15880.x

Banerjee, S., & Olejak, A. 2024, arXiv e-prints, arXiv:2411.15112, doi:10.48550/arXiv.2411.15112

Barkat, Z., Rakavy, G., & Sack, N. 1967, Physical Review Letters, 18, 379, doi:10.1103/PhysRevLett.18.379

Barrett, J. W., Gaebel, S. M., Neijssel, C. J., et al. 2018, Mon. Not. Roy. Astron. Soc., 477, 4685, doi:10.1093/mnras/sty908

Bastian, N., Covey, K. R., & Meyer, M. R. 2010, ARA&A, 48, 339, doi:10.1146/annurev-astro-082708-101642

Bavera, S. S., Fishbach, M., Zevin, M., Zapartas, E., & Fragos, T. 2022, A&A, 665, A59, doi:10.1051/0004-6361/202243724

Bavera, S. S., Fragos, T., Qin, Y., et al. 2020, Astron. Astrophys., 635, A97, doi:10.1051/0004-6361/201936204

Bavera, S. S., Fragos, T., Zevin, M., et al. 2021, A&A, 647, A153, doi:10.1051/0004-6361/202039804

Behroozi, P., Wechsler, R. H., Hearin, A. P., & Conroy, C. 2019, MNRAS, 488, 3143, doi:10.1093/mnras/stz1182

Behroozi, P. S., Wechsler, R. H., & Wu, H.-Y. 2012a, The Astrophysical Journal, 762, 109, doi:10.1088/0004-637x/762/2/109

Behroozi, P. S., Wechsler, R. H., Wu, H.-Y., et al. 2012b, The Astrophysical Journal, 763, 18, doi:10.1088/0004-637x/763/1/18

Belczynski, K., Bulik, T., Fryer, C. L., et al. 2010a, ApJ, 714, 1217, doi:10.1088/0004-637X/714/2/1217

Belczynski, K., Dominik, M., Bulik, T., et al. 2010b, ApJL, 715, L138, doi:10.1088/2041-8205/715/2/L138

Belczynski, K., Heger, A., Gladysz, W., et al. 2016, Astronomy and Astrophysics, 594, 1, doi:10.1051/0004-6361/201628980

Belczynski, K., Klencki, J., Fields, C. E., et al. 2020, A&A, 636, A104, doi:10.1051/0004-6361/201936528

Belczynski, K., Romagnolo, A., Olejak, A., et al. 2022, ApJ, 925, 69, doi:10.3847/1538-4357/ac375a

Bestenlehner, J. M., Gräfener, G., Vink, J. S., et al. 2014, Astronomy & Astrophysics, 570, A38

Bestenlehner, J. M., Crowther, P. A., Caballero-Nieves, S. M., et al. 2020, Monthly Notices of the Royal Astronomical Society, 499, 1918

Bethe, H. A., & Brown, G. E. 1998, Astrophys. J., 506, 780, doi:10.1086/306265

Bethe, H. A., & Wilson, J. R. 1985, ApJ, 295, 14, doi:10.1086/163343

Bhattacharya, M., Goodman, C. D., & Garcia, A. 2009, Phys. Rev. C, 80, 055501, doi:10.1103/PhysRevC.80.055501

Bhattacharyya, J., Adhikari, S., Banerjee, A., et al. 2022, ApJ, 932, 30, doi:10.3847/1538-4357/ac68e9

Bionta, R. M., Blewitt, G., Bratton, C. B., et al. 1987, in The Standard Model. The Supernova 1987A, ed. J. Tran Thanh Van, 735–738

Bionta, R. M., et al. 1987, Phys. Rev. Lett., 58, 1494, doi:10.1103/PhysRevLett.58.1494

Blumenthal, G. R., Faber, S. M., Primack, J. R., & Rees, M. J. 1984, Nature, 311, 517, doi:10.1038/311517a0

Boco, L., Mapelli, M., Sander, A. A. C., et al. 2025, A&A, 703, A243, doi:10.1051/0004-6361/202556187

Boddy, K. K., Feng, J. L., Kaplinghat, M., Shadmi, Y., & Tait, T. M. P. 2014, Phys. Rev. D, 90, 095016, doi:10.1103/PhysRevD.90.095016

Boesky, A., Broekgaarden, F. S., & Berger, E. 2024a. `https://arxiv.org/abs/2405.01630`

—. 2024b. `https://arxiv.org/abs/2405.01623`

Bombaci, I. 1996, A&A, 305, 871

Bondi, H. 1952, MNRAS, 112, 195, doi:10.1093/mnras/112.2.195

Botticella, M. T., Riello, M., Cappellaro, E., et al. 2008, A&A, 479, 49, doi:10.1051/0004-6361:20078011

Bouffanais, Y., Mapelli, M., Santoliquido, F., et al. 2021, Mon. Not. Roy. Astron. Soc., 507, 5224, doi:10.1093/mnras/stab2438

Boylan-Kolchin, M. 2023, Nature Astronomy, 7, 731, doi:10.1038/s41550-023-01937-7

Boylan-Kolchin, M., Bullock, J. S., & Kaplinghat, M. 2011, MNRAS, 415, L40, doi:10.1111/j.1745-3933.2011.01074.x

—. 2012, MNRAS, 422, 1203, doi:10.1111/j.1365-2966.2012.20695.x

Brainerd, T. G., Blandford, R. D., & Smail, I. 1996, ApJ, 466, 623, doi:10.1086/177537

Bratton, C. B., et al. 1988, Phys. Rev. D, 37, 3361, doi:10.1103/PhysRevD.37.3361

Bressan, A., Marigo, P., Girardi, L., et al. 2012, Monthly Notices of the Royal Astronomical Society, 427, 127, doi:10.1111/j.1365-2966.2012.21948.x

Briel, M. M., Eldridge, J. J., Stanway, E. R., Stevance, H. F., & Chrimes, A. A. 2022, Mon. Not. Roy. Astron. Soc., 514, 1315, doi:10.1093/mnras/stac1100

Brinchmann, J., Charlot, S., White, S. D. M., et al. 2004, MNRAS, 351, 1151, doi:10.1111/j.1365-2966.2004.07881.x

Broekgaarden, F. S. 2021, BHBH simulations from: Impact of Massive Binary Star and Cosmic Evolution on Gravitational Wave Observations II: Double Compact Object Mergers, 1, Zenodo, doi:10.5281/zenodo.5651073

Broekgaarden, F. S., Banagiri, S., & Payne, E. 2024, Astrophys. J., 969, 108, doi:10.3847/1538-4357/ad4709

Broekgaarden, F. S., & Berger, E. 2021, Astrophys. J. Lett., 920, L13, doi:10.3847/2041-8213/ac2832

Broekgaarden, F. S., Stevenson, S., & Thrane, E. 2022, arXiv e-prints, arXiv:2205.01693. `https://arxiv.org/abs/2205.01693`

Broekgaarden, F. S., Stevenson, S., & Thrane, E. 2022, Astrophys. J., 938, 45, doi:10.3847/1538-4357/ac8879

Broekgaarden, F. S., Justham, S., de Mink, S. E., et al. 2019, Mon. Not. Roy. Astron. Soc., 490, 5228, doi:10.1093/mnras/stz2558

Broekgaarden, F. S., Berger, E., Neijssel, C. J., et al. 2021, MNRAS, 508, 5028, doi:10.1093/mnras/stab2716

Broekgaarden, F. S., Berger, E., Stevenson, S., et al. 2022, MNRAS, 516, 5737, doi:10.1093/mnras/stac1677

Brooks, A. M., & Zolotov, A. 2014, ApJ, 786, 87, doi:10.1088/0004-637X/786/2/87

Brown, L. M. 1978, Physics Today, 31, 23, doi:10.1063/1.2995181

Bruel, T., Lamberts, A., Rodriguez, C. L., et al. 2025. https://arxiv.org/abs/2503.03810

Budding, E. 1986, Ap&SS, 118, 241, doi:10.1007/BF00651132

Burbidge, E. M., Burbidge, G. R., Fowler, W. A., & Hoyle, F. 1957, Reviews of Modern Physics, 29, 547, doi:10.1103/RevModPhys.29.547

Burrows, A., & Goshy, J. 1993, ApJL, 416, L75, doi:10.1086/187074

Burrows, A., & Vartanyan, D. 2021, Nature, 589, 29, doi:10.1038/s41586-020-03059-w

Byrd, R. H., Lu, P., Nocedal, J., & Zhu, C. 1995, SIAM J. Sci. Comput., 16, 1190, doi:10.1137/0916069

Callister, T. A. 2024. https://arxiv.org/abs/2410.19145

Cameron, A. G. W. 1957, PASP, 69, 201, doi:10.1086/127051

Carignan, C., & Freeman, K. C. 1988, ApJL, 332, L33, doi:10.1086/185260

Carniani, S., Hainline, K., D'Eugenio, F., et al. 2024, Nature, 633, 318–322, doi:10.1038/s41586-024-07860-9

Castellano, M., Fontana, A., Treu, T., et al. 2022, The Astrophysical Journal Letters, 938, L15, doi:10.3847/2041-8213/ac94d0

Castellano, M., et al. 2025. https://arxiv.org/abs/2504.05893

Castor, J. I., Abbott, D. C., & Klein, R. I. 1975, ApJ, 195, 157, doi:10.1086/153315

Chen, M.-H., Li, L.-X., Chen, Q.-H., Hu, R.-C., & Liang, E.-W. 2024, MNRAS, 529, 1154, doi:10.1093/mnras/stae475

Chen, Y., Bressan, A., Girardi, L., et al. 2015, Monthly Notices of the Royal Astronomical Society, 452, 1068, doi:10.1093/mnras/stv1281

Chowdhury, S. R., & Santra, D. 2024. https://arxiv.org/abs/2411.11902

Chruślińska, M. 2024, Annalen Phys., 536, 2200170, doi:10.1002/andp.202200170

Chruślińska, M., Jeřábková, T., Nelemans, G., & Yan, Z. 2020, A&A, 636, A10, doi:10.1051/0004-6361/202037688

Chruslinska, M., & Nelemans, G. 2019, MNRAS, 488, 5300, doi:10.1093/mnras/stz2057

Chruslinska, M., Nelemans, G., & Belczynski, K. 2019, Mon. Not. Roy. Astron. Soc., 482, 5012, doi:10.1093/mnras/sty3087

Chruślińska, M., Nelemans, G., Boco, L., & Lapi, A. 2021, MNRAS, 508, 4994, doi:10.1093/mnras/stab2690

Cirelli, M., Strumia, A., & Zupan, J. 2024, arXiv e-prints, arXiv:2406.01705, doi:10.48550/arXiv.2406.01705

Claeys, J. S. W., Pols, O. R., Izzard, R. G., Vink, J., & Verbunt, F. W. M. 2014, A&A, 563, A83, doi:10.1051/0004-6361/201322714

Clowe, D., Bradač, M., Gonzalez, A. H., et al. 2006, ApJL, 648, L109, doi:10.1086/508162

Cole, S., & Lacey, C. 1996, MNRAS, 281, 716, doi:10.1093/mnras/281.2.716

Cole, S., Percival, W. J., Peacock, J. A., et al. 2005, MNRAS, 362, 505, doi:10.1111/j.1365-2966.2005.09318.x

Colgate, S. A., & White, R. H. 1966, ApJ, 143, 626, doi:10.1086/148549

Conselice, C. J., et al. 2024. `https://arxiv.org/abs/2407.14973`

Cowan, Jr., C. L., Reines, F., Harrison, F. B., Kruse, H. W., & McGuire, A. D. 1956, Science, 124, 103, doi:10.1126/science.124.3212.103

Cowan, J. J., Sneden, C., Lawler, J. E., et al. 2021, Reviews of Modern Physics, 93, 015002, doi:10.1103/RevModPhys.93.015002

Crawford, J. A. 1955, ApJ, 121, 71, doi:10.1086/145965

Croon, D., McDermott, S. D., & Sakstein, J. 2020, Phys. Rev. D, 102, 115024, doi:10.1103/PhysRevD.102.115024

Croon, D., & Sakstein, J. 2024, Phys. Rev. D, 109, 103021, doi:10.1103/PhysRevD.109.103021

Crowther, P. A., Schnurr, O., Hirschi, R., et al. 2010, Mon. Not. Roy. Astron. Soc., 408, 731, doi:10.1111/j.1365-2966.2010.17167.x

Curti, M., Maiolino, R., Curtis-Lake, E., et al. 2024, A&A, 684, A75, doi:10.1051/0004-6361/202346698

Curtis-Lake, E., Carniani, S., Cameron, A., et al. 2023, Nature Astronomy, 7, 622, doi:10.1038/s41550-023-01918-w

D'Agostini, G. 1995, Nuclear Instruments and Methods in Physics Research A, 362, 487, doi:10.1016/0168-9002(95)00274-X

—. 2010, arXiv e-prints, arXiv:1010.0632, doi:10.48550/arXiv.1010.0632

Dasgupta, B., & Beacom, J. F. 2011, Phys. Rev. D, 83, 113006, doi:10.1103/PhysRevD.83.113006

Dave, R., Spergel, D. N., Steinhardt, P. J., & Wandelt, B. D. 2001, Astrophys. J., 547, 574, doi:10.1086/318417

de Blok, W. J. G., McGaugh, S. S., Bosma, A., & Rubin, V. C. 2001, ApJL, 552, L23, doi:10.1086/320262

De Luca, V., Desjacques, V., Franciolini, G., Pani, P., & Riotto, A. 2021, Phys. Rev. Lett., 126, 051101, doi:10.1103/PhysRevLett.126.051101

de Mink, S. E., Langer, N., Izzard, R. G., Sana, H., & de Koter, A. 2013, ApJ, 764, 166, doi:10.1088/0004-637X/764/2/166

de Mink, S. E., Pols, O. R., & Hilditch, R. W. 2007, A&A, 467, 1181, doi:10.1051/0004-6361:20067007

Dehghani, R., Salucci, P., & Ghaffarnejad, H. 2020, A&A, 643, A161, doi:10.1051/0004-6361/201937079

Di Cintio, A., Brook, C. B., Macciò, A. V., et al. 2014, MNRAS, 437, 415, doi:10.1093/mnras/stt1891

Diehl, R., et al. 2006, Nature, 439, 45, doi:10.1038/nature04364

Diemer, B. 2018, ApJS, 239, 35, doi:10.3847/1538-4365/aaee8c

Disberg, P., & Mandel, I. 2025, ApJL, 989, L8, doi:10.3847/2041-8213/adf286

Dominik, M., Belczynski, K., Fryer, C., et al. 2013, Astrophys. J., 779, 72, doi:10.1088/0004-637X/779/1/72

Dominik, M., Berti, E., O'Shaughnessy, R., et al. 2015, Astrophys. J., 806, 263, doi:10.1088/0004-637X/806/2/263

Donnan, C. T., McLure, R. J., Dunlop, J. S., et al. 2024, MNRAS, 533, 3222, doi:10.1093/mnras/stae2037

Driver, S. P., et al. 2018, Mon. Not. Roy. Astron. Soc., 475, 2891, doi:10.1093/mnras/stx2728

Drlica-Wagner, A., Bechtol, K., Rykoff, E. S., et al. 2015, ApJ, 813, 109, doi:10.1088/0004-637X/813/2/109

Dubinski, J., & Carlberg, R. G. 1991, ApJ, 378, 496, doi:10.1086/170451

Durrant, A., Crain, R. A., Lacey, C. G., et al. 2026, arXiv e-prints, arXiv:2607.16404, doi:10.48550/arXiv.2607.16404

Edelman, B., Farr, B., & Doctor, Z. 2023, Astrophys. J., 946, 16, doi:10.3847/1538-4357/acb5ed

Einstein, A. 1918, Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften, 154

Elbert, O. D., Bullock, J. S., & Kaplinghat, M. 2018, Monthly Notices of the Royal Astronomical Society, 473, 1186, doi:10.1093/mnras/stx1959

Erb, D. K., Shapley, A. E., Pettini, M., et al. 2006, ApJ, 644, 813, doi:10.1086/503623

Escudero, H. G., Kuo, J.-L., Keeley, R. E., & Abazajian, K. N. 2022, Phys. Rev. D, 106, 103517, doi:10.1103/PhysRevD.106.103517

Essig, R., McDermott, S. D., Yu, H.-B., & Zhong, Y.-M. 2019, Phys. Rev. Lett., 123, 121102, doi:10.1103/PhysRevLett.123.121102

Farah, A. M., Edelman, B., Zevin, M., et al. 2023. `https://arxiv.org/abs/2301.00834`

Farah, A. M., Fishbach, M., & Holz, D. E. 2024, Astrophys. J., 962, 69, doi:10.3847/1538-4357/ad0558

Farr, W. M., Sravan, N., Cantrell, A., et al. 2011, Astrophys. J., 741, 103, doi:10.1088/0004-637X/741/2/103

Farr, W. M., Stevenson, S., Coleman Miller, M., et al. 2017, Nature, 548, 426, doi:10.1038/nature23453

Farrell, E. J., Groh, J. H., Hirschi, R., et al. 2021, Mon. Not. Roy. Astron. Soc., 502, L40, doi:10.1093/mnrasl/slaa196

Feng, J. L., Kaplinghat, M., Tu, H., & Yu, H.-B. 2009, JCAP, 2009, 004, doi:10.1088/1475-7516/2009/07/004

Feng, W.-X., Yu, H.-B., & Zhong, Y.-M. 2021, ApJL, 914, L26, doi:10.3847/2041-8213/ac04b0

Feng, W.-X., Yu, H.-B., & Zhong, Y.-M. 2025. `https://arxiv.org/abs/2506.17641`

Fernández-García, E., Betancort-Rijo, J. E., Prada, F., et al. 2026, A&A, 707, L4, doi:10.1051/0004-6361/202558431

Fields, B. D., Molaro, P., & Sarkar, S. 2014, arXiv e-prints, arXiv:1412.1408, doi:10.48550/arXiv.1412.1408

Finkelstein, S. L., Bagley, M. B., Arrabal Haro, P., et al. 2022, ApJL, 940, L55, doi:10.3847/2041-8213/ac966e

Fiorillo, D. F. G., Heinlein, M., Janka, H.-T., et al. 2023, Phys. Rev. D, 108, 083040, doi:10.1103/PhysRevD.108.083040

Fishbach, M., & Kalogera, V. 2021, Astrophys. J. Lett., 914, L30, doi:10.3847/2041-8213/ac05c4

Fishbach, M., Doctor, Z., Callister, T., et al. 2021, Astrophys. J., 912, 98, doi:10.3847/1538-4357/abee11

Flores, R. A., & Primack, J. R. 1994, ApJL, 427, L1, doi:10.1086/187350

Fowler, W. A., & Hoyle, F. 1964, The Astrophysical Journal Supplement Series, 9, 201, doi:10.1086/190103

Frenk, C. S., & White, S. D. M. 2012, Annalen der Physik, 524, 507, doi:10.1002/andp.201200212

Fryer, C. L., Belczynski, K., Wiktorowicz, G., et al. 2012, Astrophys. J., 749, 91, doi:10.1088/0004-637X/749/1/91

Fryer, C. L., & Kalogera, V. 2001, ApJ, 554, 548, doi:10.1086/321359

Fryer, C. L., Woosley, S. E., & Heger, A. 2001, Population (English Edition), 20, 372

Fukuda, Y., Hayakawa, T., Ichihara, E., et al. 1998, Phys. Rev. Lett., 81, 1562, doi:10.1103/PhysRevLett.81.1562

Fuller, J., & Ma, L. 2019, ApJL, 881, L1, doi:10.3847/2041-8213/ab339b

Furlong, M., Bower, R. G., Theuns, T., et al. 2015, MNRAS, 450, 4486, doi:10.1093/mnras/stv852

Gallazzi, A., Charlot, S., Brinchmann, J., White, S. D., & Tremonti, C. A. 2005, Monthly Notices of the Royal Astronomical Society, 362, 41, doi:10.1111/j.1365-2966.2005.09321.x

García, F., Bunzel, A. S., Chaty, S., Porter, E., & Chassande-Mottin, E. 2021, Astron. Astrophys., 649, A114, doi:10.1051/0004-6361/202038357

Geha, M., Brown, T. M., Tumlinson, J., et al. 2013, ApJ, 771, 29, doi:10.1088/0004-637X/771/1/29

Gennaro, M., Tchernyshyov, K., Brown, T. M., et al. 2018, ApJ, 855, 20, doi:10.3847/1538-4357/aaa973

Gerosa, D., Berti, E., O'Shaughnessy, R., et al. 2018, Phys. Rev. D, 98, 084036, doi:10.1103/PhysRevD.98.084036

Gerosa, D., & Fishbach, M. 2021, Nature Astron., 5, 749, doi:10.1038/s41550-021-01398-w

Gerosa, D., Kesden, M., Berti, E., O'Shaughnessy, R., & Sperhake, U. 2013, Phys. Rev. D, 87, 104028, doi:10.1103/PhysRevD.87.104028

Giacobbo, N., & Mapelli, M. 2018, Mon. Not. Roy. Astron. Soc., 480, 2011, doi:10.1093/mnras/sty1999

Gommers, R., Virtanen, P., Haberland, M., et al. 2024, scipy/scipy: SciPy 1.13.1, v1.13.1, Zenodo, doi:10.5281/zenodo.11255513

Governato, F., Brook, C., Mayer, L., et al. 2010, Nature, 463, 203, doi:10.1038/nature08640

Greene, J. E., Labbe, I., Goulding, A. D., et al. 2024, ApJ, 964, 39, doi:10.3847/1538-4357/ad1e5f

Hahn, O., & Abel, T. 2011, Monthly Notices of the Royal Astronomical Society, 415, 2101–2121, doi:10.1111/j.1365-2966.2011.18820.x

Harikane, Y., Nakajima, K., Ouchi, M., et al. 2024, ApJ, 960, 56, doi:10.3847/1538-4357/ad0b7e

Harikane, Y., Zhang, Y., Nakajima, K., et al. 2023a, ApJ, 959, 39, doi:10.3847/1538-4357/ad029e

Harikane, Y., Ouchi, M., Oguri, M., et al. 2023b, ApJS, 265, 5, doi:10.3847/1538-4365/acaaa9

Harries, T. J., Hilditch, R. W., & Howarth, I. D. 2003, MNRAS, 339, 157, doi:10.1046/j.1365-8711.2003.06169.x

Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, Nature, 585, 357, doi:10.1038/s41586-020-2649-2

Heger, A., Fryer, C. L., Woosley, S. E., Langer, N., & Hartmann, D. H. 2003, Astrophys. J., 591, 288, doi:10.1086/375341

Hempel, M., & Schaffner-Bielich, J. 2010, Nucl. Phys. A, 837, 210, doi:10.1016/j.nuclphysa.2010.02.010

Hendriks, D. D., van Son, L. A. C., Renzo, M., Izzard, R. G., & Farmer, R. 2023, Mon. Not. Roy. Astron. Soc., 526, 4130, doi:10.1093/mnras/stad2857

Hendriks, D. D., van Son, L. A. C., Renzo, M., Izzard, R. G., & Farmer, R. 2023, MNRAS, 526, 4130, doi:10.1093/mnras/stad2857

Hinshaw, G., Larson, D., Komatsu, E., et al. 2013, ApJS, 208, 19, doi:10.1088/0067-0049/208/2/19

Hirata, K., Kajita, T., Koshiba, M., et al. 1987, in Neutrino Masses and Neutrino Astrophysics (Including Supernova 1987a), ed. V. Barger, F. Halzen, M. Marshak, & K. Olive, 51–62

Hirata, K., et al. 1987, Phys. Rev. Lett., 58, 1490, doi:10.1103/PhysRevLett.58.1490

Hirata, K. S., et al. 1988, Phys. Rev. D, 38, 448, doi:10.1103/PhysRevD.38.448

Hirschi, R., Goodman, K., Meynet, G., et al. 2025, MNRAS, 543, 2796, doi:10.1093/mnras/staf1470

Hobbs, G., Lorimer, D. R., Lyne, A. G., & Kramer, M. 2005, MNRAS, 360, 974, doi:10.1111/j.1365-2966.2005.09087.x

Hoekstra, H., Yee, H. K. C., & Gladders, M. D. 2004, ApJ, 606, 67, doi:10.1086/382726

Holdom, B. 1986a, Physics Letters B, 178, 65, doi:10.1016/0370-2693(86)90470-3

—. 1986b, Physics Letters B, 166, 196, doi:10.1016/0370-2693(86)91377-8

Homma, D., Chiba, M., Okamoto, S., et al. 2016, ApJ, 832, 21, doi:10.3847/0004-637X/832/1/21

Hopkins, P. F. 2015, Mon. Not. Roy. Astron. Soc., 450, 53, doi:10.1093/mnras/stv195

Hotokezaka, K., & Piran, T. 2017, ApJ, 842, 111, doi:10.3847/1538-4357/aa6f61

Hou, Z., Reichardt, C. L., Story, K. T., et al. 2014, ApJ, 782, 74, doi:10.1088/0004-637X/782/2/74

Hoyle, F. 1946, MNRAS, 106, 343, doi:10.1093/mnras/106.5.343

Hoyle, F., & Lyttleton, R. A. 1939, Proceedings of the Cambridge Philosophical Society, 35, 405, doi:10.1017/S0305004100021150

Hunter, J. D. 2007, Computing in Science & Engineering, 9, 90, doi:10.1109/MCSE.2007.55

Hurley, J. R., Pols, O. R., & Tout, C. A. 2000, Mon. Not. Roy. Astron. Soc., 315, 543, doi:10.1046/j.1365-8711.2000.03426.x

Hurley, J. R., Pols, O. R., & Tout, C. A. 2000, MNRAS, 315, 543, doi:10.1046/j.1365-8711.2000.03426.x

Hurley, J. R., Tout, C. A., & Pols, O. R. 2002, Mon. Not. Roy. Astron. Soc., 329, 897, doi:10.1046/j.1365-8711.2002.05038.x

Hutter, A., Cueto, E. R., Dayal, P., et al. 2025, A&A, 694, A254, doi:10.1051/0004-6361/202452460

Iorio, G., Mapelli, M., Costa, G., et al. 2023a, MNRAS, 524, 426, doi:10.1093/mnras/stad1630

—. 2023b, MNRAS, 524, 426, doi:10.1093/mnras/stad1630

Ivanova, N., Justham, S., Chen, X., et al. 2013, A&A Rev., 21, 59, doi:10.1007/s00159-013-0059-2

Janka, H.-T. 2012, Ann. Rev. Nucl. Part. Sci., 62, 407, doi:10.1146/annurev-nucl-102711-094901

Janka, H.-T. 2017, in Handbook of Supernovae, ed. A. W. Alsabti & P. Murdin, 1575, doi:10.1007/978-3-319-21846-5-4

Janka, H.-T., & Hillebrandt, W. 1989, A&A, 224, 49

Jiang, F., Jia, Z., Zheng, H., et al. 2025. https://arxiv.org/abs/2503.23710

Kalogera, V. 2000, ApJ, 541, 319, doi:10.1086/309400

Kalogera, V., & Baym, G. 1996, Astrophys. J. Lett., 470, L61, doi:10.1086/310296

Karathanasis, C., Mukherjee, S., & Mastrogiovanni, S. 2023, Mon. Not. Roy. Astron. Soc., 523, 4539, doi:10.1093/mnras/stad1373

Kauffmann, G., White, S. D. M., & Guiderdoni, B. 1993, MNRAS, 264, 201, doi:10.1093/mnras/264.1.201

Keil, M. T., Raffelt, G. G., & Janka, H.-T. 2003, ApJ, 590, 971, doi:10.1086/375130

Kiendrebeogo, R. W., et al. 2023, Astrophys. J., 958, 158, doi:10.3847/1538-4357/acfcb1

King, I. R. 1963, ARA&A, 1, 179, doi:10.1146/annurev.aa.01.090163.001143

Kinugawa, T., Nakamura, T., & Nakano, H. 2021, Mon. Not. Roy. Astron. Soc., 501, L49, doi:10.1093/mnrasl/slaa191

Kippenhahn, R., & Weigert, A. 1967, ZAp, 65, 251

Kirby, E. N., Cohen, J. G., Guhathakurta, P., et al. 2013, Astrophysical Journal, 779, doi:10.1088/0004-637X/779/2/102

Klencki, J., Nelemans, G., Istrate, A. G., & Chruslinska, M. 2021, Astron. Astrophys., 645, A54, doi:10.1051/0004-6361/202038707

Klencki, J., Podsiadlowski, P., Langer, N., et al. 2026, A&A, 706, A296, doi:10.1051/0004-6361/202555500

Kluyver, T., Ragan-Kelley, B., Pérez, F., et al. 2016, in Positioning and Power in Academic Publishing: Players, Agents and Agendas, ed. F. Loizides & B. Schmidt, IOS Press, 87–90

Klypin, A. A., Trujillo-Gomez, S., & Primack, J. 2011, ApJ, 740, 102, doi:10.1088/0004-637X/740/2/102

Kocevski, D. D., Onoue, M., Inayoshi, K., et al. 2023, ApJL, 954, L4, doi:10.3847/2041-8213/ace5a0

Kocevski, D. D., Finkelstein, S. L., Barro, G., et al. 2025, ApJ, 986, 126, doi:10.3847/1538-4357/adbc7d

Koda, J., & Shapiro, P. R. 2011, Monthly Notices of the Royal Astronomical Society, 415, 1125, doi:10.1111/j.1365-2966.2011.18684.x

Kokorev, V., Caputi, K. I., Greene, J. E., et al. 2024, ApJ, 968, 38, doi:10.3847/1538-4357/ad4265

Konopinski, E. J., & Uhlenbeck, G. E. 1935, Phys. Rev., 48, 7, doi:10.1103/PhysRev.48.7

Krishnan, J. R., & Abazajian, K. N. 2026, Phys. Rev. D, 113, 083007, doi:10.1103/gjf5-3r89

Kroupa, P. 2001, Monthly Notices of the Royal Astronomical Society, 322, 231, doi:10.1046/j.1365-8711.2001.04022.x

Kudritzki, R. P., Pauldrach, A., & Puls, J. 1987, A&A, 173, 293

Kuijken, K., & Gilmore, G. 1989, MNRAS, 239, 605, doi:10.1093/mnras/239.2.605

Kumamoto, J., Fujii, M. S., & Tanikawa, A. 2020, MNRAS, 495, 4268, doi:10.1093/mnras/staa1440

Kushnir, D., Zaldarriaga, M., Kollmeier, J. A., & Waldman, R. 2016, MNRAS, 462, 844, doi:10.1093/mnras/stw1684

Kuzio de Naray, R., Martinez, G. D., Bullock, J. S., & Kaplinghat, M. 2010, ApJL, 710, L161, doi:10.1088/2041-8205/710/2/L161

Kuzio de Naray, R., McGaugh, S. S., & de Blok, W. J. G. 2008, ApJ, 676, 920, doi:10.1086/527543

Labbé, I., van Dokkum, P., Nelson, E., et al. 2023, Nature, 616, 266, doi:10.1038/s41586-023-05786-2

Labbe, I., Greene, J. E., Bezanson, R., et al. 2025, ApJ, 978, 92, doi:10.3847/1538-4357/ad3551

Lacey, C., & Cole, S. 1993, MNRAS, 262, 627, doi:10.1093/mnras/262.3.627

Lalleman, M., Turbang, K., Callister, T., & van Remortel, N. 2025, A&A, 698, A85, doi:10.1051/0004-6361/202553941

Lattimer, J. M., & Schramm, D. N. 1974, ApJL, 192, L145, doi:10.1086/181612

Lattimer, J. M., & Swesty, D. F. 1991, Nuclear Physics A, 535, 331, doi:10.1016/0375-9474(91)90452-C

Lee, H., Skillman, E. D., Cannon, J. M., et al. 2006, ApJ, 647, 970, doi:10.1086/505573

Lee, N., Sanders, D. B., Casey, C. M., et al. 2015, ApJ, 801, 80, doi:10.1088/0004-637X/801/2/80

Li, P., Lelli, F., McGaugh, S., & Schombert, J. 2020, ApJS, 247, 31, doi:10.3847/1538-4365/ab700e

Li, S. W., Roberts, L. F., & Beacom, J. F. 2021, Phys. Rev. D, 103, 023016, doi:10.1103/PhysRevD.103.023016

Li, W., Chornock, R., Leaman, J., et al. 2011, MNRAS, 412, 1473, doi:10.1111/j.1365-2966.2011.18162.x

LIGO Scientific Collaboration, Virgo Collaboration, & KAGRA Collaboration. 2025, GWTC-4.0: Population Properties of Merging Compact Binaries, Zenodo, doi:10.5281/zenodo.16911563

Liu, B., & Bromm, V. 2020, Astrophys. J. Lett., 903, L40, doi:10.3847/2041-8213/abc552

Lu, S., Frenk, C. S., Bose, S., et al. 2025, MNRAS, 536, 1018, doi:10.1093/mnras/stae2646

Lucy, L. B., & Solomon, P. M. 1970, ApJ, 159, 879, doi:10.1086/150365

Lynden-Bell, D., & Wood, R. 1968, MNRAS, 138, 495, doi:10.1093/mnras/138.4.495

Ma, X., Hopkins, P. F., Faucher-Giguere, C.-A., et al. 2016, Mon. Not. Roy. Astron. Soc., 456, 2140, doi:10.1093/mnras/stv2659

Mace, C., Zeng, Z. C., Peter, A. H. G., et al. 2024, Phys. Rev. D, 110, 123024, doi:10.1103/PhysRevD.110.123024

Madau, P., & Dickinson, M. 2014, Ann. Rev. Astron. Astrophys., 52, 415, doi:10.1146/annurev-astro-081811-125615

Madau, P., & Fragos, T. 2017, The Astrophysical Journal, 840, 39, doi:10.3847/1538-4357/aa6af9

Mandel, I., & Broekgaarden, F. S. 2022, Living Reviews in Relativity, 25, 1, doi:10.1007/s41114-021-00034-3

Mandel, I., & Farmer, A. 2022, Phys. Rept., 955, 1, doi:10.1016/j.physrep.2022.01.003

Mannucci, F., Cresci, G., Maiolino, R., et al. 2009, MNRAS, 398, 1915, doi:10.1111/j.1365-2966.2009.15185.x

Mapelli, M. 2016, Mon. Not. Roy. Astron. Soc., 459, 3432, doi:10.1093/mnras/stw869

Mapelli, M., Colpi, M., & Zampieri, L. 2009, MNRAS, 395, L71, doi:10.1111/j.1745-3933.2009.00645.x

Mapelli, M., Ripamonti, E., Zampieri, L., Colpi, M., & Bressan, A. 2010, MNRAS, 408, 234, doi:10.1111/j.1365-2966.2010.17048.x

Mapelli, M., Santoliquido, F., Bouffanais, Y., et al. 2021, Symmetry, 13, 1678, doi:10.3390/sym13091678

Mapelli, M., Spera, M., Montanari, E., et al. 2020, The Astrophysical Journal, 888, 76, doi:10.3847/1538-4357/ab584d

Marciano, W. J., & Parsa, Z. 2003, J. Phys. G, 29, 2629, doi:10.1088/0954-3899/29/11/013

Marigo, P., Girardi, L., Bressan, A., et al. 2017, The Astrophysical Journal, 835, 77, doi:10.3847/1538-4357/835/1/77

Marks, M., Kroupa, P., Dabringhausen, J., & Pawlowski, M. S. 2012, Mon. Not. Roy. Astron. Soc., 422, 2246, doi:10.1111/j.1365-2966.2012.20767.x

Matthee, J., Naidu, R. P., Brammer, G., et al. 2024, ApJ, 963, 129, doi:10.3847/1538-4357/ad2345

McLeod, D. J., Donnan, C. T., McLure, R. J., et al. 2024, MNRAS, 527, 5004, doi:10.1093/mnras/stad3471

Metzger, B. D. 2017, arXiv e-prints, arXiv:1710.05931, doi:10.48550/arXiv.1710.05931

Meynet, G., Maeder, A., Schaller, G., Schaerer, D., & Charbonnel, C. 1994, A&AS, 103, 97

Miller, M. C., Lamb, F. K., Dittmann, A. J., et al. 2019, ApJL, 887, L24, doi:10.3847/2041-8213/ab50c5

Minor, Q. E., Kaplinghat, M., & Li, N. 2017, ApJ, 845, 118, doi:10.3847/1538-4357/aa7fee

Mirizzi, A., Tamborra, I., Janka, H.-T., et al. 2016, Riv. Nuovo Cim., 39, 1, doi:10.1393/ncr/i2016-10120-8

Moe, M., & Di Stefano, R. 2017, ApJS, 230, 15, doi:10.3847/1538-4365/aa6fb6

Morales, J. L., & Nocedal, J. 2011, ACM Trans. Math. Softw., 38, 7, doi:10.1145/2049662.2049669

Morton, D. C. 1960, ApJ, 132, 146, doi:10.1086/146908

Mould, M., Gerosa, D., Broekgaarden, F. S., & Steinle, N. 2022, MNRAS, 517, 2738, doi:10.1093/mnras/stac2859

Mukherjee, S. 2022, Mon. Not. Roy. Astron. Soc., 515, 5495, doi:10.1093/mnras/stac2152

Mukherjee, S., & Moradinezhad Dizgah, A. 2022, Astrophys. J. Lett., 937, L27, doi:10.3847/2041-8213/ac903b

Nadler, E. O., Yang, D., & Yu, H.-B. 2023, Astrophys. J. Lett., 958, L39, doi:10.3847/2041-8213/ad0e09

Naidu, R. P., Oesch, P. A., Setton, D. J., et al. 2022, arXiv e-prints, arXiv:2208.02794, doi:10.48550/arXiv.2208.02794

Nakajima, K., Ouchi, M., Isobe, Y., et al. 2023, ApJS, 269, 33, doi:10.3847/1538-4365/acd556

Nandez, J. L. A., & Ivanova, N. 2016, MNRAS, 460, 3992, doi:10.1093/mnras/stw1266

Napolitano, L., Castellano, M., Pentericci, L., et al. 2025, Astronomy & Astrophysics, 693, A50, doi:10.1051/0004-6361/202452090

Navarro, J. F., Eke, V. R., & Frenk, C. S. 1996a, MNRAS, 283, L72, doi:10.1093/mnras/283.3.L72

Navarro, J. F., Frenk, C. S., & White, S. D. M. 1996b, ApJ, 462, 563, doi:10.1086/177173

—. 1997, ApJ, 490, 493, doi:10.1086/304888

Neijssel, C. J., Vigna-Gómez, A., Stevenson, S., et al. 2019a, Monthly Notices of the Royal Astronomical Society, 490, 3740, doi:10.1093/mnras/stz2840

Neijssel, C. J., Vigna-Gómez, A., Stevenson, S., et al. 2019b, Monthly Notices of the Royal Astronomical Society, 490, 3740–3759, doi:10.1093/mnras/stz2840

Nguyen, C. T., Costa, G., Girardi, L., et al. 2022, A&A, 665, A126, doi:10.1051/0004-6361/202244166

Nishikawa, H., Boddy, K. K., & Kaplinghat, M. 2020, Phys. Rev. D, 101, 063009, doi:10.1103/PhysRevD.101.063009

Oh, S.-H., Hunter, D. A., Brinks, E., et al. 2015, AJ, 149, 180, doi:10.1088/0004-6256/149/6/180

O'Leary, R. M., Rasio, F. A., Fregeau, J. M., Ivanova, N., & O'Shaughnessy, R. W. 2006, Astrophys. J., 637, 937, doi:10.1086/498446

Olejak, A., & Belczynski, K. 2021, Astrophys. J. Lett., 921, L2, doi:10.3847/2041-8213/ac2f48

Olejak, A., Klencki, J., Xu, X.-T., et al. 2024, A&A, 689, A305, doi:10.1051/0004-6361/202450480

Oman, K. A., Marasco, A., Navarro, J. F., et al. 2019, MNRAS, 482, 821, doi:10.1093/mnras/sty2687

Oort, J. H. 1932, Bull. Astron. Inst. Netherlands, 6, 249

O'Shaughnessy, R., Kalogera, V., & Belczynski, K. 2010, Astrophys. J., 716, 615, doi:10.1088/0004-637X/716/1/615

Ott, C. D., O'Connor, E. P., Gossan, S., et al. 2013, Nuclear Physics B Proceedings Supplements, 235, 381, doi:10.1016/j.nuclphysbps.2013.04.036

Outmezguine, N. J., Boddy, K. K., Gad-Nasr, S., Kaplinghat, M., & Sagunski, L. 2023, Mon. Not. Roy. Astron. Soc., 523, 4786, doi:10.1093/mnras/stad1705

Ozel, F., Psaltis, D., Narayan, R., & McClintock, J. E. 2010, Astrophys. J., 725, 1918, doi:10.1088/0004-637X/725/2/1918

Paczyński, B. 1971, ARA&A, 9, 183, doi:10.1146/annurev.aa.09.090171.001151

Palubski, I., Slone, O., Kaplinghat, M., Lisanti, M., & Jiang, F. 2024, JCAP, 2024, 074, doi:10.1088/1475-7516/2024/09/074 pandas development team, T. 2026, pandas-dev/pandas: Pandas, v3.0.2, Zenodo, doi:10.5281/zenodo.19340003

Pauldrach, A., Puls, J., & Kudritzki, R. P. 1986, A&A, 164, 86

Peebles, P. J. E. 1982, ApJL, 263, L1, doi:10.1086/183911

—. 1984, ApJ, 277, 470, doi:10.1086/161714

Perez, F., & Granger, B. E. 2007, Computing in Science and Engineering, 9, 21, doi:10.1109/MCSE.2007.53

Pérez-González, P. G., et al. 2025. `https://arxiv.org/abs/2503.15594`

Perley, D. A., et al. 2020, Astrophys. J., 904, 35, doi:10.3847/1538-4357/abbd98

Peters, P. C. 1964, Phys. Rev., 136, B1224, doi:10.1103/PhysRev.136.B1224

Plavec, M. 1968, Ap&SS, 1, 239, doi:10.1007/BF00717920

PlotDigitizer. 2026, PlotDigitizer. `https://plotdigitizer.com`

Pols, O. R., Schröder, K.-P., Hurley, J. R., Tout, C. A., & Eggleton, P. P. 1998, MNRAS, 298, 525, doi:10.1046/j.1365-8711.1998.01658.x

Pontzen, A., & Governato, F. 2012, MNRAS, 421, 3464, doi:10.1111/j.1365-2966.2012.20571.x

Popesso, P., Concas, A., Morselli, L., et al. 2019, MNRAS, 483, 3213, doi:10.1093/mnras/sty3210

Popesso, P., Concas, A., Cresci, G., et al. 2023, MNRAS, 519, 1526, doi:10.1093/mnras/stac3214

Portegies Zwart, S. F., & McMillan, S. L. W. 2000, ApJL, 528, L17, doi:10.1086/312422

Poulin, V., Smith, T. L., Karwal, T., & Kamionkowski, M. 2019, Phys. Rev. Lett., 122, 221301, doi:10.1103/PhysRevLett.122.221301

Press, W. H., & Schechter, P. 1974, ApJ, 187, 425, doi:10.1086/152650

Qin, Y., Fragos, T., Meynet, G., et al. 2018, A&A, 616, A28, doi:10.1051/0004-6361/201832839

Renzo, M., & Götberg, Y. 2021, ApJ, 923, 277, doi:10.3847/1538-4357/ac29c5

Renzo, M., Zapartas, E., Justham, S., et al. 2023, ApJL, 942, L32, doi:10.3847/2041-8213/aca4d3

Renzo, M., Zapartas, E., de Mink, S. E., et al. 2019, A&A, 624, A66, doi:10.1051/0004-6361/201833297

Rezzolla, L., Most, E. R., & Weih, L. R. 2018, Astrophys. J. Lett., 852, L25, doi:10.3847/2041-8213/aaa401

Rhee, G., Valenzuela, O., Klypin, A., Holtzman, J., & Moorthy, B. 2004, ApJ, 617, 1059, doi:10.1086/425565

Riley, T. E., Watts, A. L., Bogdanov, S., et al. 2019, ApJL, 887, L21, doi:10.3847/2041-8213/ab481c

Rinaldi, S., Del Pozzo, W., Mapelli, M., Lorenzo-Medina, A., & Dent, T. 2024, Astron. Astrophys., 684, A204, doi:10.1051/0004-6361/202348161

Rinaldi, S., Liang, Y., Demasi, G., Mapelli, M., & Del Pozzo, W. 2025. https://arxiv.org/abs/2506.05929

Roberts, M. G., Braff, L., Garg, A., Profumo, S., & Jeltema, T. 2025. https://arxiv.org/abs/2507.03230

Robertson, A., Massey, R., & Eke, V. 2017, Monthly Notices of the Royal Astronomical Society, 467, 4719, doi:10.1093/mnras/stx463

Robertson, B. E., et al. 2023, Nature Astron., 7, 611, doi:10.1038/s41550-023-01921-1

Robitaille, T., Tollerud, E., Aldcroft, T., et al. 2023, astropy/astropy: v5.3.4, v5.3.4, Zenodo, doi:10.5281/zenodo.8408438

Roepke, F. K., & De Marco, O. 2023, Liv. Rev. Comput. Astrophys., 9, 2, doi:10.1007/s41115-023-00017-x

Romagnolo, A., Belczynski, K., Klencki, J., et al. 2023, MNRAS, 525, 706, doi:10.1093/mnras/stad2366

Romagnolo, A., Broekgaarden, F. S., Antoniadis, K., & Gormaz-Matamala, A. C. 2026a, arXiv e-prints, arXiv:2601.02263, doi:10.48550/arXiv.2601.02263

—. 2026b, arXiv e-prints, arXiv:2601.02263, doi:10.48550/arXiv.2601.02263

Romagnolo, A., Gormaz-Matamala, A. C., & Belczynski, K. 2024, Astrophys. J. Lett., 964, L23, doi:10.3847/2041-8213/ad2fbe

Romagnolo, A., Gormaz-Matamala, A. C., & Belczynski, K. 2024, ApJL, 964, L23, doi:10.3847/2041-8213/ad2fbe

Romagnolo, A., Klencki, J., Vigna-Gómez, A., & Belczynski, K. 2025, A&A, 693, A137, doi:10.1051/0004-6361/202452169

Rozwadowska, K., Vissani, F., & Cappellaro, E. 2021, New A, 83, 101498, doi:10.1016/j.newast.2020.101498

Rubin, V. C., & Ford, Jr., W. K. 1970, ApJ, 159, 379, doi:10.1086/150317

Sadiq, J., Dent, T., & Gieles, M. 2024, Astrophys. J., 960, 65, doi:10.3847/1538-4357/ad0ce6

Safarzadeh, M., & Ramirez-Ruiz, E. 2021. `https://arxiv.org/abs/2105.08746`

Sakstein, J., Croon, D., McDermott, S. D., Straight, M. C., & Baxter, E. J. 2020, Phys. Rev. Lett., 125, 261105, doi:10.1103/PhysRevLett.125.261105

Sales, L. V., Wetzel, A., & Fattahi, A. 2022, Nature Astronomy, 6, 897, doi:10.1038/s41550-022-01689-w

Salpeter, E. E. 1955, ApJ, 121, 161, doi:10.1086/145971

Sana, H., Gosset, E., Naze, Y., Rauw, G., & Linder, N. 2008, Mon. Not. Roy. Astron. Soc., 386, 447, doi:10.1111/j.1365-2966.2008.13037.x

Sana, H., De Mink, S. E., De Koter, A., et al. 2012, Science, 337, 444, doi:10.1126/science.1223344

Sand, C., Ohlmann, S. T., Schneider, F. R. N., Pakmor, R., & Röpke, F. K. 2020, A&A, 644, A60, doi:10.1051/0004-6361/202038992

Sanders, R. L., Shapley, A. E., Kriek, M., et al. 2015, ApJ, 799, 138, doi:10.1088/0004-637X/799/2/138

Santoliquido, F., Mapelli, M., Bouffanais, Y., et al. 2020, Astrophys. J., 898, 152, doi:10.3847/1538-4357/ab9b78

Schiebelbein-Zwack, A., & Fishbach, M. 2024. `https://arxiv.org/abs/2403.17156`

Schneider, F., Izzard, R., Langer, N., & de Mink, S. 2015, The Astrophysical Journal, 805, 20

Scholberg, K. 2012, Ann. Rev. Nucl. Part. Sci., 62, 81, doi:10.1146/annurev-nucl-102711-095006

Schwarzschild, K. 1916, in Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften zu Berlin, 424–434

Sedda, M. A., Naoz, S., & Kocsis, B. 2023, Universe, 9, 138, doi:10.3390/universe9030138

Sen, K., Langer, N., Marchant, P., et al. 2022, A&A, 659, A98, doi:10.1051/0004-6361/202142574

Sen, K., Langer, N., Pauli, D., et al. 2023, A&A, 672, A198, doi:10.1051/0004-6361/202245378

Sen, K., Renzo, M., Jin, H., et al. 2026, ApJ, 1000, 2, doi:10.3847/1538-4357/ae40fc

Sgalletta, C., Mapelli, M., Boco, L., et al. 2025, A&A, 698, A144, doi:10.1051/0004-6361/202452757

Sgalletta, C., Costa, G., Iorio, G., et al. 2026, arXiv e-prints, arXiv:2605.06807. `https://arxiv.org/abs/2605.06807`

Shen, X., Vogelsberger, M., Boylan-Kolchin, M., Tacchella, S., & Naidu, R. P. 2024, MNRAS, 533, 3923, doi:10.1093/mnras/stae1932

Silverman, M., Hussein, A., Arora, A., et al. 2026, arXiv e-prints, arXiv:2606.02566, doi:10.48550/arXiv.2606.02566

Simon, J. D. 2019, ARA&A, 57, 375, doi:10.1146/annurev-astro-091918-104453

Sipior, M. S., Portegies Zwart, S., & Nelemans, G. 2004, MNRAS, 354, L49, doi:10.1111/j.1365-2966.2004.08373.x

Smak, J. 1962, Acta Astron., 12, 28

Smith, T. 2026, Massquerade: Impacts of Mass Ratio Reversals on Binary Black Hole Merger Rates and Mass Distributions, Zenodo, doi:10.5281/zenodo.20189849

Smith, T. B., & Kaplinghat, M. 2026, ApJ, 997, 231, doi:10.3847/1538-4357/ae21d2

Speagle, J. S., Steinhardt, C. L., Capak, P. L., & Silverman, J. D. 2014, Astrophys. J. Suppl., 214, 15, doi:10.1088/0067-0049/214/2/15

Spera, M., & Mapelli, M. 2017, Monthly Notices of the Royal Astronomical Society, 470, 4739, doi:10.1093/mnras/stx1576

Spera, M., Mapelli, M., & Bressan, A. 2015, Monthly Notices of the Royal Astronomical Society, 451, 4086, doi:10.1093/mnras/stv1161

Spera, M., Mapelli, M., Giacobbo, N., et al. 2019, Monthly Notices of the Royal Astronomical Society, 485, 889, doi:10.1093/mnras/stz359

Spergel, D. N., & Steinhardt, P. J. 2000, Phys. Rev. Lett., 84, 3760, doi:10.1103/PhysRevLett.84.3760

Springel, V., Frenk, C. S., & White, S. D. M. 2006, Nature, 440, 1137, doi:10.1038/nature04805

Steidel, C. C., Rudie, G. C., Strom, A. L., et al. 2014, ApJ, 795, 165, doi:10.1088/0004-637X/795/2/165

Steiner, A. W., Hempel, M., & Fischer, T. 2013, Astrophys. J., 774, 17, doi:10.1088/0004-637X/774/1/17

Stevenson, S., & Clarke, T. 2022, doi:10.1093/mnras/stac2936

Stevenson, S., Vigna-Gómez, A., Mandel, I., et al. 2017, Nature Commun., 8, 14906, doi:10.1038/ncomms14906

Stevenson, S., Vigna-Gómez, A., Mandel, I., et al. 2017, Nat. Commun., 8, 14906, doi:10.1038/ncomms14906

Strumia, A., & Vissani, F. 2003, Phys. Lett. B, 564, 42, doi:10.1016/S0370-2693(03)00616-6

Sun, G., Faucher-Giguère, C.-A., Hayward, C. C., et al. 2023, ApJL, 955, L35, doi:10.3847/2041-8213/acf85a

Suzuki, Y. 2019, Eur. Phys. J. C, 79, 298, doi:10.1140/epjc/s10052-019-6796-2

Symbalisty, E., & Schramm, D. N. 1982, Astrophys. Lett., 22, 143

Tamborra, I., Müller, B., Hüdepohl, L., Janka, H.-T., & Raffelt, G. 2012, Phys. Rev. D, 86, 125031, doi:10.1103/PhysRevD.86.125031

Tang, J., Bressan, A., Rosenfield, P., et al. 2014, Monthly Notices of the Royal Astronomical Society, 445, 4287, doi:10.1093/mnras/stu2029

Tang, P. N., Eldridge, J. J., Stanway, E. R., & Bray, J. C. 2020, Mon. Not. Roy. Astron. Soc., 493, L6, doi:10.1093/mnrasl/slz183

Tanikawa, A., Kinugawa, T., Yoshida, T., Hijikawa, K., & Umeda, H. 2021, Mon. Not. Roy. Astron. Soc., 505, 2170, doi:10.1093/mnras/stab1421

Tanikawa, A., Yoshida, T., Kinugawa, T., et al. 2022, Astrophys. J., 926, 83, doi:10.3847/1538-4357/ac4247

Tauris, T. M., & Sennels, T. 2000, A&A, 355, 236, doi:10.48550/arXiv.astro-ph/9909149

Tegmark, M., Blanton, M. R., Strauss, M. A., et al. 2004, ApJ, 606, 702, doi:10.1086/382125

Teyssier, R., Pontzen, A., Dubois, Y., & Read, J. I. 2013, MNRAS, 429, 3068, doi:10.1093/mnras/sts563

The LIGO Scientific Collaboration, the Virgo Collaboration, the KAGRA Collaboration, et al. 2026a, arXiv e-prints, arXiv:2605.27225, doi:10.48550/arXiv.2605.27225

—. 2026b, arXiv e-prints, arXiv:2605.27226, doi:10.48550/arXiv.2605.27226

—. 2026c, arXiv e-prints, arXiv:2605.27223, doi:10.48550/arXiv.2605.27223

Tiwari, V. 2022, Astrophys. J., 928, 155, doi:10.3847/1538-4357/ac589a

Toonen, S., Perets, H. B., Igoshev, A. P., Michaely, E., & Zenati, Y. 2018, A&A, 619, A53, doi:10.1051/0004-6361/201833164

Torniamenti, S., Mapelli, M., Périgois, C., et al. 2024. `https://arxiv.org/abs/2401.14837`

Tremonti, C. A., Heckman, T. M., Kauffmann, G., et al. 2004, ApJ, 613, 898, doi:10.1086/423264

Trinca, A., Schneider, R., Valiante, R., et al. 2024, MNRAS, 529, 3563, doi:10.1093/mnras/stae651

Tulin, S., & Yu, H.-B. 2018, Phys. Rept., 730, 1, doi:10.1016/j.physrep.2017.11.004

Tutukov, A., Yungelson, L., & Klayman, A. 1973, Nauchnye Informatsii, 27, 3

van den Heuvel, E. P. J. 1976, in IAU Symposium, Vol. 73, Structure and Evolution of Close Binary Systems, ed. P. Eggleton, S. Mitton, & J. Whelan, 35

Van Rossum, G., & Drake, F. L. 2009, Python 3 Reference Manual (Scotts Valley, CA: CreateSpace)

van Son, L. A. C., de Mink, S. E., Chruślińska, M., et al. 2023a, ApJ, 948, 105, doi:10.3847/1538-4357/acbf51

—. 2023b, ApJ, 948, 105, doi:10.3847/1538-4357/acbf51

van Son, L. A. C., de Mink, S. E., Callister, T., et al. 2022, Astrophys. J., 931, 17, doi:10.3847/1538-4357/ac64a3

Ventura, E. M., Qin, Y., Balu, S., & Wyithe, J. S. B. 2024, MNRAS, 529, 628, doi:10.1093/mnras/stae567

Vigna-Gómez, A., et al. 2018, Mon. Not. Roy. Astron. Soc., 481, 4009, doi:10.1093/mnras/sty2463

Vink, J. S. 2025, arXiv e-prints, arXiv:2506.06421, doi:10.48550/arXiv.2506.06421

Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, Nature Methods, 17, 261, doi:10.1038/s41592-019-0686-2

Wagg, T., Broekgaarden, F., Van-Lane, P., Wu, K., & Gültekin, K. 2025, TomWagg/software-citation-station: v1.4, v1.4, Zenodo, doi:10.5281/zenodo.17654855

Wagg, T., & Broekgaarden, F. S. 2024, arXiv e-prints, arXiv:2406.04405. `https://arxiv.org/abs/2406.04405`

Wagg, T., Johnston, C., Bellinger, E. P., et al. 2024, A&A, 687, A222, doi:10.1051/0004-6361/202449912

Wang, B., Fujimoto, S., Labbé, I., et al. 2023, ApJL, 957, L34, doi:10.3847/2041-8213/acfe07

Weatherford, N. C., Fragione, G., Kremer, K., et al. 2021, Astrophys. J. Lett., 907, L25, doi:10.3847/2041-8213/abd79c

Weaver, J. R., Davidzon, I., Toft, S., et al. 2023, A&A, 677, A184, doi:10.1051/0004-6361/202245581

Webbink, R. F. 1988, in Critical Observations Versus Physical Models for Close Binary Systems, ed. K. C. Leung, 403–446

Wechsler, R. H., & Tinker, J. L. 2018, ARA&A, 56, 435, doi:10.1146/annurev-astro-081817-051756

Wellstein, S., Langer, N., & Braun, H. 2001, A&A, 369, 939, doi:10.1051/0004-6361:20010151

Wes McKinney. 2010, in Proceedings of the 9th Python in Science Conference, ed. Stéfan van der Walt & Jarrod Millman, 56–61, doi:10.25080/Majora-92bf1922-00a

Whitaker, K. E., Franx, M., Leja, J., et al. 2014, ApJ, 795, 104, doi:10.1088/0004-637X/795/2/104

White, S. D. M., & Rees, M. J. 1978, MNRAS, 183, 341, doi:10.1093/mnras/183.3.341

Williams, C. C., Alberts, S., Ji, Z., et al. 2024, ApJ, 968, 34, doi:10.3847/1538-4357/ad3f17

Willman, B., Dalcanton, J. J., Martinez-Delgado, D., et al. 2005, ApJL, 626, L85, doi:10.1086/431760

Woosley, S. E. 2017, The Astrophysical Journal, 836, 244, doi:10.3847/1538-4357/836/2/244

—. 2019, The Astrophysical Journal, 878, 49, doi:10.3847/1538-4357/ab1b41

Woosley, S. E., Blinnikov, S., & Heger, A. 2007, Nature, 450, 390, doi:10.1038/nature06333

Woosley, S. E., & Heger, A. 2021, Astrophys. J. Lett., 912, L31, doi:10.3847/2041-8213/abf2c4

Xiao, M., Oesch, P. A., Elbaz, D., et al. 2024, Nature, 635, 311, doi:10.1038/s41586-024-08094-5

Xu, X.-J., & Li, X.-D. 2010, ApJ, 716, 114, doi:10.1088/0004-637X/716/1/114

Yang, D., & Yu, H.-B. 2022, JCAP, 09, 077, doi:10.1088/1475-7516/2022/09/077

Yang, S., Du, X., Zeng, Z. C., et al. 2023, The Astrophysical Journal, 946, 47, doi:10.3847/1538-4357/acbd49

Yoshida, T., Takiwaki, T., Kotake, K., et al. 2019, ApJ, 881, 16, doi:10.3847/1538-4357/ab2b9d

Yung, L. Y. A., Somerville, R. S., & Iyer, K. G. 2025, MNRAS, 543, 3802, doi:10.1093/mnras/staf1699

Zahid, H. J., Geller, M. J., Kewley, L. J., et al. 2013, ApJL, 771, L19, doi:10.1088/2041-8205/771/2/L19

Zahid, H. J., Kewley, L. J., & Bresolin, F. 2011, ApJ, 730, 137, doi:10.1088/0004-637X/730/2/137

Zapartas, E., de Mink, S. E., Justham, S., et al. 2021, A&A, 645, A6, doi:10.1051/0004-6361/202037744

Zapartas, E., Fox, O. D., Su, J., et al. 2026, MNRAS, 546, staf2208, doi:10.1093/mnras/staf2208

Zemp, M., Gnedin, O. Y., Gnedin, N. Y., & Kravtsov, A. V. 2011, The Astrophysical Journal Supplement Series, 197, 30, doi:10.1088/0067-0049/197/2/30

Zevin, M., & Bavera, S. S. 2022, ApJ, 933, 86, doi:10.3847/1538-4357/ac6f5d

Zevin, M., Romero-Shaw, I. M., Kremer, K., Thrane, E., & Lasky, P. D. 2021, ApJL, 921, L43, doi:10.3847/2041-8213/ac32dc

Zhang, Z.-Y., Romano, D., Ivison, R. J., Papadopoulos, P. P., & Matteucci, F. 2018, Nature, 558, 260, doi:10.1038/s41586-018-0196-x

Zhu, C., Byrd, R. H., Lu, P., & Nocedal, J. 1997, ACM Trans. Math. Softw., 23, 550, doi:10.1145/279232.279236

Ziegler, J., & Freese, K. 2021, Phys. Rev. D, 104, 043015, doi:10.1103/PhysRevD.104.043015

—. 2024, Phys. Rev. D, 109, 103042, doi:10.1103/PhysRevD.109.103042

Zolotov, A., Brooks, A. M., Willman, B., et al. 2012, ApJ, 761, 71, doi:10.1088/0004-637X/761/1/71

Zwicky, F. 1933, Helvetica Physica Acta, 6, 110