A Galaxy of Binaries: The impact of binary interactions on gravitational-wave sources, asteroseismic pulsations and stellar kinematics

Thomas Wagg

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

Doctor of Philosophy

University of Washington

2025

Reading Committee:

Eric Bellm, Chair

Julianne Dalcanton

Emily Levesque

Program Authorized to Offer Degree:
Astronomy

University of Washington

Abstract

A Galaxy of Binaries: The impact of binary interactions on gravitational-wave sources, asteroseismic pulsations and stellar kinematics

Thomas Wagg

Chair of the Supervisory Committee:
Eric Bellm
Astronomy

Massive stars are extremely impactful across a range of astrophysics, providing essential feedback for galaxy evolution and producing transients such as supernovae and gravitational-wave mergers. The vast majority of these stars are formed in binaries and higher-order multiple systems, a large fraction of which will interact during their lifetimes and significantly alter their subsequent evolution. Despite the importance of massive stars, and the prevalence of massive *binary* stars, many aspects of binary stellar evolution remain uncertain across several orders of magnitude.

In this dissertation I describe my work in exploring the impact of binary interactions on a variety of massive stellar populations, both through rapid population synthesis and detailed 1D stellar evolution models. In particular, I consider the sensitivity of these results to major uncertainties in binary physics, with a view towards placing constraints on binary physics parameters. The culmination of this work is not only a series of scientific results and predictions, but also two new open-source codes (`LEGWORK` and `cogsworth`) that enable future community-driven investigations into these matters.

In Chapter 1, I introduce the importance of massive binary stars, and outline the major uncertainties that remain, as well as some methodologies for addressing these uncertainties. After this introduction, I consider the population of gravitational-wave sources that will be detectable by the future spaced-based detector *LISA* (Chapters 2 & 3). Our new open-

source code, `LEGWORK`, provides the community with a reliable resource for calculating the detectability of *LISA* sources and computing their evolution due to gravitational-wave emission. We use `LEGWORK` to explore how the rates and demographics of the galactic population of black hole and neutron star binaries are sensitive to different aspects of binary physics.

From here, I move from gravitational-waves to gravity-mode pulsations and the asteroseismology of massive stars that have accreted material from a companion (Chapter 4). With our proof-of-principle analysis, we established for the first time that mass transfer can leave a significant imprint on the asteroseismic signals of accretor stars. We showed that the rejuvenation of the convective core as a result of mass transfer leaves an imprint on the chemical composition gradient in the star. This change in gradient influences all pulsations that are sensitive to that region of the star, shifting their pulsation periods. We showed that typical asteroseismic methods for estimating the mass and age of the star (which assume single star evolution) can be significantly erroneous for accretor stars.

In the subsequent chapters, I focus on the new techniques I developed for predicting the *extrinsic* positions and kinematics of massive binaries, in addition to their intrinsic properties such as rates, masses and orbital periods. I present a new open-source code, `cogsworth`, which can be used to perform self-consistent population synthesis and galactic dynamics simulation (Chapter 5). We used `cogsworth` to demonstrate how the positions of massive stellar populations can inform our understanding of binary interactions. From here, I examine how ejection velocities of stars are relatively insensitive to the strength of supernova natal kicks and compare the prescriptions for these ejections across three different population synthesis codes (Chapter 6).

In Chapter 7, I use `cogsworth` to explore how binary interactions impact the timing and location of core-collapse supernovae. Current models for supernova feedback in hydrodynamical simulations assume that all massive stars are formed as single stars. We show that binary interactions result in both late-time and spatially-displaced supernovae. We find that more than a quarter of supernovae occur after the time of the final single star supernova, while 13% explode more than 100pc from their parent star cluster. We assessed the robustness of these predictions to a plethora of variations in binary physics, initial conditions and galaxy parameters and found these results are surprisingly insensitive to uncertainties in the parameters we considered. Given that these supernovae could have a significantly different impact on galactic feedback, we developed a new analytic model for core-collapse supernova feedback. This model includes physically-motivated metallicity-dependent transitions and reproduces the timing and velocity distributions of supernovae to within 1% and 4% respectively. Our model can be used in future hydrodynamical simulations to better account for binary evolution. This may be particularly relevant for high redshift galaxies in which the spatial extent of the galaxy is reduced, while the low-metallicity environment produces even later and more distributed supernovae.

Finally in Chapter 8, I discuss the future directions of this research and consider the potential avenues for constraining binary evolution that this thesis has enabled. In Chapter 9, I summarise this thesis, concluding with the hope that the findings I presented, as well as the new open-source codes that we released, will drive forward the field of massive binary evolution.

DEDICATION

To my incredible mother Emma Wagg. Not satisfied with the workload of raising three young children, she decided to attend university online and train to become the kind of teacher of which every student dreams. It's hard to imagine a stronger demonstration of what you can achieve with the right motivation driving you forwards, she is truly a paragon of compassion and tenacity. Since then, through her love, dedication and generosity she's gone on to inspire countless students to be their very best selves…me most of all.

I'd never have made it this far without you mum, truly thank you for everything.

ACKNOWLEDGEMENTS

> *I am glad you are here with me. Here at the end of all things, Sam.*
> – Frodo Baggins in *The Return of the King*

Although I'm not at the end of *all* things like Frodo, it's certainly the end of *some* things, and I so dearly appreciate the company I've had along the way. These past four years have been quite the journey, in many ways a somewhat meandering path and one that I have only been able to navigate because of the help, guidance, love and compassion from a myriad of wonderful people around me.

My mentors — I've benefited from the help and support of a variety of mentors over the course of my degree, but in particular I must highlight two stalwart companions: **Katie Breivik** and **Mathieu Renzo**. Katie and I first got in touch about re-deriving the *LISA* gravitational-wave SNR equations (you can look forward to the fruits of that labour in Chapter 2) and have been sharing our passion for open-source software and binaries ever since. At a similar time, Mathieu and I met through Selma de Mink, and I quickly began to appreciate not only Mathieu's encyclopedic knowledge of binary stellar evolution literature, but also his vehemently strong opinions. In the past several years I have frequently turned to them for advice and help and, without fail, have been met with sage wisdom and constant kindness. Their unerring belief in me and unwavering support of my science have meant the world to me, and given me the confidence to keep going even at the most difficult moments. In those moments it could have been easy to feel adrift, but it was these two selfless people far across the country who were my anchors, providing that essential motivating guidance throughout. I don't think I could have managed this without them, I am truly very grateful.

Eric Bellm took on the mantle of my PhD advisor before there was any evidence these research directions could succeed. Despite how often he might label it as "free-associating", I've learnt a lot from Eric about the world of academia and the reality of being a faculty member. It was with his help that I was able to apply to NASA for the funding that drove much of the research in this thesis. And I have no doubt that my success in acquiring that funding was in no small part as a result of the guidance of **Julianne Dalcanton**. Julianne has taught me so much about moving beyond answering questions. From her I've learnt to think more deeply about which questions to ask, how to answer them, and why they need answering right now[1]. In short, she has made me a much better scientist. And perhaps someday I may forgive her from leaving UW the moment I arrived...

I'm also very thankful for the advice and wisdom of **Selma de Mink** over the past few years. After I completed my undergraduate thesis with Selma, I was lucky enough to spend my gap year working with her, during which she taught me much of what it means to be a

---

[1] Also how Oxford commas work

good scientist. Indeed my enthusiasm for binary stars was first kindled in Selma's office when hearing her describe some of the wonders of massive stellar evolution during our "30 minute" meetings...which often went on for two hours because we were having so much fun. Despite my constant refusal to move to Germany, Selma has always been generous with her time and efforts — as she is with so many early career scientists — and I am deeply appreciative of everything I've learned from her. I look forward to many more years of collaboration ahead!

I have continued to be inspired by **Floor Broekgaarden**'s compassionate mentoring since she supervised my undergraduate thesis with Selma. Floor has a deep understanding of how good science can only be done by people who feel welcome, included and supported. I can only hope to be such an excellent mentor to my students some day.

I've been lucky to connect my work to several other subfields throughout my PhD. **Adrian Price-Whelan** welcomed me into the delightful Nearby Universe group while I was a pre-doc at the CCA and taught me pretty much everything I know about galactic dynamics. I can't wait to join the CCA and keep learning more together. In Chapter 4 you'll hear about my escapades into asteroseismology one summer under the tutelage of **Cole Johnston** and **Earl Bellinger**. Cole and Earl were a delight to work with and made sure I never felt out of my depth regardless of the fact I was very new to the idea of stellar pulsations. I definitely suffered withdrawal from the puns once we finished that summer in Munich.

My department — Though most people mentioned above are located rather far from Seattle, I would be remiss not to acknowledge the people here at UW who've brightened my PhD experience. **Eric Agol** was kind enough to take me on as a student and not resent the fact that I quickly realised exoplanet observations were not for me, instead lending me help in applying for the CCA pre-doc position that helped me to continue working on massive stars. I'm thankful to **Jess Werk** (a.k.a. Eric C) for her relentless dedication to this department and its people, I can think of no one better suited, or that I would trust more, to be the leader of UW astronomy. It's been great to discuss all sorts of things with fellow massive star enthusiast **Emily Levesque** and I'm very glad to have her as part of my thesis committee. I've also very much appreciated **Jim Davenport**'s open door and understanding, as well as **Andy Connolly**'s sense of humour and honest conversations. I'm thankful to everyone in the front office for all that they do in keeping things running, and particularly to **Tyneshia Valdez** and **Liza Young** for their countless hours spent looking out for grad students.

My community — Somehow it never felt like I was living on the opposite coast from **Alex**. He's the best friend one can hope for, always ready to listen to my stories, and gently nudge me in the right direction. In much the same way, **Timmy** is always there when I need him and ready to offer some fresh perspective on what's really important (plus an easy target in a pool match). **Connor** and I have may have mostly been disembodied voices to one another these past few years, but his ability to say what everyone's thinking (but too embarrassed to say) has been a constant joy. I'm always grateful for **Shelly**'s empathy in all situations and for all the support he gives me. Now speaking of support, I seem to recall the promise of research funding from him the *last* time I wrote thesis acknowledgements, we'll have to follow up on that. One of the happiest parts of the last year has been meeting **Emily**. Her genuine enthusiasm for my work has constantly reminded me of my own fascination with astronomy. My day is constantly brightened each time I see her, I feel extremely lucky that we'll both be moving to New York this year.

The experience of this PhD would have been utterly different without the community I found in my fellow grad students in the astronomy department. Our bonds have certainly been forged in fire at times! I am so glad to have been able to share the start of my experience with such a wonderful cohort, **The Fab Five**$^{\mathrm{TM}}$, you all immediately made me feel at home in Seattle. I don't have the space to express my thanks for every moment of joy and support I've shared with every grad, but I shall name a few that come to mind: **Andy** for arriving at my apartment within minutes of me calling, just when I felt like my life was falling apart, **Megan** and **Rudy** for providing just the distractions I needed in B329, **Sam** for inspiring the quotes channel, **David** for bringing me along on all sorts of adventures, **v** for your insightful questions and poignant conversations, **Tobin** for the best 4th of July I've ever had and pulling me out of that class five rapid that one time, **Jake** for your level-headed wisdom and endless compassion, and **Tyler** for enduring my constant jokes.

Everyone in both the Seattle Field Hockey Club and Seattle Jets has provided a wonderful home outside of academia with our matches twice a week. I can't mention everyone but I'd particularly like to thank **Sarah** for her patience in transporting many of us to and from Saturday hockey in her "party bus". Hopefully this acknowledgement can give her another reason to use her favourite nickname for me. **Aps** has been a steadfast leader and friend and always so supportive. It's also a delight to have some proper Englishmen about with **Taz** and **Mark**, and I've got to acknowledge Mark for the best Sunday roast I've had in a long while.

My family — Of course I could never have come this far without my family. You can read more about **mum** in my dedication, so I shall simply say here that I owe her everything. **Dad** taught me many things, but most of all tenacity and that life is a marathon, not a sprint. I'm lucky to have a plethora of siblings with **Joe**, **Lucy**, **Dan**, and **Jordan**. I'm so proud of their continuing successes and all of the ways that they have made me a better person. Spending time in everyone's favourite English country cottage with **Jamie, Rebecca, Freddie and Verity** is always a highlight of my year, and I'm very lucky to have such a wonderful role model of a relationship and family to aim for. I'm very thankful to my grandmothers, **Jo** and **Mary**, for their love and kindness and cherished family dinners. I'm looking forward to my next wander around Eccelshall with Nan Jo where she can parade Dr. Wagg about! Last, and most definitely not least, **Granddad Terry**'s enthusiasm for my science was a constant driving force in my childhood. I fondly remember many hours spent talking in his workshop and garden together. I wish that he was here to read this thesis as he did my last one, I still miss him terribly and I hope that he would be proud of the research he's inspired.

It is such a joy being able to read back over this at the veritable conglomerate I have amassed, it very much takes a village. I love you all and am very excited for the next steps in this scientific adventure. Without any further ado, on with the show!

CONTENTS

CHAPTER 1

INTRODUCTION

The Wheel of Time turns, and ages come and pass, leaving memories that become legend. Legends fade to myth, and even myth is long forgotten when the Age that gave it birth comes again. In one Age, called the third age by some, an Age yet to come, an age long pass, a wind rose in the Mountains of Mist. The wind was not the beginning. There are neither beginnings or endings to the turning of the Wheel of Time. But it was a beginning.

– Robert Jordan in *The Eye of the World*

1 The central role of massive stars in astronomy

Massive stars are inherently rare. Let us define massive stars to be those that end their lives with a supernova, such that they are generally more massive than $\sim 7\,\mathrm{M}_\odot$ at solar metallicity. Given the stellar initial mass function, these represent only around 0.5% of all stars (Kroupa, 2001). Moreover, the nuclear timescale of a more massive star is significantly reduced, such that a $20\,\mathrm{M}_\odot$ star typically lives only 10 Myr, compared to a $1\,\mathrm{M}_\odot$ star that lives for 10 Gyr.

Despite their rare and short-lived nature, massive stars play a critical role in an array of astrophysical processes, from individual events to extragalactic scales. Massive stars are the drivers behind a variety of astrophysical transients, including core-collapse supernovae (e.g. Janka, 2012), gamma-ray bursts (GRBs, e.g. **?**Bloom et al., 2002) — both short GRBs as neutron star remnants and long GRBs as collapsing massive stars — and gravitational-wave events as the merger of two compact objects (e.g Abbott et al., 2023).

Beyond individual events, massive stars produce strong feedback that shapes their galactic environment, regulates star formation and enriches their surroundings (e.g. Dekel & Silk, 1986; Hopkins et al., 2012). Feedback from supernovae and strong stellar winds heats the interstellar medium and prevents further collapse of molecular clouds, thereby halting star formation, whilst also depositing enriched material. Without this massive stellar feedback, simulated galaxies tend to form too many stars and become too dense compared to observed galaxies (e.g. Katz et al., 1996; Hopkins et al., 2018a).

An understanding of massive stars is also essential for characterising higher redshift galaxies, in which spectra are found to be dominated by the high ionisation features associated with massive stars (e.g. Pettini et al., 2000; Shapley et al., 2003; Bunker et al., 2004). The strong ionisation provided by massive stars is also necessary to reproduce the re-ionisation of the Universe (Haiman & Loeb, 1997).

Thus it is imperative for a range of astrophysics that we produce clear and well-constrained models for massive stellar evolution. Yet each of the phenomena above, whether by the rate of short gamma-ray bursts and gravitational wave events or the necessity of significant ionising radiation, additionally hints at another requirement for massive stars: their binarity.

2 The ubiquity of binary massive stars

In recent years it has become clear than an understanding of massive stars *requires* an understanding of binaries. Massive binaries are not a rare phenomenon; the majority of massive stars are formed orbiting a companion in binary and multiple star systems (e.g., Kroupa, 1995; Duchêne & Kraus, 2013; Moe & Di Stefano, 2017; Oliva & Kuiper, 2020; Offner et al., 2023). Binary systems are constructed in preferential manners, with Moe & Di Stefano (2017) showing that there are correlations in the mass ratio and orbital period of binary systems, as well as an excess of twin stars. Moreover, in the most recent iteration of these studies, Offner et al. (2023) demonstrate the well-known strong dependence of multiplicity on stellar mass, as shown in Figure 1.1. The bold lines in this figure show the "non-single" fraction and thus one can note that only half of solar-type stars are single, whilst almost no O-type star is formed without a companion.

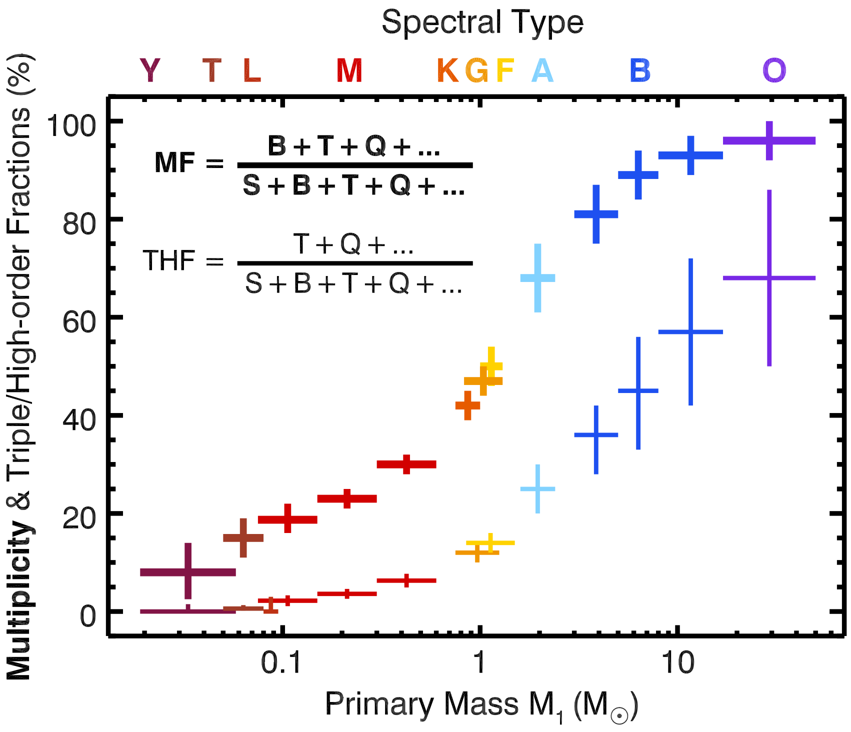

Figure 1.1: Binarity and multiplicity increase a function of stellar mass (Figure 1 from Offner et al., 2023). Almost all O-type stars are formed in multiple star systems, whilst even solar-type stars are formed in multiple star systems approximately half of the time. Bold lines show total multiplicity and thinner lines indicate triple and higher order fractions.

Furthermore, these are not simply wide binaries in which the stars rarely interact. Sana et al. (2012) established that the vast majority (∼70%) of O-type stars are formed in *interacting* binaries that exchange mass in some way. The subsequent evolution of these massive stars is significantly altered as a result of interactions with their companion, often resulting in the stripping of a star's envelope, the accretion of (potentially enriched) material and stellar mergers (e.g., Sana et al., 2012; de Mink et al., 2014). Figure 1.2 shows the fraction of O-type stars that experience these different types of interactions as a function of initial binary orbital period from Sana et al. (2012). As one may expect, the closest systems result in mergers, with slightly wider systems experiencing accretion and stripping and the widest systems have no interaction at all.

Binary interactions provide several new pathways of evolution for massive stars (Marchant & Bodensteiner, 2024), changing both their intrinsic properties, such through stripping (e.g., Kippenhahn & Weigert, 1967; Götberg et al., 2023; Drout et al., 2023) or rapid spin up via accretion (e.g., Packet, 1981; Renzo & Götberg, 2021), as well as their positions and kinematics within their galaxies (e.g., Blaauw, 1961; Renzo et al., 2019).

We have discussed how massive stars are incredibly impactful across the Universe and that most are found in interacting binaries, which have processes that can dramatically change their evolution and thus their impact. The reader may therefore find themselves concerned to learn quite how uncertain many of the processes in binary stellar evolution remain.

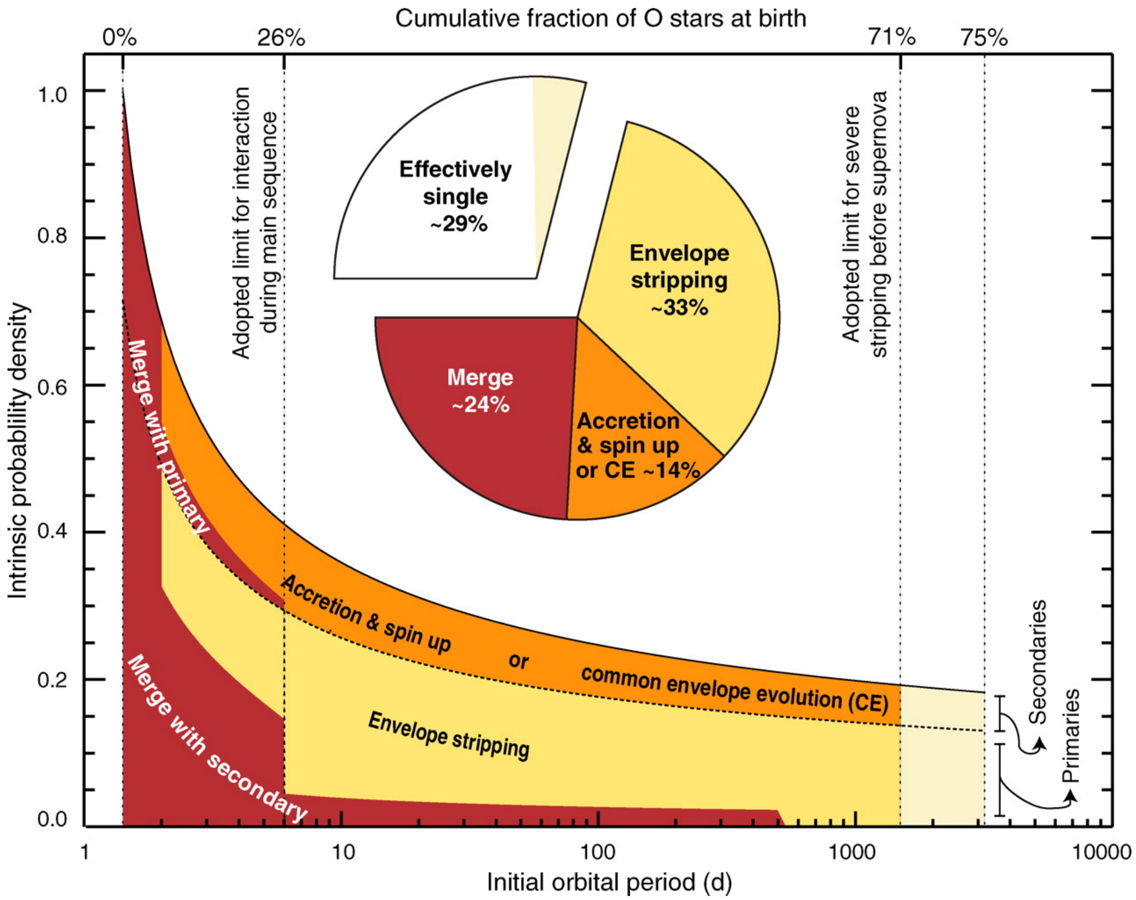

Figure 1.2: The majority of massive stars are formed in interacting binaries, experiencing some sort of stripping, accretion or merger (Figure 2 from Sana et al., 2012, reprinted with permission from AAAS). Cumulative histogram of evolutionary pathways followed by O stars in a given initial binary orbital period.

3 Uncertainties in binary stellar evolution

Many key aspects of binary stellar evolution remain extremely uncertain. In particular, several binary physics parameters are still unconstrained both theoretically and observationally across several orders of magnitude (e.g., Ivanova et al., 2013; Katsuda et al., 2018). Beyond this, single massive stellar physics such as stellar winds and remnant mass prescriptions are also poorly understood (e.g., Beasor et al., 2020; Fryer et al., 2012).

In Figure 1.3, we illustrate the major uncertainties that remain in binary stellar evolution through a simple schematic of a potential evolutionary pathway of a binary. The initially more massive star evolves more quickly, losing mass to stellar winds and expanding over the course of its evolution. The primary star expands to such an extent that not all material remains bound to it, initiating mass transfer. This mass transfer becomes unstable, leading to a common envelope, which is successfully ejected. Finally, the primary star reaches core collapse and unbinds the binary with its supernova. Each of these processes has some degree of uncertainty (annotated in Figure 1.3, which we outline below.

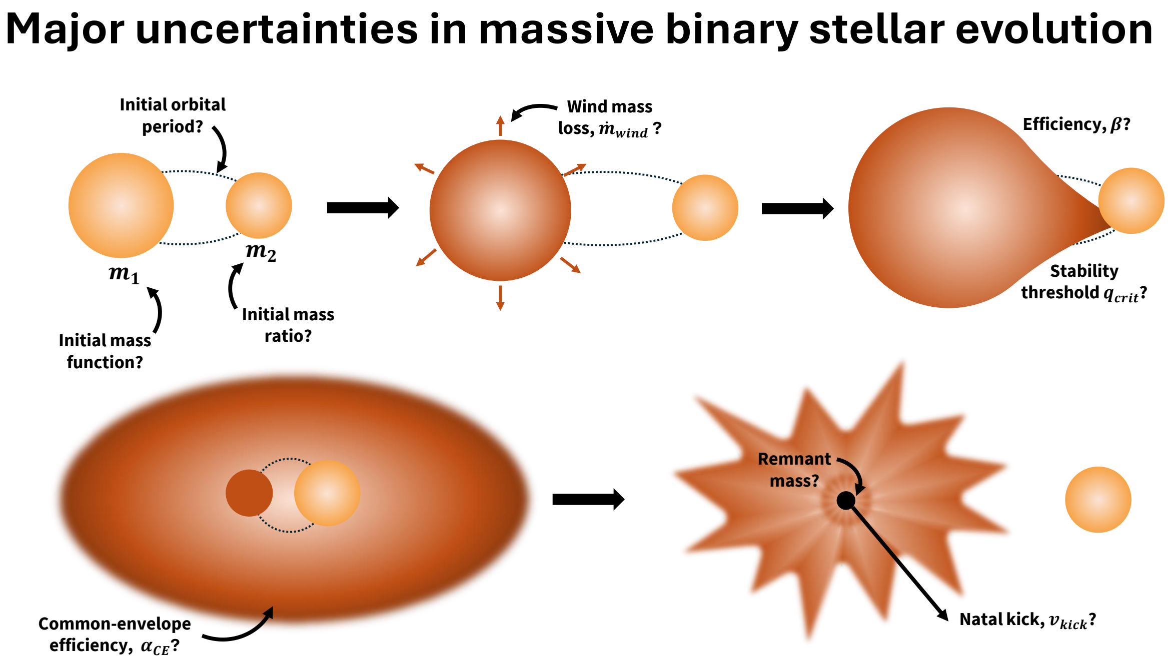

Figure 1.3: A schematic of the major uncertainties in binary stellar evolution, illustrated through the potential evolutionary pathway of a binary. After the binary starts with some initial parameters, the primary star evolves, losing mass to stellar wind and initiating mass transfer after Roche-lobe overflow. This mass transfer eventually becomes unstable, leading to a common-envelope which is ejected, allowing the primary start to reach core collapse and unbind the binary orbit. The major uncertainties in these processes are annotated with arrows and discussed in detail in Section 3.

In the following subsections we examine each the primary areas of uncertainty in binary interactions. For each area, we first introduce our current understanding of the phenomenon, then explain its importance in massive stellar evolution and highlight the major uncertainties in the process.

3.1 Initial populations

Before even examining how the stars in a binary evolve, it is worth considering the distribution of the *initial* populations. How many massive stars are formed? What are the relative masses of their companions, and the separations of the two? These distributions can have a significant impact on not only the prevalence of massive stars, but also the frequency at which they are found in interacting binaries.

3.1.1 Initial mass function

The phenomenon — The initial mass function (IMF) defines the frequency at which stars of a given mass are formed. It is generally written in the form

$$\zeta(m) \propto m^{-\alpha_{\mathrm{IMF}}}, \tag{1.1}$$

where m is the mass of the star and α_{IMF} is a mass-dependent constant. Generally, less massive stars are much more common than their massive counterparts. It was first measured in the solar neighbourhood by Salpeter (1955) and modelled as a two-part power law. This measurement was later expanded by Kroupa (2001) to a four-part broken power law.

Impact on massive stars — The slope of the high-mass end of the IMF is incredibly impactful for massive stars. A more top-heavy IMF, in which more stars are formed with high masses, leads to an increased number of binary transients, supernovae and therefore also feedback in galaxies. It also has important implications for the amount of ionising radiation present in a galaxy given that the primary source of this is self-stripping massive stars or stripped cores from binaries, both of which are more common in a top-heavy IMF.

Primary uncertainties — Given the impact of the high-mass end of the IMF, improving constraints on this value is of immediate importance. However, it has been shown that the literature values of uncertainties on this value have been consistently underestimated, such that ∼3/4 of quoted uncertainties are smaller than the theoretical lower limit (Weisz et al., 2013). Moreover, the universality of the IMF is a highly debated and controversial topic (e.g., Bastian et al., 2010). Though many early works argued that the IMF is constant across environments (Salpeter, 1955; Kroupa, 2001), more recent observations have put pressure on this assertion. Observations of 30 Doradus showed an excess of massive stars (Schneider et al., 2018a,b), while there was already evidence that the IMF may be dependent on metallicity (Marks et al., 2012) or the local star formation rate (Gunawardhana et al., 2011; Jeřábková et al., 2018).

3.1.2 Initial orbital periods and mass ratios

The phenomenon — Similar to the initial mass function, initial orbital periods and mass ratios are typically defined as a power law. Orbital periods are assumed to be distributed as $p(P) \propto (\log_{10}(P/\,\mathrm{days}))^{\pi}$ and mass ratios as $p(q) \propto q^{\kappa}$, where $q \equiv m_2/m_1$, and π and κ are constants. Earlier works measured these power laws independently, finding values for κ (**?**Sana et al., 2012; Kobulnicky & Fryer, 2007) and π (Kobulnicky & Fryer, 2007; Mason et al., 2009; Sana et al., 2012) separately. However, more recently evidence has been presented that these distributions are strongly correlated, and may not follow a single power law each (Moe & Di Stefano, 2017).

Impact on massive stars — Changes to the orbital period distribution of binaries directly alters the fraction of systems that will interact. A wider orbital period distribution reduces the number of instances of mass transfer and hence de-emphasises the impact of binary evolution. The initial mass ratio distribution can heavily impact the *type* of mass transfer that occurs. As we will discuss in Section 3.2, the mass ratio of a system is very important for determining the stability of mass transfer. And therefore a distribution that is skewed towards more equal mass ratios is more likely to produce stable mass transfer systems, and avoid creating a large number of tight binaries from common-envelope events.

Primary uncertainties — The precise values of κ and π are still not well constrained, particular in terms of their correlations. Moe & Di Stefano (2017) also highlights that there is an excess of "twin" systems with mass ratios close to unity and hence a true power law does not apply fully. Moreover, although mass ratios have a constrained range, the upper limit limit of orbital period distributions are less certain. For example, while Sana et al. (2012) argue an upper limit of $10^{3.5}$ day, de Mink & Belczynski (2015) extend this to 10^5 day. Furthermore, there is potentially a mass dependence to the orbital period distribution, where Renzo et al. (2019) argue that the power law slope should only be used for massive stars (with $m_1 > 15\,\mathrm{M}_\odot$), while lower masses use a flat-in-log distribution for orbital periods.

3.2 Stable mass transfer

A mass transfer event occurs between two stars when one overflows its Roche lobe as a result of expansion[1], such that the material is no longer bound to the expanding 'donor' star (Roche, 1849, 1850, 1851; Eggleton, 1983). The expansion of the donor star generally occurs in three distinct phases, each of which may initiate mass transfer depending on the initial orbital period of the binary (e.g., de Mink et al., 2008). The first phase is the slower expansion during the main sequence, which is referred to as case A mass transfer (Kippenhahn & Weigert, 1967). After the star's main sequence, it will expand again, rapidly, as it crosses the Hertzsprung gap and ascends the first giant branch. Mass transfer initiated at this stage is referred to as case B (Kippenhahn & Weigert, 1967). Lastly, during a star's asymptotic giant branch phase, after the onset of core helium burning, rapid and significant expansion occurs and can initiate case C mass transfer (Lauterborn, 1970). Each type of mass transfer begins with Roche Lobe overflow, which is followed by mass streaming through the inner Lagrange point (L1) and transferring to the companion 'accretor' star. There are several uncertainties regarding this process, most importantly regarding its stability, efficiency and rate of angular momentum loss, which we discuss in the subsections below.

3.2.1 Mass transfer stability

The phenomenon — The stability of mass transfer is driven by the response of the Roche lobe and radius of the donor star to the loss of mass (e.g., Benson, 1970; Paczynski, 1976). In the case of stable mass transfer, the donor star contracts slightly after transferring some mass, requiring the donor star to re-expand to continue transferring mass, following its nuclear timescale. In contrast, dynamically unstable mass transfer occurs when the donor star expands as it transfers mass, creating a runaway mass transfer scenario, most often leading to a common-envelope (see Section 3.3).

Impact on massive stars — Stable mass transfer and common envelopes have vastly different effects on binaries. Stable mass transfer can alter the mass ratio of a system and often widen the binary, whereas a common-envelope event can lead to a merger or significantly shrink the binary. Therefore, changing the criteria for the stability mass transfer for a given

---

[1] Binaries that evolve in dense environments or higher-order multiples may also experience Roche-lobe overflow without stellar expansion as a result of external processes (such as Von Zeipel-Kozai-Lidov oscillations) bringing the stars closer.

binary system can vastly alter the outcome of the system. On a population level this can affect the number of luminous red novae we may expect to detect from stellar mergers (e.g., Stegmann et al., 2024), as well as alter the rates of gravitational-wave mergers detectable by *LIGO-Virgo-KAGRA* (e.g., Broekgaarden et al., 2022).

Primary uncertainties — The response of the donor star and its Roche lobe to mass transfer depends on several factors, particularly the mass ratio of the system (more unequal mass ratios lead to more instability) and the evolutionary state of the donor (more rapidly expanding donors are more unstable, e.g. case A is generally more stable). In practice, the stability of mass transfer in population synthesis models is determined by tables of "critical mass ratios", which have different values given the stellar type of the donor and have been tabulated in various works (e.g., Hjellming & Webbink, 1987; Soberman et al., 1997; Belczynski et al., 2008; Claeys et al., 2014; Neijssel et al., 2019). The comparison of the mass ratio of the binary to the critical mass ratio at the onset of Roche lobe overflow is used to determine whether the mass transfer will proceed stably. The precise values of these critical mass ratio values are still uncertain and have a strong impact of which binaries proceed through common-envelopes, which changes the fraction of close binaries and stellar mergers. Furthermore, the use of critical mass ratios as a criterion for stability may over-simplify the problem and be overly granular in the definitions of different stellar evolutionary states.

3.2.2 Mass transfer efficiency

The phenomenon — The efficiency of mass transfer, $\beta \equiv \Delta M_{\rm accretor}/\Delta M_{\rm donor}$, is defined as the fraction of mass that is transferred by the donor ($\Delta M_{\rm donor}$) which is accreted by the companion ($\Delta M_{\rm accretor}$). $\beta = 1$ indicates fully conservative mass transfer, in which all mass is accreted by the companion, whilst $\beta = 0$ is fully non-conservative, in which all mass is lost from the system. The efficiency of mass transfer is determined by the accretor's ability to accept incoming mass (e.g., Paczynski, 1976). At first this is governed by comparing the mass transfer rate to the Kelvin-Helmholtz timescale of the accretor, since this timescale dictates the rate at which the star can thermally readjust to the new mass. Yet as the accretor gains angular momentum, it spins up and the efficiency of mass transfer becomes determined by the rate of the dominated rotationally-enhanced mass loss once it nears critical rotation (e.g., Petrovic et al., 2005). In the extreme case of mass transfer onto a compact object, the accretion may become Eddington-limited instead, meaning that the accretion is suppressed by the radiation pressure from an accretion disc (Eddington, 1926).

Impact on massive stars — The loss of mass from the binary system in this manner can significantly alter the subsequent evolution of the system. With lower values of β, the accretor star will become less massive which may prevent a massive star from reaching core collapse and hence β has a strong impact on the rate of supernovae (see Chapter 7). Alternatively, increasing β could lead the accretor star to expand more, triggering a second (reverse) phase of mass transfer, which could cause the two stars to merge. Beyond the changes from the mass lost from the system, angular momentum is also carried away, as we will discuss in Section 3.2.3. For these reasons, constraints on the value of β for different evolutionary phases is essential to understanding the evolution of interacting binaries.

Primary uncertainties — Most population synthesis works assume that β is set such that the amount of mass accreted during Roche-lobe overflow is limited to 10x the thermal rate of the accretor for main sequence, Hertzsprung gap and core helium burning stars and unlimited for giant branch stars (e.g. Kippenhahn & Weigert, 1967; Schneider et al., 2015). However, it is possible the value should depend on more than the stellar type of the binary and the thermal timescale. Moreover, in practice the inclusion of the criteria related to the rotation of the star can prove difficult. In detailed 1D stellar evolution models (discussed in more detail in Section 4.1) the conditions for spin-up and rotation modelling can be difficult since it does not allow for 3D turbulent motion which can play a role in these areas. Moreover, in population synthesis models (see Section 4.2), stellar rotation is usually not tracked at all.

3.2.3 Angular momentum loss during mass transfer

The phenomenon — For non-conservative mass transfer, some of the transferred mass is lost from the system and with it, its angular momentum. This angular momentum loss can occur through several different 'modes' (Soberman et al., 1997). When mass is lost rapidly from the vicinity of the donor star during mass transfer this is known as the Jeans (or "fast") mode (Jeans, 1924, 1925; Huang, 1963). Alternatively, the mass can be lost from the accretor as an isotropic wind during mass transfer, where this enhanced stellar wind is driven by the accretion of material. This mode is currently known as isotropic re-emission, but has been referred to as type B mass loss in the past (e.g., Kuiper, 1941). Another possibility is that as mass is transferred to the accretor, a fraction is lost through the outer lagrangian point (L2), forming an accretion ring or disc around the binary (e.g., Struve, 1941; Frank et al., 1985). This disc can introduce tidal effects onto the donor star, which further transfers angular momentum (e.g., Lin & Papaloizou, 1979; Artymowicz & Lubow, 1994; Zapartas et al., 2017). Each of these modes may occur during non-conservative mass transfer, and in many cases a combination may be possible (Soberman et al., 1997). In population synthesis, it is not possible to track the exact nature of mass transfer and instead several approximations are made. A common method for account for angular momentum loss is to introduce γ, which is defined to be the fraction of the specific angular momentum of the binary that is carried away by the ejected matter. In this way, γ can be defined in terms of the mass of the two stars for each mode (as well as the relative size of the accretion disc and the binary orbit).

Impact on massive stars — Each mode of mass transfer strongly affects the way that angular momentum is lost and hence the post-mass-transfer orbital separation of the binary. Therefore angular momentum loss is particularly important for investigations concerned with the separations of binary stars or compact objects. For example, changes to the binary separation can alter the delay time between binary black hole formation and a gravitational-wave merger. Moreover, the ejection velocities of runaway stars produced via the binary supernova scenario are almost entirely determined by their orbital velocity and hence are highly dependent on which mode of mass transfer occurred in the past.

Primary uncertainties — Population synthesis models typically choose a fixed prescription for the mode of mass transfer that occurs in binaries. This is clearly not true in reality, where a combination of different modes may occur for different types of mass transfer, which requires further work to investigate. However, there is currently no clear prescription for selecting the particular mode of mass transfer that occurs for a given system, so population synthesis is limited to constraining the *dominant* mode of mass transfer.

3.2.4 Rejuvenation

The phenomenon — Massive stars that accrete mass from a companion need to compensate for the additional material deposited on their surfaces. The accretion of material drives the star out of thermal equilibrium, increasing its luminosity. The additional luminosity allows a greater region of the star to sustain convection, thereby increases the size of the convective core. As the core expands it draws in fresh hydrogen-rich material to the potentially depleted core, which can extend the lifetime of the star. This process is referred to as rejuvenation and has been detailed in many previous works (e.g., Neo et al., 1977; Hellings, 1983; Braun & Langer, 1995; Cantiello et al., 2007; Schneider et al., 2016; Staritsin, 2019; Renzo et al., 2023; Lau et al., 2024).

Impact on massive stars — The altering of stellar lifetimes delays has important implications for exactly when future interactions or transient events within a binary will occur. For example, a significantly rejuvenated companion may reach core collapse at a later time, giving the star longer to disperse to less dense regime in the ISM and have a very different impact on the feedback in the galaxy.

Primary uncertainties — The exact extent to which rejuvenation can impact the future lifetime of a star is debated (Neo et al., 1977; Schneider et al., 2016; Renzo et al., 2023), where some authors argue the effect is negligible, whilst others suggest it could range up to several millions of years. These differences are often a result of differing choices for convective boundary mixing and overshooting prescriptions (e.g., Renzo et al., 2023).

3.3 Common-envelope evolution

The phenomenon — Dynamically unstable mass transfer leads to a runaway process in which a giant star (the evolved primary star in a binary) may engulf its companion within its envelope material. This engulfment results in a short-lived phase in which the core of the primary star and its companion orbit together within a common-envelope (CE) (e.g., Paczynski, 1976). The extreme drag forces dramatically reduce the orbital separation, depositing energy into the envelope material. In a successful CE event this energy is sufficient to eject the envelope and create a tightly bound binary. Alternatively, a CE may "fail" and instead lead to a stellar merger.

Impact on massive stars — A clear understanding of CEs is essential to explaining a variety of binaries involving compact objects. From cataclysmic variables and type Ia supernovae, to X-ray binaries and gravitational-wave mergers (e.g., Webbink, 1984; Wang, 2018; Tauris & van den Heuvel, 2014; Abbott et al., 2016), each requires a present day orbital separation significantly smaller than the combined zero-age main-sequence radii of their progenitors. The solution to this separation problem is the common envelope. A common envelope can quickly reduce binary orbital separations late in the evolution of the system. Moreover, failed CEs, in which a stellar merger occurs, are also important parts of binary evolution

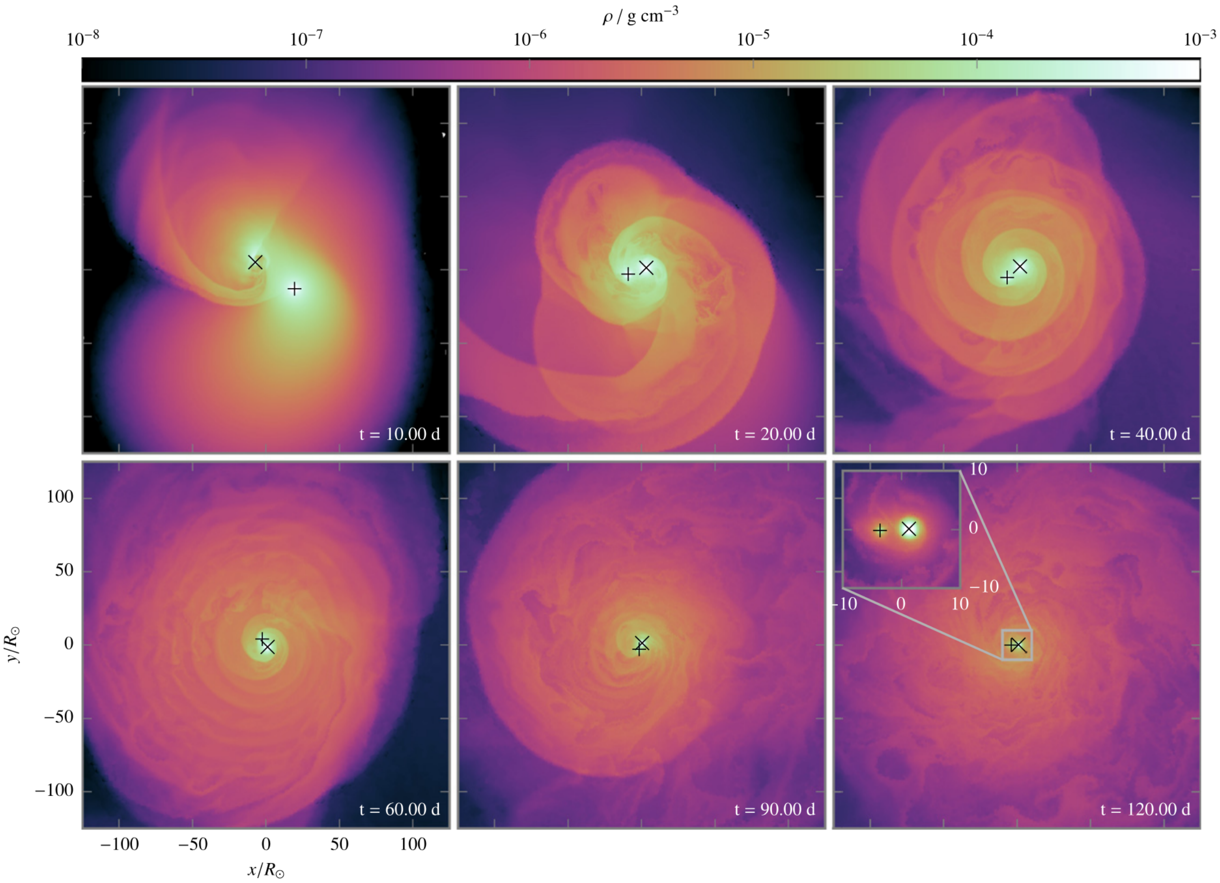

Figure 1.4: A series of snapshots of the density profile of a common-envelope event between a red giant donor (plus marker) and main sequence star (cross marker) evolved in the moving-mesh code `AREPO` (Figure 3 from Ohlmann et al., 2016).

since they produce luminous red novae (e.g., Munari et al., 2002; Soker & Tylenda, 2003) and the resulting merger products can form interesting astrophysical populations such as Thorne–Żytkow objects (Thorne & Zytkow, 1975). Given the range of phenomena that rely upon this phase, understanding common envelopes is arguably the most important outstanding task in binary evolution.

Primary uncertainties — Despite the wide-ranging implications of this phase, common envelopes remain poorly understood. It is agreed that the common envelope is initiated by a dynamical timescale plunge-in and inspiral, which is followed by a thermal timescale self-regulating phase until the successful ejection of the envelope or a merger occurs (e.g., Röpke & De Marco, 2023). However, an array of uncertainties remain, including: the exact timescales involved, fraction of orbital energy used to unbind the envelope, thermal effects from heating in the envelope, the binding energy of the envelope, and how prior accretion events may alter the envelope (e.g., Ivanova et al., 2013, 2020; Röpke & De Marco, 2023; Renzo et al., 2023; Landri et al., 2025).

Detailed 3D numerical simulations represent the state-of-the-art understanding of common-envelope events, but remain too expensive to run at the scale of populations of binaries (e.g., Röpke & De Marco, 2023). An alternative prescription used in population synthesis studies is the α-λ prescription (Webbink, 1984; de Kool, 1990). In this prescription, the CE is parameterised in terms of α, the fraction of initial orbital energy available to unbind the envelope, and λ, is a structure parameter that modifies the binding energy of the envelope. This method allows one to calculate whether a CE will succeed based on the pre-CE parameters, and if so, what it's post-CE binary configuration will be. This parameterisation likely oversimplifies the problem, the symmetry of a 1D prescription clearly fails to encapsulate the detail of the problem visible in Figure 1.4. Yet this method provides a useful tool for performing population level studies and gaining insight on the problem through constraints on α and λ.

The common-envelope α parameter in particular is poorly constrained. Although one may expect that it ranges from 0 to 1, it can formally exceed 1 to account for the fact that other energy sources beyond the orbital energy alone may contribute to unbinding the envelope (e.g., Ivanova et al., 2013). Previous works have claimed a range of α values are necessary to produce observations, with some finding it must be below $0.6 - 1$ (Iaconi & De Marco, 2019), while others suggest it is $\alpha = 5$ (Fragos et al., 2019) or even as large as $\alpha = 14$ (El-Badry et al., 2023). Altering this parameter over this range has drastic effects on a range of populations, including the distribution and rate of binary black hole mergers (e.g., Broekgaarden et al., 2022). Further constraints on common envelope events from population synthesis that cover a range of binary populations are essential to our understanding of massive binary stars.

3.4 Remnant mass prescriptions

The phenomenon — Low and intermediate mass stars end their lives as white dwarfs, whilst more massive stars may form neutron stars or black holes as their stellar remnants. The theoretical mapping of a progenitor's final core mass to the mass of its stellar remnant is often referred to as a remnant mass function. There are currently no direct observations that can explicitly establish a relationship for the link between core masses and remnant masses and thus we must rely upon simple analytic models or indirectly constrained relationships.

Of particular uncertainty is the delineation between the formation of a neutron star or a black hole from a massive star. Many introductory texts imply that there are simply two contiguous zero-age main-sequence mass ranges that form neutron stars ($8 \leq M_{\rm ZAMS}/\,\mathrm{M}_\odot \leq 20$) and black holes ($M_{\rm ZAMS} > 20\,\mathrm{M}_\odot$). However, the true picture is more complicated, with the final remnant outcome showing a stronger dependence on the presupernova core's "compactness" (e.g., O'Connor & Ott, 2011; Sukhbold & Woosley, 2014). This compactness parameter is defined as

$$\xi_M = \frac{M/\,\mathrm{M}_\odot}{R(M)/1000\,\mathrm{km}}, \tag{1.2}$$

where M is the mass coordinate where the compactness is evaluated (typically $2.5\,\mathrm{M}_\odot$ to ensure the entire iron core is encapsulated) and R is the radius at a given mass coordinate. It has been shown through numerous works that the compactness parameter has a non-monotonic relationship with core mass, and as such a more massive (but critically, less compact) may produce a neutron star, whilst a less massive core creates a black hole. Very recently, it has been argued that this non-monotonicity is a result of the timing and energetics of the last nuclear burning phases (e.g., Laplace et al., 2025).

Impact on massive stars — Unsurprisingly, changes in the remnant mass prescription have important effects on compact object populations, particularly those detectable via gravitational waves (e.g., Broekgaarden et al., 2022). These observed populations are potentially constraining for our different prescriptions, especially since some models predict stark features in the black hole mass distribution. For example, the Fryer *rapid* prescription predicts a gap in the mass distribution between $\sim 2-5\,\mathrm{M}_\odot$, whilst the *delayed* prescription fills this region in parameter space. This gap has potentially been observed in galactic x-ray binaries (e.g., Farr et al., 2011) and wide neutron star-star binaries detected with *Gaia* (e.g., El-Badry et al., 2024, Figure 18). However, more data is necessary to conclusively remark upon the existence of this gap.

Primary uncertainties — In population synthesis simulations, it is necessary to use a remnant mass prescription, which approximates how compactness may influence the final remnant mass (since the detailed stellar structure of each star is not simulated). These prescriptions are often a simple one-to-one relationship between pre-supernova core masses and remnant masses (e.g., Fryer et al., 2012), though there have been attempt to create less deterministic prescriptions which better capture some of the non-monotonicity of stellar compactness (e.g., Mandel & Müller, 2020). Figure 1.5 illustrates a series of prescriptions that are implemented in the `COMPAS` rapid binary population synthesis code (Riley et al., 2022). This figure highlights the importance of metallicity on the remnant mass (top right panel) as well as the stark differences between different prescriptions (bottom panel), which emphasises the high uncertainty we have on remnant masses.

3.5 Supernova natal kicks

The phenomenon — Observations of pulsars in the Milky Way necessitate that they have a high velocity dispersion and hence that these objects acquire high velocities upon formation (e.g., Gunn & Ostriker, 1970; Hobbs et al., 2005; Igoshev, 2020). For compact objects formed via isolated binary evolution, these velocities are expected to be imparted at the moment of core-collapse. Even for a symmetric explosion, the instantaneous mass loss from a supernova changes the post-supernova orbital parameters (Blaauw, 1961; Boersma, 1961). This mass loss is typically insufficient to disrupt the orbit, since the ejecta mass is often limited by loss during previous mass transfer phases, but in some cases (∼16%) mass loss can unbind the binary and result in both binary constituents being ejected with their pre-supernova orbital velocity (Renzo et al., 2019).

Supernova natal kicks arise due to asymmetry in the explosion, which imparts a significant velocity kick on the newly formed compact object (e.g., Shklovskii, 1970; Lyne & Lorimer, 1994; Janka, 2013, 2017), which in the majority of cases is strong enough to disrupt a binary (e.g., Eldridge et al., 2011; Renzo et al., 2019). Similar asymmetry in neutrino emission can result in additional natal kicks, which can dominate in the case of full fallback (e.g., Janka & Mueller, 1994; Burrows & Hayes, 1996; Scheck et al., 2006; Wongwathanarat et al., 2013; Coleman & Burrows, 2022; Vigna-Gómez et al., 2024).

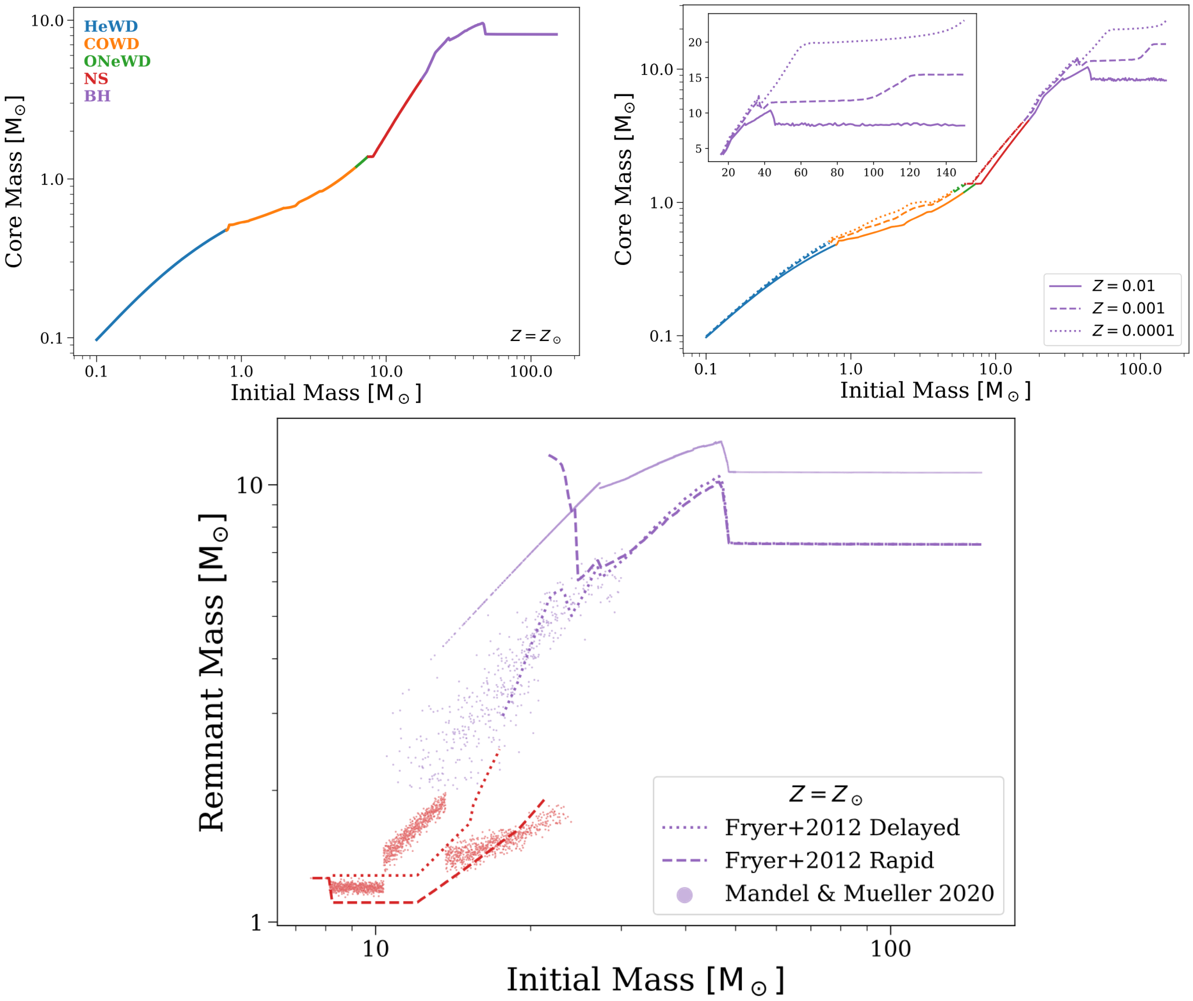

Figure 1.5: Various prescriptions for remnant masses exist, each of which have dependences on metallicity (Figure 8 from Riley et al., 2022). **Top left:** The relation between the initial (ZAMS) mass of single stars at solar metallicity (Z = 0.0142) and either the WD remnant mass or the CO core mass at the moment of compact object formation for NSs and BHs, assuming the COMPAS default model, where colours denote the remnant type. **Top right**: The same relation at Z = 0.01 (solid), Z = 0.001 (dashed) and Z = 0.0001 (dotted) to show the metallicity dependence. The inset panel shows the same quantities as the main panel, but with linear axes, focusing on BH masses. **Bottom:** The initial–compact object mass relation at $Z_\odot$, plotted for different remnant mass models: the delayed (dotted line) and rapid (dashed line) prescriptions from Fryer et al. (2012), and the stochastic model from Mandel & Müller (2020) (points converging to a solid line).

Impact on massive stars — Supernova natal kicks significantly alter the spatial distribution of compact objects (e.g., Sweeney et al., 2022) and determine the fraction of binaries that remain bound (e.g., Renzo et al., 2019). Even for binaries that are not unbound by a supernova natal kick, their post-supernova orbital parameters can be significantly altered. This phenomenon has been investigated by a myriad of earlier works (Flannery & van den Heuvel, 1975; Sutantyo, 1978; Hills, 1983; Dewey & Cordes, 1987; Wijers et al., 1992; Brandt & Podsiadlowski, 1995; Kalogera, 1996; Tauris & Takens, 1998; Hurley et al., 2002; Pfahl et al., 2002; Belczynski et al., 2008; Kiel & Hurley, 2009). These investigations account for a variety of effects, with the most extensive handling eccentric pre-supernova systems, including the effect of instantaneous symmetric mass loss, the impact of the natal kick on both the compact object and disruption of the system, the supernova blast wave impulse on the secondary star and collisions between the compact object and star. For a detailed derivation, which explains how each of these effects impacts a binary in clear and compact vector notation, we refer the reader to Appendix B of Pfahl et al. (2002).

Primary uncertainties — The magnitude and distribution of supernova natal kicks remains very uncertain. Earlier estimates found that the natal kick distribution follows a single maxwellian distribution with $\sigma = 265\,\mathrm{km\,s^{-1}}$ (e.g., Hobbs et al., 2005). More recently, evidence has been found for a bimodal distribution with two maxwellians, where the additional component has $\sigma = 30\,\mathrm{km\,s^{-1}}$ (e.g., Igoshev, 2020). This dichotomy is predicted theoretically to be populated by electron-capture supernovae (ECSN) and ultra-stripped supernovae (USSN), which are found to have reduced asymmetry and therefore weaker kicks (e.g., Podsiadlowski et al., 2004; Ivanova et al., 2008; Suwa et al., 2015).

Natal kick distributions are primarily constrained by the motions of pulsars and therefore may only be applicable to neutron stars. It is unclear whether black holes follow the same distribution given the small population of observed black holes (which are almost all in bound binaries). Some argue that black holes should have the same *momentum* distribution (thus scaling the kick down in magnitude due to the higher mass black holes) (e.g., Mandel & Müller, 2020; Riley et al., 2022). Others suggest it is important to account for how mass falling back into the black hole after core collapse may negate some of the asymmetry and therefore the kick a black hole receives (e.g., Fryer et al., 2012).

Overall, given that supernova natal kicks strongly influence the demographics of post-supernovae binaries, as well as the spatial distribution of compact objects, there is a clear opportunity to further constrain our models for these events via binary populations.

3.6 Stellar winds

The phenomenon — Massive stars can lose significant amounts of material to stellar winds, particularly through line-driven winds. The mechanism behind line-driven winds can be understood as follows: heavily ionised metals (e.g., carbon, nitrogen, oxygen, iron) are accelerated through repeated absorption of photons, which transfers momentum from the radiation to the gas. Even though only a small fraction of ions are accelerated, their collisions and interactions with other particles drags the gas along in tandem (Sobolev, 1960; Morton, 1967; Lucy & Solomon, 1970; Castor et al., 1975). This process is strongest in massive stars since stellar radiation is peaked in the UV, where metal ions have their strongest spectral lines

(Castor et al., 1975). Given that these winds are dependent on metal ions, there is also a strong dependence on the metallicity of the star.

Impact on massive stars — These stellar winds are extremely important for the future evolution and fate of the star. For the most massive stars, particularly at high metallicity (since metal ions are more abundant), stellar winds can dramatically decrease the total mass of a star, and the exact strength of those winds may change the type of remnant that the star leaves behind. The mass lost also carries away angular momentum, and hence can change not only the spin-down rate of the star but also the separation of two stars in a binary system. Changes in separation can of course affect whether a star experiences mass transfer and whether it is stable (see Section 3.2). It is therefore clear that understanding stellar winds is critical to a clear picture of binary stellar evolution.

Primary uncertainties — Recent observations have improved our constraints on stellar winds. Stars in different evolutionary phases have different stellar winds and the major uncertainty that remains for binary stellar evolution is winds for stripped helium stars, in addition to the metallicity dependence of winds. Observations of stripped stars in the Magellanic Clouds provide stronger constraints on these winds, though still with uncertainties of nearly a factor of 5 in strength (Vink, 2017; Shenar et al., 2019; Hamann et al., 2019; Sander & Vink, 2020; Yang et al., 2023).

The comparison of population synthesis simulations to observations may provide an additional indirect constraint on the strength of stellar winds. Population synthesis codes use prescriptions for stellar winds (e.g., Vink et al., 2001) and consistently update to the latest models (e.g., Krtička & Kubát, 2017, 2018; Beasor et al., 2020; Sander & Vink, 2020; Krtička et al., 2024). Many predictions have shown that the rate of gravitational-wave sources are strongly dependent on the strength of stellar winds (e.g., Broekgaarden et al., 2021, 2022; Wagg et al., 2022c). Given that winds can affect not only the rate of formation of these objects but also the orbital properties due to angular momentum loss, there are many further opportunities to understand this aspect of binary evolution through population synthesis.

4 Methods for constraining binary evolution models

In the sections above, we have discussed how massive stars play a central role in an array of astrophysical processes, are almost always formed in binaries, but that their binary evolution remains uncertain at several key stages. We now examine the avenues that exist for constraining binary evolution models.

In this thesis, we employ two main methodologies for constraining binary evolution: detailed 1D stellar evolution models and population synthesis. In addition to these techniques, one can also apply full 3D hydrodynamical simulations of stars and stellar explosions (e.g., Janka, 2012). This method is extremely expensive and can take many months to compute the core collapse of a massive star, but allows for asymmetrical effects that are present in massive stars. For the remainder of this thesis we will focus on the former two techniques.

4.1 Detailed 1D stellar evolution models

The computation of stellar evolution can be accomplished reasonably efficiently under several assumptions: spherical symmetry, hydrostatic equilibrium, an equation of state and that the star is fuelled by nuclear fusion. Under these assumptions one can write out the stellar structure equations, which account for mass continuity, hydrostatic equilibrium, conservation of energy and energy transport. We refer the reader to Chapter 10 of Levesque & Lamers (2017) for a detailed introduction to these equations and their boundary conditions.

In 1950s and 1960s several methods were presented for solving the stellar structure equation and computing the evolution of a star. Schwarzschild (1958) proposed a "shooting" method in which the equations are iteratively integrated from the core to the surface and back. Henyey et al. (1959, 1964) improved upon this by expressing the structure equation as a set of linear finite difference equations, which was more stable, converged faster and better handles nonlinearities in the structure. Over the next decade several codes were developed to work on stellar evolution (Iben & Ehrman, 1962; Hofmeister et al., 1964; Kippenhahn & Weigert, 1967). Many improvements were made to the numerical stability of these models as well as the efficiency of tracking shell burning in low-mass stars and the advanced stages of burning in massive stars (Sugimoto, 1970; Eggleton, 1971; Weaver et al., 1978).

Aside 4.1: The Sounds of Stellar Evolution

I was lucky to meet Emmi Meyer-Hofmeister (the Hofmeister of Hofmeister et al. 1964) at the Max Planck Institute in Munich in the summer of 2023. She was a student working with Rudolf Kippenhahn during the early development of large-scale stellar evolution computation, which was done with computers large enough to fill a room. She regaled me with an amusing story about Kippenhahn, who would sometimes help with debugging their evolution. During the computation, which could last for many hours, Kippenhahn would often return home. So, if something was confusing in the output, they would need to call Kippenhahn on the telephone. They would then hold the receiver up to the computer, since he could tell what stage of stellar evolution was occurring based on the sounds it was making and diagnose the problem!

Now, at present day, there are many stellar evolution codes which have been used to compute detailed evolution tracks, such as the Cambridge STARS code (Eggleton, 1971; Pols et al., 1995), PARSEC (Bressan et al., 2012; Nguyen et al., 2022) and the Geneva code GENEC (Eggenberger et al., 2008).

Of particular note is Modules for Experiments in Stellar Astrophysics (`MESA`, Paxton et al., 2011, 2013, 2015, 2018, 2019; Jermyn et al., 2023), which is an *open-source* 1D stellar evolution code which now has a community of thousands of users. One can use this code to evolve everything from low-mass stars to extremely massive binaries, including the effects of rotation and using complex nuclear reaction networks. It is also coupled to `GYRE` (Townsend & Teitler, 2013; Townsend et al., 2018; Goldstein & Townsend, 2020; Sun et al., 2023), which allows one to compute the properties of pulsations in stars which are evolved in `MESA`. We employ these codes to investigate stellar pulsations in Chapter 4.

4.1.1 Advantages and limitations

Internal structure information — 1D stellar evolution codes give you access to the internal structure of the star at any moment during its evolution. As such, one can compute the extent of a convective core, the fraction of mass contained in hydrogen shell-burning region or track the abundance of elements involve in the CNO cycle. This level of detail is not possible in simpler population synthesis models and enables a series of more detailed investigations.

For example, Laplace et al. (2021) investigates how binary stripping alters not only the envelope structure, but also the core of a star. They use `MESA` models to compute the interior composition of a variety of stars at the onset of core collapse and compare how binary stripping changes these compositions. Figure 1.6 highlights the level of detail that's available in these models. One can track the interior composition of a variety of important layers in the star, which can be influenced by binary interactions. They show that binary interactions lead to an extended carbon-oxygen gradient at the outer edge of the oxygen-rich layer (labelled as II), as well as a greater abundance of carbon in the core. They used these detailed models to predict that the supernova yields of binary-stripped progenitors would be significantly altered.

Flexibility of physics assumptions — `MESA` has an *extensive* array of settings for evolving stars. These can be as general as the efficiency of mass transfer (which we discussed in Section 3.2), or as detailed as the convective core overshooting parameters. Modern population synthesis codes generally rely upon a single set of stellar tracks that made fixed choices for many of these more detailed parameters[2].

Long runtimes — A disadvantage of these 1D stellar evolution codes is the runtime. Evolving a single binary star with a single phase of mass transfer can still take several hours even when using tens of cores on a computing cluster. As a result, investigating large populations of sources can become intractable when using these more detailed models.

4.2 Rapid binary population synthesis

Given the present-day computational limitations, large-scale population-level simulations require a more approximate approach to stellar evolution to allow for reduced runtimes. Various attempts were made in the past to use lookup tables and simple approximations (e.g., Whyte & Eggleton, 1985; Eggleton et al., 1989; Schaller et al., 1992), but the true revolution in population synthesis occurred with the analytic fitting formulae provided by Hurley et al. (2000), reproducing the single stellar tracks provided in Pols et al. (1998). These formulae can be used to extremely rapidly compute how the structure of a star changes during its evolution, approximating the 1D stellar model to a reasonable level of accuracy. The initial formulae were quickly followed by the release of `BSE` (Hurley et al., 2002), which adjusted these results to account for binary interactions such as mass transfer. This code was able to rapidly evolve millions of *binary* stars in seconds, a stark contrast to the runtime of more detailed models. This reduced runtime enabled an entirely new approach for constraining binary evolution models by focusing on the bulk properties of a population of systems rather

---

[2] Some newer codes could, in principle, allow for changes to these assumptions (e.g., Fragos et al., 2023; Andrews et al., 2024; Agrawal et al., 2020, 2023), but this have not yet been investigated.

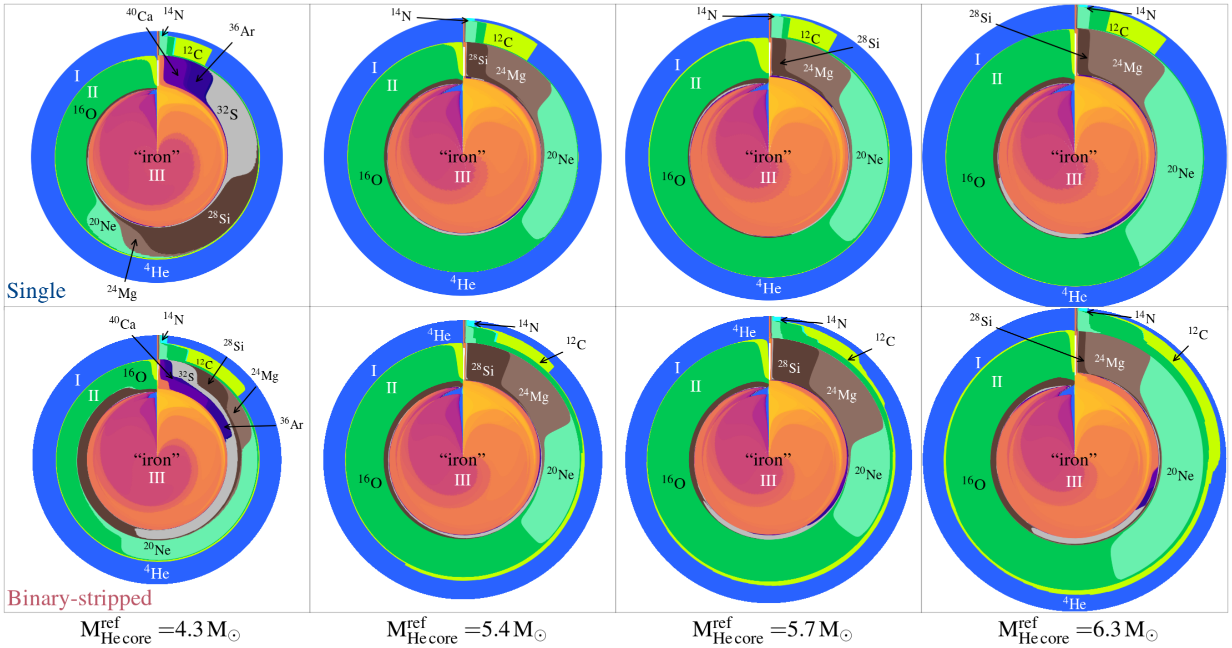

Figure 1.6: Detailed stellar evolution models allow the investigation of interior compositions and can demonstrate how binary interactions can alter the pre-supernova core structure of massive stars (Figure 2 from Laplace et al., 2021, reproduced with permission from Astronomy & Astrophysics, ©ESO). Composition profiles of single (top row) and binary-stripped (bottom row) star models at the onset of core collapse. Each colour represents an isotope and the surface area spanned by the colour is proportional to the mass of this isotope in the star. The radius of each diagram is proportional to the final helium core mass. The hydrogen-rich envelope of single stellar models is not shown. Three prominent regions are marked with roman numerals: (I) a helium-rich layer, (II) an oxygen-rich layer, and (III) an inner iron-rich zone. Binary-stripped star models contain an extended carbon-oxygen gradient at the edge of the oxygen-rich region (II) that is absent in single-star models.

than aiming to reproduce a specific system.

The release of `BSE` quickly led to the development of several related codes that adapted the `BSE` equations to new regimes or use the same methodology. These include `binary_c` (Izzard et al., 2004, 2006, 2009, 2018; Izzard & Jermyn, 2023), `SeBa` (Portegies Zwart & Verbunt, 1996; Toonen et al., 2012), and `Startrack` (Belczynski et al., 2008). In addition, other codes move beyond fitting formulae entirely, instead interpolating between updated underlying stellar tracks, such as `SEVN` (Spera et al., 2022; Iorio et al., 2023) and `POSYDON` (Fragos et al., 2023; Andrews et al., 2024), or combining with stellar evolution models for synthetic atmospheric spectra, such as `BPASS` (Eldridge et al., 2017; Byrne et al., 2022). The codes using interpolation tend to bridge the gap between detailed stellar evolution codes and population synthesis, sacrificing some of the improvement in runtime to regain some of the detail associated with full 1D models.

Two codes in particular to highlight are `COMPAS` (Riley et al., 2022) and `COSMIC` (Breivik et al., 2020b), because we use these in the later chapters of this thesis. Both codes are extensions to the original `BSE` code that are fully open-source, where `COMPAS` is generally focused on massive gravitational-wave sources, while `COSMIC` has additional features for dynamic timesteps and reproducibility.

The power of binary population synthesis comes in offering constraints on uncertain binary physics parameters from large-scale populations. With the rapid runtimes of these codes, one can reproduce simulations of millions of binaries for a range of choices in these parameters and constrain the parameter space that is permitted by observations. Moreover, population synthesis studies can also highlight when certain predictions are *insensitive* to certain uncertainties in binary evolution. These same predictions would be computationally intractable with a full 1D stellar evolution code such as `MESA`.

4.2.1 Advantages and limitations

Rapid runtime — The critical advantage of rapid binary population synthesis codes is their runtime. Modern codes are able to evolve millions of binaries in under an hour on a typical laptop. This runtime is an incredible improvement relative to the several hours it may take to evolve *just one* binary with a 1D stellar evolution code. The reduced runtime enables a more detailed exploration of a range of binary physics parameters.

Population-level inference — The predictions from a rapid population code are able to reproduce the bulk properties of stellar populations in spite of approximating many of the processes in stellar evolution. In Figure 1.7, we illustrate the level of detail in single stellar tracks in a modern population synthesis code. This figure shows the evolution tracks in `COMPAS` for a variety of initial masses and two different metallicities. The tracks are realistic, though clearly approximate in some areas (e.g., consider the extremely straight lines during Hertzsprung gap (HG) evolution). These tracks can therefore be compared well to existing populations and used to constrain binary evolution. Given the advent of Big Data and huge catalogues of stellar populations from observatories such as the *Vera C. Rubin Observatory*, *Gaia*, and *Chandra*, rapid binary population synthesis is well-poised to take advantage of these observations (see Section 4.2.2).

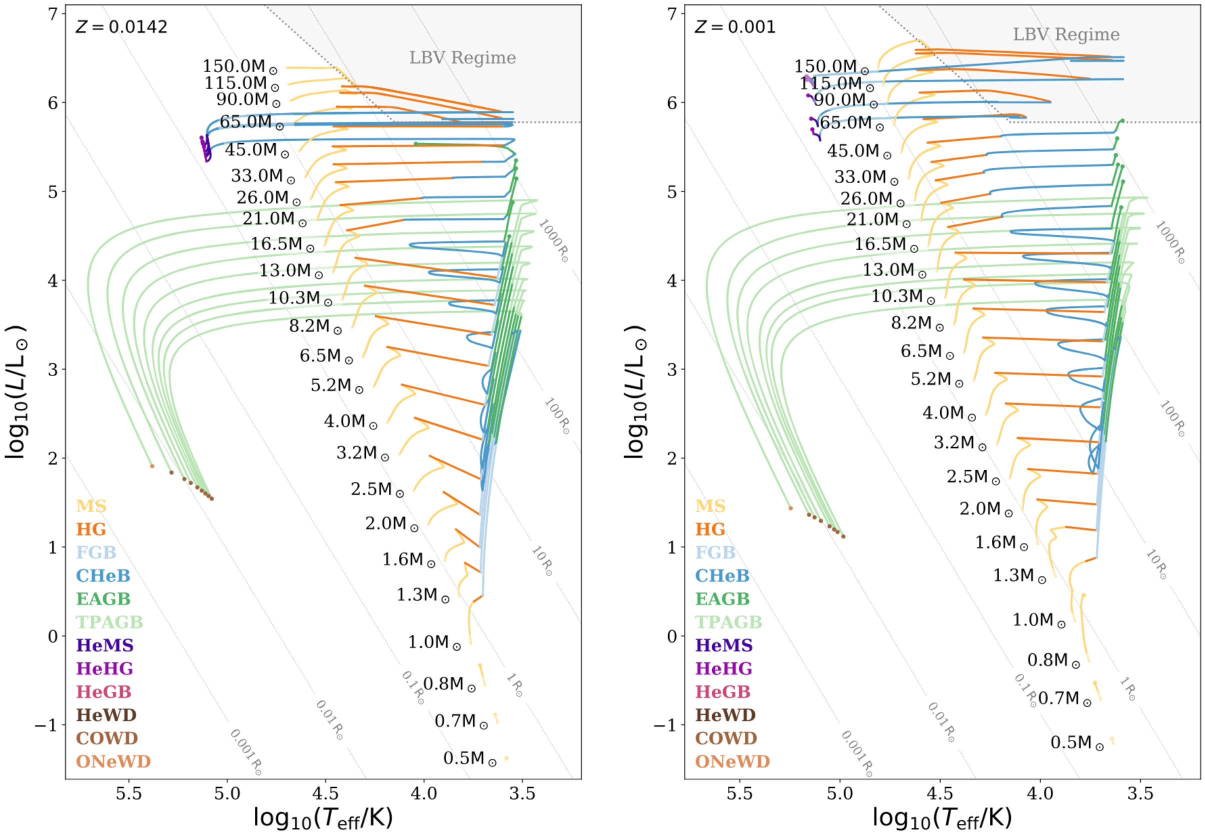

Figure 1.7: An example Hertzsprung-Russell diagram from the `COMPAS` code, which demonstrates the approximate stellar tracks used for rapid population synthesis (Figure 5 from Riley et al., 2022). Stellar tracks are shown for a series of different initial masses and points are coloured by their stellar type (abbreviations are translated in Table 1.1). The left panel is at solar metallicity, whilst the right is a lower metallicity example. The regime in which luminous blue variable winds are applied (the Humphrey-Davidson limit) are shown in a grey shaded region in the upper right. Grey dotted lines of constant stellar radii are shown as a guide.

Lack of internal structure — Rapid binary population synthesis only approximates the overall evolution of a star and thus at best tracks the mass of the helium core and the metallicity of a star. In this way, one cannot estimate the extent or abundance profile of a stellar envelope or predict in detail the yield from a given supernova.

Reliance on underlying stellar tracks — The accuracy and flexibility of a population synthesis code is directly limited by the underlying stellar tracks that it uses. For example, `BSE`-based codes rely upon the tracks from Pols et al. (1998), despite the improvements that have been made to our understanding of stellar evolution since they were computed. Moreover, the choices made in those tracks such as core overshooting parameters and the level of rejuvenation are fixed and as such cannot be altered as other parameters may be in population synthesis. This rigidity highlights the need to consider a variety of underlying stellar tracks, and use more up-to-date models, as suggested by the authors of `METISSE` and `POSYDON`.

Stellar evolution phase	Abbreviation
Main sequence, $M < 0.7\ \mathrm{M}_\odot$	MS
Main sequence, $M > 0.7\ \mathrm{M}_\odot$	MS
Hertzsprung gap	HG
First giant branch	FGB
Core helium burning	CHeB
Early asymptotic giant branch	EAGB
Thermally pulsing asymptotic giant branch	TPAGB
Helium main sequence	HeMS
Helium hertzsprung gap	HeHG
Helium giant branch	HeGB
Helium white dwarf	HeWD
Carbon-oxygen white dwarf	COWD
Oxygen-Neon white dwarf	ONeWD
Neutron star	NS
Black hole	BH
Massless remnant	MR

Table 1.1: Stellar phases and abbreviations used in most `BSE` based population synthesis codes (Hurley et al., 2000).

4.2.2 Potential observables

As noted above, rapid binary population synthesis is well-suited to comparison with existing observations. Let us briefly consider the range of potential observables that could be used for comparison with population synthesis.

Gravitational-wave sources — Both the final stages of inspiral and merger of two compact objects are detectable as gravitational-waves. The merger of two black holes was first detected a decade ago (Abbott et al., 2016) and since then more than one hundred gravitational-wave events have been detected (**?**). In the next two decades we expect to detect significantly more, particularly tens of thousands of galactic sources with the upcoming *LISA* mission (e.g., Amaro-Seoane et al., 2017; Nelemans et al., 2001; Korol et al., 2017; Breivik et al., 2020b).

Double compact objects represent an extreme tail of binary evolution, in which both stars have reached core collapse and remained bound to one another. The close separations necessary to reach a merger via gravitational-wave emission also imply that many systems will need to experience a common-envelope, or at least stable mass transfer. As such, the rates and demographics of these objects are highly sensitive to uncertainties in binary physics (e.g., Broekgaarden et al., 2021, 2022). This sensitivity means that they offer an excellent probe into our understanding of binary evolution.

Transient events — Massive (binary) stars lead to a variety of transient events, including supernovae, luminous red novae (as mergers), and short gamma-ray bursts (as the merger of two neutron stars) (e.g., Munari et al., 2002; Berger, 2014). The upcoming *Vera C. Rubin Observatory* will rapidly increase our populations of these transients. These populations are each sensitive to changes in binary physics. The strength of supernova natal kicks affects how many binary neutron stars can remain bound to one another, while changes to the stability of mass transfer can vastly change the rate of luminous red novae. Moreover, the timing and location of supernovae is strongly impacted by binary interactions, which can delay them by many millions of years (e.g., Zapartas et al., 2017).

Runaway stars and spatial distributions — Runaway stars are produced when the primary star in a binary reaches core collapse and unbinds the binary with its supernova natal kick. The secondary star is ejected with its pre-supernova orbital velocity and is defined as a runaway star when it is ejected at more than $30\,\mathrm{km\,s^{-1}}$ (e.g., Blaauw, 1961; Boersma, 1961; Renzo et al., 2019). Many runaway stars have been detected, particularly in the vicinity of star forming regions (e.g., Sana et al., 2022). The rates of these objects are sensitive to the strength of supernova natal kicks, while their velocities are strongly affected by aspects of binary evolution that change the orbital period of a binary (such as mass transfer and stellar winds).

5 The structure and outline of this thesis

At the end of this introduction you should now be familiar with the importance and prevalence of massive binary stars. We have considered how uncertain many binary processes remain, and discussed two methods by which we can improve our understanding. Over the course of this we consider the impact of binary interactions on a variety of astrophysical populations, from gravitational-wave sources, to pulsating massive stars, and supernovae using the methods outlined in Section 4. Through our investigations we aim to test the dependence of our results on the major uncertainties in binary physics that we detailed in Section 3. As a result our predictions can be compared against observed populations and used to constrain binary evolution models. The subsequent chapters are structured as follows:

- **Chapter 2-3: Galactic gravitational-wave sources**
  We first consider galactic gravitational-wave sources that will be detectable with the future spaced-based gravtiational-wave detector *LISA*. In Chapter 2 we introduce a new open-source code we developed for assessing the detectability of sources in *LISA*. Then in Chapter 3 we combine this code with rapid binary population synthesis to make predictions for the population of black hole and neutron star binaries that will be detectable by *LISA*. We investigate how these results depend on different aspects of binary physics and consider the prospects for using detections to constrain binary physics.

- **Chapter 4: Asteroseismology of accretor stars**
  Next, we transition from gravitational waves to gravity-wave pulsations in massive stars. We use the `MESA` 1D stellar evolution code to investigate how the internal structure and pulsations of a star change after accreting material from a comparison. We demonstrate for the first time that mass transfer leaves a significant imprint on the asteroseismic properties of a star. This imprint may shift with different types of mass transfer and thus pulsations could be used to better understand binary mass transfer physics. This result motivates a series of future studies expanding the simulation parameter space.

- **Chapter 5: Self-consistent population synthesis and galactic dynamics**
  Having considered the influence of binary interactions on the intrinsic properties of double compact objects and internal structure of massive stars, we next turn our focus to the *extrinsic* positions and kinematics of sources. We present `cogsworth`, a new open-source code for self-consistent population synthesis and galactic dynamics simulations. This code allows one to predict the positions and kinematics of massive binary products while accounting for their galactic context. We outline a series of potential use cases, including tracking the constituents of a dispersing star cluster and creating a synthetic colour-magnitude diagram.

- **Chapter 6: Stellar ejection velocities and supernova natal kicks**
  In this brief chapter, we highlight the counter-intuitively weak dependence of stellar ejection velocities on supernova natal kick magnitudes. We compare how stellar ejection velocities are calculated in three different population synthesis codes for a representative binary and confirm that they are in agreement. We additionally explain how the angle of the supernova natal kick can slightly affect the ejection velocity, as well as amend a prescription for supernova kicks from an earlier work.

- **Chapter 7: Timing and location of core-collapse supernovae**
  In the next chapter, we apply `cogsworth` to considering how binary interactions change the timing and location of core-collapse supernovae. We synthesise a suite of populations based on a hydrodynamical zoom-in simulation, varying binary physics, initial conditions and galaxy parameters. We show that the significant impacts of binary interactions are robust to these variations and conclude with a new analytic model for core-collapse supernova feedback to be used in hydrodynamical simulations.

- **Chapter 8: Future directions**
  In the penultimate chapter we ponder some potential future directions of this research. We focus on the timely opportunities to constrain `cogsworth` simulations with upcoming data on large-scale binary populations, particularly with *Gaia* DR4.

- **Chapter 9: Conclusions & Summary**
  Finally, we briefly summarise the main conclusions of this thesis.

> *Now this is not the end. It is not even the beginning of the end. But it is, perhaps, the end of the beginning.*
>
> – Winston Churchill

CHAPTER 2

`LEGWORK`: PREDICTING *LISA* STELLAR-ORIGIN GRAVITATIONAL-WAVE SOURCES

> *No one is finally dead until the ripples they cause in the world die away, until the clock wound up winds down, until the wine she made has finished its ferment, until the crop they planted is harvested. The span of someone's life is only the core of their actual existence.*
>
> – Terry Pratchett in *Reaper Man*

1 Introduction

The planned space-based gravitational wave detector LISA (Laser Interferometer Space Antenna) will present an entirely new view of gravitational waves by focusing on lower frequencies ($10^{-5} < f/\,\mathrm{Hz} < 10^{-1}$) than ground-based detectors. This will enable the study of many new source classes including mergers of supermassive black holes (e.g. Begelman et al., 1980; Klein et al., 2016; Bellovary et al., 2019), extreme mass ratio inspirals (e.g. Berti et al., 2006; Barack & Cutler, 2007; Babak et al., 2017; Moore et al., 2017), and cosmological GW backgrounds (e.g. Bartolo et al., 2016; Caprini et al., 2016; Caldwell et al., 2019). However, this frequency regime is also of interest for the detection of local stellar-mass binaries during their inspiral phase. LISA is expected to detect Galactic stellar-origin binaries containing combinations of white dwarfs, neutron stars, and black holes, ranging from the numerous double white dwarf population, to the rare but loud double black hole population.

The potential to detect stellar-origin sources with LISA has been studied in the past by various groups (e.g. Nelemans et al., 2001; Liu, 2009; Liu & Zhang, 2014; Ruiter et al., 2010; Belczynski et al., 2010; Yu & Jeffery, 2010; Nissanke et al., 2012). The recent direct detection of gravitational waves with ground-based detectors has led to renewed interest in this topic (e.g. Korol et al., 2017, 2018, 2019, 2020; Christian & Loeb, 2017; Kremer et al., 2017, 2018; Lamberts et al., 2018, 2019; Fang et al., 2019; Andrews et al., 2020; Lau et al., 2020; Breivik et al., 2020b,a; Roebber et al., 2020; Chen et al., 2020; Sesana et al., 2020; Shao & Li, 2021).

Each of these studies require making estimates of the signal-to-noise ratio of individual binary systems and possibly the slow gravitational wave inspiral that lead to the present-day parameters. So far, most studies made use of custom made codes which have not been made publicly available.

We believe that the large renewed interest in LISA and the stellar-origin sources it may detect will lead to many more studies in the near future that would need similar computations. This leads to a significant amount of redundancy which, at best results in extra work for each individual and at worst leads to an increased chance of introducing mistakes and inconsistencies when translating the necessary expressions to software.

`LEGWORK` is an open-source Python package designed to streamline the process of making predictions of LISA detection rates for stellar-origin binaries such that it is as fast, reliable and simple as possible. With `LEGWORK` one can evolve the orbits of a binary or a collection of binaries and calculate their strain amplitudes for any range of frequency harmonics. One can compute the sensitivity curve for LISA or other future gravitational wave detectors (e.g. TianQin's curve, or that of a custom instrument) and use it to compute the signal-to-noise ratio of a collection of sources. Furthermore, `LEGWORK` provides tools to visualise all of the results with easy-to-use plotting functions. Finally, `LEGWORK` is fully tested to check for consistency in the derivations described below.

Specifically, we implement the post-Newtonian expressions by Peters & Mathews (1963) and Peters (1964) for the evolution of binary orbits due to the emission of gravitational waves, equations for the strain amplitudes and signal-to-noise ratios of binaries from various papers (e.g. Flanagan & Hughes, 1998; Finn & Thorne, 2000; Cornish & Larson, 2003; Barack & Cutler, 2004; Moore et al., 2015) and approximations for the LISA and TianQin sensitivity curves given in Robson et al. (2019) and Huang et al. (2020) respectively. The post-Newtonian expressions are approximately of order 0.5 as they account for the orbital evolution from energy loss due to the emission of GWs, but don't include other effects such as spin-orbit coupling. We find this is an excellent approximation for stellar-origin sources in the LISA band as any higher order terms are either zero or negligible.

The open-source nature of the project means that new users as well as seasoned experts in the field can work together in a collaborative setting to consider new features and enhancements to the package as well as check the implementation. At the same time, with our thorough online documentation, derivations and tutorials, we hope LEGWORK can make this functionality more accessible to the broader scientific community.

We note that LEGWORK is not the only initiative of this kind. We highlight the "Gravitational Wave Universe Toolbox" presented by Yi et al. (2022) which was developed to simulate observations on the GW universe with different detectors covering the full gravitational-wave spectrum and source classes. We also highlight "GWPlotter" by (Moore et al., 2015), which provides an interactive plotting tool to compare the sensitivity of different gravitational wave detectors. LEGWORK differs from the tools listed here as we have tried to provide tools that are optimised to rapidly make predictions for large populations of stellar-origin sources and we have focused on space-based detectors.

LEGWORK has been developed with stellar-origin binary population studies in mind. We highlight two recent papers that use our package. Wagg et al. (2022c) investigates several populations of potential LISA sources (double black holes, black hole neutron stars and double neutron stars). They use LEGWORK to evolve each synthesised source over the age of the Milky Way, make predictions about the LISA detectable population and explore how it varies with different binary physics assumptions. Additionally, **?** examines the implications of assuming a metallicity-dependent binary fraction for the formation of close double white dwarfs (WDWDs) in the LISA frequency band. They use LEGWORK to calculate the SNR of WDWDs and, in particular, apply a custom noise model that combines the LISA sensitivity curve from Robson et al. (2019) with a new fit for the Galactic confusion noise based on their WDWD population.

LEGWORK can be installed with pip or obtained from the GitHub repository [1]. All examples shown in this paper and code to reproduce the figures are available in the repository. Instructions for installation and basic usage are provided in the online documentation[2] which contain the most up-to-date instructions.

This chapter is organised as follows. In Section 2 we give an overview of the capabilities of the LEGWORK package and its various modules. We detail the derivations of the equations relevant for LEGWORK in Section 3. In Section 4 we outline some example use cases of LEGWORK to demonstrate its use. Finally, in Section 5 we conclude and summarise our work.

2 Package Overview

The LEGWORK package is composed of seven modules that each focus on a particular aspect of calculations useful for gravitational wave sources that are detectable by space-based de-

---

[1] `https://github.com/TeamLEGWORK/LEGWORK`

[2] `https://legwork.readthedocs.io/en/latest/`

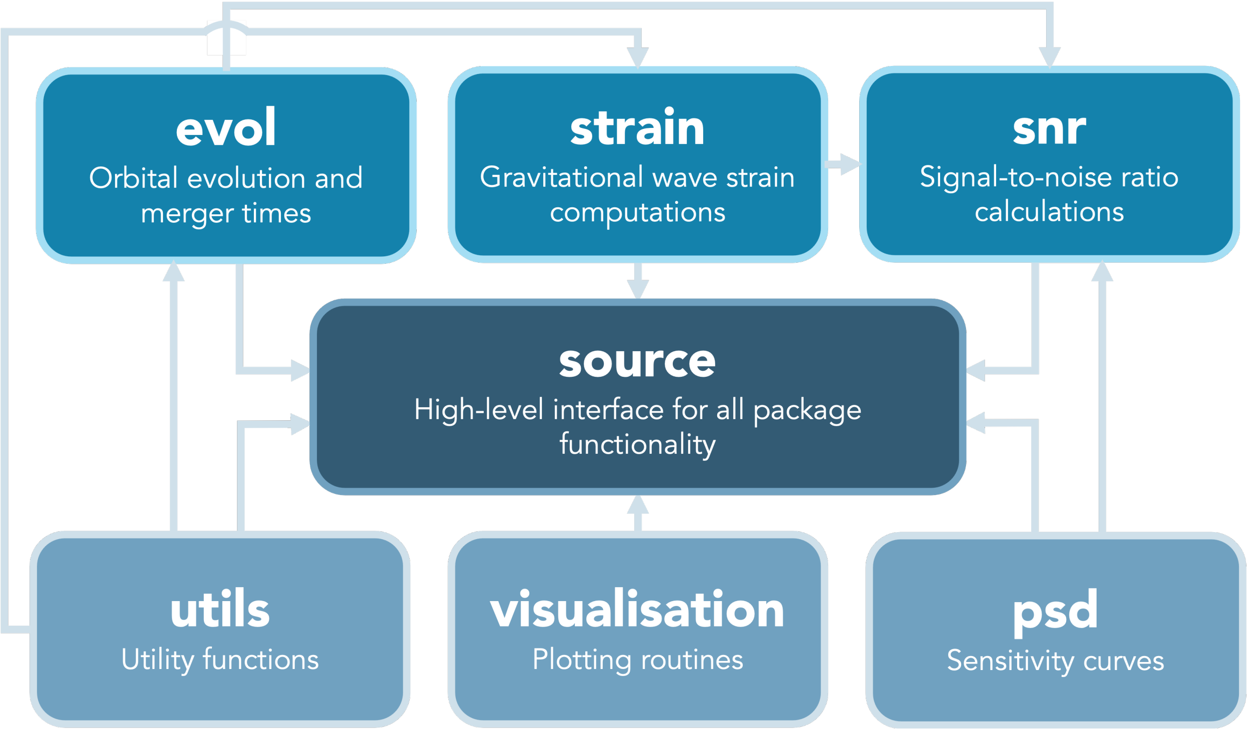

Figure 2.1: Package structure of `LEGWORK`. Each box represents a module and describes its function. The arrows indicate the inter-dependencies of the modules.

tectors. In Figure 2.1, we illustrate the general structure of the package with each of its modules. The `source` module is the central module of the package and provides an simple interface to the functions in the rest of the modules. For more complex analyses, users may want to interact directly with individual modules, particularly those in the top row of Figure 2.1 as they comprise the core functionality of `LEGWORK`. Below we explain the capabilities of each module in detail.

`evol` handles the orbital evolution of a binary due to the emission of gravitational waves. It includes functions for computing the merger times of both circular and eccentric binaries. In addition, you can use this module to evolve binary orbit parameters forward in time with any number or arrangement of timesteps. We discuss the relevant equations in Section 3.2.

`strain` contains two functions that compute a binary's gravitational wave strain and characteristic strain amplitude respectively. Each of these functions is capable of computing the strain for an array of binaries at any number of timesteps and evaluated at any number of frequency harmonics. We discuss the relevant equations in Section 3.3.

`psd` is used for evaluating the effective noise power spectral density of a detector at different frequencies. The module currently contains the LISA and TianQin sensitivity curves that can be tweaked by adjusting parameters such as the observation time, response function and even the arm length. For the Galactic confusion noise, a user can choose one of three models (Robson et al., 2019; Huang et al., 2020; **?**), use their own custom model, or turn it off entirely. Additionally, this module allows the user to specify a custom detector sensitivity curve. We discuss the relevant equations in Section 3.4.

`snr` uses the functions in `evol`, `strain` and `psd` to compute the signal-to-noise ratio of sources. It contains four functions that cover the permutations of whether a source is circular or eccentric and stationary in frequency space on the timescale of the mission or evolving. We discuss the relevant equations in Section 3.5.

`visualisation` contains several wrappers for plotting 1- and 2-dimensional distributions with histograms, scatter plots, and kernel density estimator (KDE) plots in order to quickly analyse a collection of sources. In addition, it provides functions for plotting sources directly onto a sensitivity curve.

`source` provides a direct and simple interface to the functions in other modules through the `Source` Class. You can instantiate this Class with an array of sources and use it compute their strains or signal-to-noise ratios directly. Moreover, depending on the user's choice of allowed gravitational wave luminosity error, the Class dynamically decides on the number of frequency harmonics needed to capture the full signal of each binary and at what eccentricity to no longer consider a binary circular. This Class also provides a quick means of evolving the sources, visualising the parameters of each source and allows you to plot the binaries on the sensitivity curve.

`utils` is a collection of miscellaneous utility functions mainly consisting of conversions between variables as well as constants and expressions from Peters (1964). We discuss the relevant equations in Section 3.1.

2.1 Units and automated testing

To ensure stability with physical units, all quantities included in `LEGWORK` use the `astropy.units` module (Astropy Collaboration et al., 2013, 2018). This means that all inputs to `LEGWORK` can be given in the units of the user's choice and will be automatically converted.

Furthermore, all of the source code in `LEGWORK` is fully tested with continuous integration in the `LEGWORK` GitHub repository. We employ several unit tests to ensure consistency between each of the use cases described below. For example, we require that the SNR calculation for circular and stationary binaries produces consistent output whether `LEGWORK` uses the stationary and circular approximations or not. Similarly, we verify that the antenna patterns described below produce the expected values when averaged over source positions, inclinations, and polarisations.

2.2 Optimistations

We developed `LEGWORK` with an emphasis on increasing the efficiency of these computations in order to make simulations of large populations of systems tractable. We ensured that the entirety of `LEGWORK` is vectorised and thus scales well with larger populations. In addition, we made a several specific optimisations to further increase the speed of calculations.

Firstly, we find that the runtime of calculating strains and SNRs for large populations of sources is mainly limited by the computation of (1) the relative gravitational-wave power in each harmonic for eccentric systems (see Eq. 2.5) (2) the sensitivity curve of the given detector. Therefore, in order to significantly reduce the runtime of strain and SNR computations, by default `LEGWORK` automatically interpolates these functions upon instantiation of any `Source` class with a large number of sources. Thereafter all functions use the tabulated values instead of calculating them exactly.

In certain cases one can apply approximations in place of the general SNR calculations. Although it is possible to use each of these approximations directly through the `snr` module, the `source` module will automatically apply the most appropriate function for each individual source. `LEGWORK` dynamically classifies each source as one of four types, which cover the permutations of whether a binary is effectively circular or eccentric and whether or not it is stationary in frequency space on the timescale of the LISA mission. Then when computing the SNR, it applies a different function to each type of the source. This avoids computing unnecessary, time intensive integrals.

Moreover, for increasingly eccentric sources, gravitational waves are emitted in an increasing number of higher frequency harmonics. Although the total SNR is formally calculated as a the sum of the SNR over an infinite number of harmonics, in practice it is sufficient to only consider a subset. In the `source` module, the user can provide `gw_lum_tol`, a maximum allowable tolerance for the accuracy of the gravitational wave luminosity. Given this tolerance, `LEGWORK` automatically calculates the required number of harmonics to satisfy this tolerance and thus minimise the computation time.

3 Derivations

In this section, we present a derivation of the equations used in `LEGWORK`. We emphasise that these are not new derivations, conversely, they are in fact given frequently in the literature. However, they are often incomplete, unclear and, in some cases, contain spurious constant factors that arise from invalid combinations of previous work. Here, we aim to present a clear, clean and concise explanation of the expressions we use in `LEGWORK`.

A </> symbol in an equation directly links to the relevant online `LEGWORK` documentation for the implementation of that equation, which additionally contains a link to exact code used to reproduce the equation. These derivations are also given in more detail in the `LEGWORK` documentation, where we show each of the intervening steps in more detail.

3.1 Conversions and definitions (`utils`)

We start these derivations by defining some useful conversions and definitions. The chirp mass of a binary is the mass quantity measured by LISA and is given by

$$\mathcal{M}_c \equiv \frac{(m_1 m_2)^{3/5}}{(m_1 + m_2)^{1/5}}, \tag*{</> (2.1)}$$

where m_1 and m_2 are the primary and secondary masses of the binary.

It is often convenient to convert between orbital frequency, $f_{\rm orb}$, and the semi-major axis, a, of a binary and this can be accomplished with Kepler's third law.

$$a = \left(\frac{G(m_1 + m_2)}{(2\pi f_{\rm orb})^2} \right)^{1/3}, \tag*{</> (2.2)}$$

where G is the gravitational constant. Inversely,

$$f_{\rm orb} = \frac{1}{2\pi}\sqrt{\frac{G(m_1+m_2)}{a^3}}. \tag{2.3}$$

For circular binaries, gravitational-wave emission occurs at twice the orbital frequency ($f_{\rm GW} = 2 f_{\rm orb}$). However, for eccentric binaries, we need to consider all frequency harmonics of gravitational-wave emission. These are defined such that the $n^{\rm th}$ harmonic frequency is

$$f_n \equiv n \cdot f_{\rm orb}. \tag{2.4}$$

It will be important to know the relative gravitational-wave power radiated into the $n^{\rm th}$ harmonic for a binary with eccentricity e for the strain and SNR calculations. This is given by Peters & Mathews (1963, Eq. 20)

$$\begin{aligned} g(n,e) = \frac{n^4}{32}\Bigg\{&\Big[J_{n-2}(ne) - 2eJ_{n-1}(ne) + \frac{2}{n}J_n(ne) + 2eJ_{n+1}(ne) - J_{n+2}(ne)\Big]^2 \\ &+ (1-e^2)\Big[J_{n-2}(ne) - 2J_n(ne) + J_{n+2}(ne)\Big]^2 + \frac{4}{3n^2}\Big[J_n(ne)\Big]^2\Bigg\}, \end{aligned} \tag{2.5}$$

where $J_n(v)$ is the Bessel function of the first kind. Thus, the *sum* of $g(n,e)$ over all harmonics gives the factor by which the gravitational-wave emission is stronger for a binary of eccentricity e over an equivalent circular binary. This enhancement factor is (Peters & Mathews, 1963, Eq. 17)

$$F(e) \equiv \sum_{n=1}^{\infty} g(n,e) = \frac{1 + \frac{73}{24}e^2 + \frac{37}{96}e^4}{(1-e^2)^{7/2}}. \tag{2.6}$$

A useful rule of thumb is that $F(0.5) \approx 5.0$, or in words, a binary with eccentricity 0.5 loses energy to gravitational waves at 5 times the rate of an equivalent circular binary.

3.2 Orbital evolution (`evol`)

3.2.1 Circular binaries

For a circular binary, the orbital evolution due to gravitational-wave emission can be calculated analytically, as the rate at which the separation of the binary shrinks is simply a function of its mass and the current separation (Peters, 1964, Eq. 5.6)

$$\frac{\mathrm{d}a}{\mathrm{d}t}_{e=0} = -\frac{\beta}{a^3}, \tag{2.7}$$

where the constant β is defined as

$$\beta(m_1, m_2) \equiv \frac{64}{5}\frac{G^3}{c^5} m_1 m_2 (m_1 + m_2), \tag{2.8}$$

where c is the speed of light, m_1 is the primary mass and m_2 is the secondary mass. This gives the semi-major axis of a circular binary as a function of time, t, as (Peters, 1964, Eq. 5.9)

$$a(t, m_1, m_2) = [a_0^4 - 4t\beta(m_1, m_2)]^{1/4}, \tag{2.9}$$

where a_0 is the initial semi-major axis. Moreover, we can solve for the merger time, the time until the binary will merge, by setting the final semi-major axis in Eq. 2.9 to zero.

$$t_{\rm merge,circ} = \frac{a_0^4}{4\beta}. \tag{2.10}$$

3.2.2 Eccentric binaries

The orbital evolution is more complex for eccentric binaries since the semi-major axis and eccentricity both evolve simultaneously and depend on one another. The final expression cannot be solved analytically and require numerical integration. The semi major axis, a, and eccentricity, e, are related as (Peters, 1964, Eq. 5.11)

$$a(e) = c_0 \frac{e^{12/19}}{(1 - e^2)} \left(1 + \frac{121}{304} e^2\right)^{870/2299}, \tag{2.11}$$

where c_0 satisfies the initial conditions such that $a(e_0) = a_0$. The time derivative of the eccentricity, e, is

$$\frac{\mathrm{d}e}{\mathrm{d}t} = -\frac{19}{12}\frac{\beta}{c_0^4} \frac{e^{-29/19}\,(1 - e^2)^{3/2}}{[1 + (121/304)e^2]^{\frac{1181}{2299}}}, \tag{2.12}$$

which we can integrate to find $e(t)$ and convert to $a(t)$ using Eq. 2.11 ().

Inverting this function and applying the fact that we know that $e \to 0$ when the binary merges gives the merger time (Peters, 1964, Eq. 5.14)

$$t_{\rm merge} = \frac{12}{19}\frac{c_0^4}{\beta} \int_0^{e_0} \frac{[1 + (121/304)e^2]^{\frac{1181}{2299}}}{e^{-29/19}\,(1 - e^2)^{3/2}} \mathrm{d}e. \tag{2.13}$$

For very small or very large eccentricities we approximate this integral using the following expressions (given in unlabelled equations after Peters, 1964, Eq. 5.14)

$$t_{\rm merge,\, e^2 \ll 1} = \frac{c_0^4}{4\beta} \cdot e_0^{48/19}, \tag{2.14}$$

$$t_{\rm merge,\, (1-e^2) \ll 1} = \frac{768}{425}\frac{a_0^4}{4\beta}(1 - e_0^2)^{7/2}. \tag{2.15}$$

The standard threshold employed by `LEGWORK` for small eccentricities is $e = 0.15$ and for large eccentricities is $e = 0.9999$ (as this approximates $t_{\rm merge}$ with an error below roughly 2%), though we note that this can be customised by the user if desired.

In addition, we implement the fit to Eq. 2.13 from Mandel (2021) that approximates the merger time as

$$t_{\rm merge} \approx t_{\rm merge,circ}(1 - e_0^2)^{7/2} \times (1 + 0.27e_0^{10} + 0.33e_0^{20} + 0.2e_0^{1000}), \tag{2.16}$$

which gives $t_{\rm merge}$ with an error below 3% for eccentricities below 0.9999. We additionally add a rudimentary polynomial fit to further reduce this error to below 0.5%. The user may specify whether to use this fit or perform the full integral when calculating merger times in `LEGWORK`.

3.3 Strains (`strain`)

3.3.1 Characteristic Strain

The strength of a gravitational wave in a detector at any one moment is determined by the strain amplitude, h_0. However, for stellar-origin sources at mHz frequencies, the signal can be present in the detector for many years. This means that, the $n^{\rm th}$ harmonic of the binary will spend approximately $f_n/\dot{f}_n$ seconds (or $f_n^2/\dot{f}_n$ cycles) in the vicinity of a frequency f_n (Finn & Thorne, 2000). This leads to the signal 'accumulating' at the frequency f_n.

Therefore, to account for the integration of the signal over the mission, we instead use the 'characteristic' strain amplitude of the $n^{\rm th}$ harmonic, $h_{c,n}$, which is the term present in the general signal-to-noise ratio equation. This can be related the to the strain amplitude in the $n^{\rm th}$ harmonic, $h_{0,n}$, as (e.g Finn & Thorne, 2000; Moore et al., 2015)[3]

$$h_{c,n}^2 = \left(\frac{f_n^2}{\dot{f}_n}\right) h_{0,n}^2, \tag{2.17}$$

The characteristic strain represents the strain measured by the detector over the duration of the mission (approximated as a single broad-band burst), whilst the strain amplitude is the strength of the GW emission at each instantaneous moment. For a stellar mass binary, the characteristic strain in the $n^{\rm th}$ harmonic is given by (e.g. Barack & Cutler, 2004, Eq. 56; Flanagan & Hughes, 1998, Eq. 5.1)

$$h_{c,n}^2 = \frac{1}{(\pi D_L)^2}\left(\frac{2G}{c^3}\frac{\dot{E}_n}{\dot{f}_n}\right), \tag{2.18}$$

where D_L is the luminosity distance to the source (note that for Milky Way sources, or any sources with redshift ~ 0, this is simply the distance to the source), $\dot{E}_n$ is the power radiated in the $n^{\rm th}$ harmonic and $\dot{f}_n$ is the rate of change of the $n^{\rm th}$ harmonic frequency.

The power radiated in the $n^{\rm th}$ harmonic can be expressed as (Peters & Mathews, 1963,

---

[3] Note that this is factor of 2 different from Finn & Thorne (2000). This is because the factor of 2 is already included in the Robson et al. (2019) sensitivity curve and so is removed here.

Eq. 19)

$$\dot{E}_n = \frac{32}{5} \frac{G^4}{c^5} \frac{m_1^2 m_2^2 \left(m_1 + m_2\right)}{a^5} g(n, e), \tag{2.19}$$

where $g(n, e)$ is given in Eq. 2.5. By substituting a for $f_{\rm orb}$ (using Eq. 2.3) and applying the definition of the chirp mass (Eq. 2.1) we obtain a more useful form for making gravitational wave predictions

$$\dot{E}_n(\mathcal{M}_c, f_{\rm orb}, e) = \frac{32}{5} \frac{G^{7/3}}{c^5} \left(2\pi f_{\rm orb} \mathcal{M}_c\right)^{10/3} g(n, e). \tag{2.20}$$

The last term needed to define the characteristic strain in Eq. 2.18 is the rate of change of the $n^{\rm th}$ harmonic frequency as a result of gravitational wave inspiral, which we can write as

$$\dot{f}_n = \frac{{\rm d} f_n}{{\rm d} a} \frac{{\rm d} a}{{\rm d} t}. \tag{2.21}$$

We can find an expression for ${\rm d} f_n \,/\, {\rm d} a$ by substituting Eq. 2.3 into Eq. 2.4 and differentiating

$$\frac{{\rm d} f_n}{{\rm d} a} = -\frac{3n}{4\pi} \frac{\sqrt{G(m_1 + m_2)}}{a^{5/2}}. \tag{2.22}$$

The rate at which the semi-major axis decreases is (Peters, 1964, Eq. 5.6)

$$\frac{{\rm d} a}{{\rm d} t} = -\frac{64}{5} \frac{G^3 m_1 m_2 \left(m_1 + m_2\right)}{c^5 a^3} F(e). \tag{2.23}$$

Substituting Eq. 2.22 and Eq. 2.23 into Eq. 2.21 gives an expression for $\dot{f}_n$

$$\dot{f}_n = \frac{48n}{5\pi} \frac{G^{7/2}}{c^5} \left(m_1 m_2 (m_1 + m_2)^{3/2}\right) \frac{F(e)}{a^{11/2}}, \tag{2.24}$$

which, as above with $\dot{E}_n$, we can recast using Kepler's third law and the definition of the chirp mass

$$\dot{f}_n = \frac{48n}{5\pi} \frac{\left(G \mathcal{M}_c\right)^{5/3}}{c^5} \left(2\pi f_{\rm orb}\right)^{11/3} F(e). \tag{2.25}$$

With definitions of both $\dot{E}_n$ and $\dot{f}_n$, we are now in a position to find an expression for the characteristic strain by plugging Eq. 2.20 and Eq. 2.25 into Eq. 2.18:

$$h_{c,n}^2 = \frac{2^{5/3}}{3\pi^{4/3}} \frac{\left(G \mathcal{M}_c\right)^{5/3}}{c^3 D_L^2} \frac{1}{f_{\rm orb}^{1/3}} \frac{g(n, e)}{n F(e)}. \tag{2.26}$$

3.3.2 Strain

In order to obtain an expression for the strain amplitude of gravitational waves in the n^{th} harmonic, we can use Eq. 2.17 and plug in Eq. 2.25 and Eq. 2.26

$$h_n^2 = \frac{2^{28/3}}{5} \frac{(G\mathcal{M}_c)^{10/3}}{c^8 D_L^2} \frac{g(n,e)}{n^2} (\pi f_{\rm orb})^{4/3}. \tag{2.27}$$

3.3.3 Amplitude modulation for orbit averaged sources

Because the LISA detectors are not stationary and instead follow an Earth-trailing orbit, the antenna pattern of LISA is not isotropically distributed or stationary. For sources that have unknown positions, inclinations, and polarisations, we use an average for the detector. However, for sources where these quantities are known we can consider the amplitude modulation of the strain due to the average motion of LISA's orbit.

We write that the position of the source on the sky is given by the ecliptic coordinates (θ, ϕ), the inclination of a source is ι and the polarisation of a source (determined by its orientation relative to the detector) is given by ψ. We follow the results of Cornish & Larson (2003) to define the amplitude modulation. However, we adapt their expression to remain in the frequency domain and follow the conventions of more recent papers (e.g. Babak et al., 2021, Eq. 67) to write the amplitude modulation as

$$A_{\rm mod}^2 = \frac{1}{4}\left(1 + \cos^2\iota\right)^2 \left\langle F_+^2 \right\rangle_{\rm orb} + \cos^2\iota \left\langle F_\times^2 \right\rangle_{\rm orb}, \tag{2.28}$$

where $\left\langle F_+^2 \right\rangle_{\rm orb}$ and $\left\langle F_\times^2 \right\rangle_{\rm orb}$, the orbit-averaged detector responses, are defined as

$$\left\langle F_+^2 \right\rangle_{\rm orb} = \frac{1}{4}\big(\cos^2 2\psi \left\langle D_+^2 \right\rangle_{\rm orb} - \sin 4\psi \left\langle D_+ D_\times \right\rangle_{\rm orb} + \sin^2 2\psi \left\langle D_\times^2 \right\rangle_{\rm orb}\big), \tag{2.29}$$

$$\left\langle F_\times^2 \right\rangle = \frac{1}{4}\big(\cos^2 2\psi \left\langle D_\times^2 \right\rangle_{\rm orb} + \sin 4\psi \left\langle D_+ D_\times \right\rangle_{\rm orb} + \sin^2 2\psi \left\langle D_+^2 \right\rangle_{\rm orb}\big), \tag{2.30}$$

and

$$\left\langle D_+ D_\times \right\rangle_{\rm orb} = \frac{243}{512} \cos\theta \sin 2\phi \left(2\cos^2\phi - 1\right) \times \left(1 + \cos^2\theta\right), \tag{2.31}$$

$$\left\langle D_\times^2 \right\rangle_{\rm orb} = \frac{3}{512}\big(120\sin^2\theta + \cos^2\theta + 162\sin^2 2\phi \cos^2\theta\big), \tag{2.32}$$

$$\left\langle D_+^2 \right\rangle_{\rm orb} = \frac{3}{2048}\left[487 + 158\cos^2\theta + 7\cos^4\theta - 162\sin^2 2\phi \left(1 + \cos^2\theta\right)^2\right]. \tag{2.33}$$

The orbital motion of LISA smears the source frequency by roughly 10^{-4} mHz due to the antenna pattern changing as the detector orbits, the Doppler shift from the motion, and the phase modulation from the + and × polarisations in the antenna pattern. Generally, the modulation reduces the strain amplitude because the smearing in frequency reduces the amount of signal build up at the true source frequency.

We note that the amplitude modulation is only implemented in `LEGWORK` for quasi-circular binaries to remain consistent with the calculation in Cornish & Larson (2003). Since the expected use case of LEGWORK is estimation of the detectability of large populations of mHz stellar-remnant binaries, for which predictions are uncertain by orders of magnitude in some cases, an extension of Cornish & Larson (2003) which includes eccentric binaries is out of the current scope of `LEGWORK`.

3.4 Sensitivity Curves (`psd`)

3.4.1 LISA

For the LISA sensitivity curve we use the equations from Robson et al. (2019). The *effective* strain spectral density of the noise is defined as

$$S_{\rm n}(f) \equiv \frac{P_n(f)}{\mathcal{R}(f)} + S_c(f), \tag{2.34}$$

where $P_{\rm n}(f)$ is the power spectral density of the detector noise and $\mathcal{R}(f)$ is the sky and polarisation averaged signal response function of the instrument. Alternatively if we expand out $P_n(f)$, approximate $\mathcal{R}(f)$ and simplify we find (Robson et al., 2019, Eq. 1)

$$\begin{aligned} S_{\rm n}(f) = \frac{10}{3L^2}&\left(P_{\rm OMS}(f) + 2\left(1 + \cos^2\left(\frac{f}{f_*}\right)\right)\frac{P_{\rm acc}(f)}{(2\pi f)^4}\right) \\ &\times\left(1 + \frac{6}{10}\left(\frac{f}{f_*}\right)^2\right) + S_c(f), \end{aligned} \tag{2.35}$$

where $L = 2.5\,{\rm Gm}$ is detector arm length, $f_* = c/2\pi L = 19.09\,{\rm mHz}$ is the transfer frequency,

$$\begin{aligned} P_{\rm OMS}(f) = &\left(1.5 \times 10^{-11}{\rm m}\right)^2 \\ &\times\left(1 + \left(\frac{2{\rm mHz}}{f}\right)^4\right){\rm Hz}^{-1}, \end{aligned} \tag{2.36}$$

is the single-link optical metrology noise (Robson et al., 2019, Eq. 10),

$$\begin{aligned} P_{\rm acc}(f) = &\left(3 \times 10^{-15}{\rm ms}^{-2}\right)^2\left(1 + \left[\frac{0.4{\rm mHz}}{f}\right]^2\right) \\ &\times\left(1 + \left[\frac{f}{8{\rm mHz}}\right]^4\right){\rm Hz}^{-1}, \end{aligned} \tag{2.37}$$

is the single test mass acceleration noise (Robson et al., 2019, Eq. 11) and

$$S_c(f) = A f^{-7/3} e^{-f^{\alpha} + \beta f \sin(\kappa f)} \times \left[1 + \tanh\left(\gamma\left(f_k - f\right)\right)\right] \mathrm{Hz}^{-1}, \tag{2.38}$$

is the galactic confusion noise (Robson et al., 2019, Eq. 14), where the amplitude A is fixed as 9×10^{-45} and the various parameters change over time and are listed in Robson et al. 2019, Table 1. `LEGWORK` allows the user to opt to use the Robson et al. (2019) confusion noise, a custom function for the confusion noise or to remove the confusion noise entirely.

3.4.2 TianQin

For the TianQin sensitivity curve we use the power spectral density given in Huang et al. (2020, Eq. 13)

$$S_N(f) = \frac{10}{3L^2}\left[\frac{4S_a}{(2\pi f)^4}\left(1 + \frac{10^{-4}\,\mathrm{Hz}}{f}\right) + S_x\right] \times \left[1 + 0.6\left(\frac{f}{f_*}\right)^2\right], \tag{2.39}$$

where $L = \sqrt{3} \times 10^5\,\mathrm{km}$ is the arm length, $S_a = 1 \times 10^{-30}\,\mathrm{m}^2\,\mathrm{s}^{-4}\,\mathrm{Hz}^{-1}$ is the acceleration noise, $S_x = 1 \times 10^{-24}\,\mathrm{m}^2\,\mathrm{Hz}^{-1}$ is the displacement measurement noise and $f_* = c/2\pi L$ is the transfer frequency. Note that Eq. 2.39 includes an extra factor of 10/3 compared to Huang et al. (2020, Eq. 13). Huang et al. (2020) absorb this factor into the waveform rather than include it in the power spectral density. We include it to match the convention used by Robson et al. (2019) for the LISA sensitivity curve (see the factor of 10/3 in Eq. 2.35) so that the sensitivity curves can be compared fairly.

3.5 SNR for LISA (`snr`)

We note that this section draws heavily from Flanagan & Hughes (1998) Section II C.

3.5.1 Defining general SNR

In order to calculate the signal to noise ratio for a given source of gravitational waves (GWs) in a 6-link LISA detector, we need to consider the following parameters:

- position of the source on the sky: (θ, ϕ)
- direction from the source to the detector: (ι, β)
- orientation of the source, which fixes the polarisation of the GW: ψ
- the distance from the source to the detector: D_L

Then, assuming a matched filter analysis of the GW signal $s(t) + n(t)$ (where $s(t)$ is the signal and $n(t)$ is the noise), which relies on knowing the shape of the signal, the signal to noise ratio, ρ, is given in the frequency domain as

$$\rho^2(D_L, \theta, \phi, \psi, \iota, \beta) = \frac{\langle s(t)^\star s(t) \rangle}{\langle n(t)^\star n(t) \rangle}, \tag{2.40}$$

$$= 2 \int_{-\infty}^{+\infty} \frac{|\tilde{s}(f)|^2}{P_{\rm n}(f)} df, \tag{2.41}$$

$$= 4 \int_{0}^{\infty} \frac{|\tilde{s}(f)|^2}{P_{\rm n}(f)} df, \tag{2.42}$$

where $\tilde{s}(f)$ is the Fourier transform of the signal, $s(t)$, and $P_{\rm n}(f)$ is the one-sided power spectral density of the noise defined as as $\langle n(t)^\star n(t) \rangle = \int_0^\infty \frac{1}{2} P_{\rm n}(f) df$ (c.f. Robson et al., 2019, Eq. 2). Here, $\tilde{s}(f)$ is implicitly also dependent on $D_L, \theta, \phi, \psi, \iota$, and β as

$$|\tilde{s}(f)|^2 = |F_+(\theta, \phi, \psi)\tilde{h}_+(f, D_L, \iota, \beta) + F_\times(\theta, \phi, \psi)\tilde{h}_\times(f, D_L, \iota, \beta)|^2, \tag{2.43}$$

where $F_{+,\times}$ are the 'plus' and 'cross' antenna patterns of the LISA detector to the 'plus' and 'cross' strains, $h_{+,\times}$. Note throughout any parameters discussed with the subscript $x_{+,\times}$ refers to both x_+ and $x_\times$.

3.5.2 Average over position and polarisation

Now, we can consider averaging over different quantities. In LISA's case, when averaged over all angles and polarisations, the antenna patterns are orthogonal and thus $\langle F_+ F_\times \rangle = 0$. This means we can rewrite Eq. 2.43 as

$$|\tilde{s}(f)|^2 = |F_+(\theta, \phi, \psi)\tilde{h}_+(f, D_L, \iota, \beta)|^2 + |F_\times(\theta, \phi, \psi)\tilde{h}_\times(f, D_L, \iota, \beta)|^2, \tag{2.44}$$

which can then be applied to Eq. 2.40 to give

$$\langle \rho \rangle^2_{\theta,\phi,\psi} = 4 \int_0^\infty df \int \frac{d\Omega_{\theta,\phi}}{4\pi} \int \frac{d\psi}{\pi} \frac{|F_+ \tilde{h}_+|^2 + |F_\times \tilde{h}_\times|^2}{P_{\rm n}(f)}. \tag{2.45}$$

From Robson et al. (2019), we can write the position and polarisation average of the signal response function of the instrument, $\mathcal{R}$, as

$$\begin{aligned} \mathcal{R} &= \langle F_+ F_+^\star \rangle = \langle F_\times F_\times^\star \rangle, \\ \text{where } \langle \mathrm{F}_{+,\times} \mathrm{F}_{+,\times}^\star \rangle &= \int \frac{d\Omega_{\theta,\phi}}{4\pi} \int \frac{d\psi}{\pi} |F_{+,\times}|^2. \end{aligned} \tag{2.46}$$

Then combining Eq. 2.45 and Eq. 2.46, we then find[4]

---

[4] Note that this is written in Flanagan & Hughes (1998) for the LIGO response function which is $\mathcal{R} = \langle F_{+,\times} \rangle^2 = 1/5$.

$$\langle\rho\rangle^2_{\theta,\phi,\psi} = 4\int_0^\infty df\mathcal{R}(f)\left(\frac{|\tilde{h}_+|^2 + |\tilde{h}_\times|^2}{P_{\rm n}(f)}\right). \tag{2.47}$$

3.5.3 Average over orientation

Now, we can average over the orientation of the source: (ι, β), noting that the averaging is independent of the distance D_L. With this in mind, we can rewrite $|\tilde{h}_+|^2 + |\tilde{h}_\times|^2$ in terms of two functions $|\tilde{H}_+|^2$ and $|\tilde{H}_\times|^2$, where $\tilde{h}_{+,\times} = \tilde{H}_{+,\times}/D_L$. Given this, averaging over the source direction gives

$$\langle\rho\rangle^2_{(\theta,\phi,\psi),(\iota,\beta)} = \frac{4}{D_L^2}\int_0^\infty df\mathcal{R}(f)\int\frac{d\Omega_{\iota,\beta}}{4\pi}\frac{|\tilde{H}_+|^2 + |\tilde{H}_\times|^2}{P_{\rm n}(f)}, \tag{2.48}$$

where we would like to express $\tilde{H}_{+,\times}(f)^2$ in terms of the energy spectrum of the GW. To do this, we note that the local energy flux of GWs at the detector is given by (e.g. Press & Thorne, 1972, Eq. 6)

$$\frac{dE}{dAdt} = \frac{1}{16\pi}\frac{c^3}{G}\overline{\left[\left(\frac{dh_+}{dt}\right)^2 + \left(\frac{dh_\times}{dt}\right)^2\right]}, \tag{2.49}$$

where the bar indicates an average over several cycles of the wave which is appropriate for LISA sources. We can transform Eq. 2.49 using Parseval's theorem, where we can write

$$\int_{-\infty}^{+\infty} dt\int dA\frac{dE}{dAdt} = \int_0^\infty df\int dA\frac{c^3}{G}\frac{\pi f^2}{2}\left(|\tilde{h}_+|^2 + |\tilde{h}_\times|^2\right). \tag{2.50}$$

Note that the factor of frequency squared comes from the Fourier transform of the square of the time derivative in Eq. 2.50. Now, since $A = D_L^2\Omega$ and $|\tilde{h}_{+,\times}|^2 = |\tilde{H}_{+,\times}|^2/D_L^2$, we find

$$|\tilde{h}_{+,\times}|^2 dA = |\tilde{H}_{+,\times}|^2 d\Omega_{\iota,\beta}, \tag{2.51}$$

then we can write Eq. 2.50 in terms of $|H_{+,\times}|^2$ as

$$\int_{-\infty}^{+\infty} dt\int dA\frac{dE}{dAdt} = \int_0^\infty df\frac{\pi f^2 c^3}{2G}\int d\Omega\left(|\tilde{H}_+|^2 + |\tilde{H}_\times|^2\right). \tag{2.52}$$

Alternatively, by using Eq. 2.51 and performing a Fourier transform we can also write that

$$\int_{-\infty}^{+\infty} dt\int dA\frac{dE}{dAdt} = \int_0^\infty df\int d\Omega\frac{dE}{d\Omega df}. \tag{2.53}$$

From inspection of Eq. 2.52 and Eq. 2.53, we can write the spectral energy flux as

$$\int d\Omega \frac{dE}{d\Omega df} = \frac{\pi f^2 c^3}{2G} \int d\Omega \left(|\tilde{H}_+|^2 + |\tilde{H}_\times|^2 \right). \tag{2.54}$$

3.5.4 Fulled averaged SNR

We are now in a position to write an expression for the fully averaged SNR. Note for brevity we write $\langle \rho \rangle^2$ when referring to $\langle \rho \rangle^2_{(\theta,\phi,\psi),(\iota,\beta)}$. The application of Eq. 2.54 to Eq. 2.48 yields

$$\langle \rho \rangle^2 = \frac{4G}{c^3 D_L^2} \int_0^\infty df \frac{1}{P_{\rm n}(f)/\mathcal{R}(f)} \int \frac{d\Omega}{4\pi} \frac{dE}{d\Omega df} \frac{2}{\pi f^2}. \tag{2.55}$$

This simplifies nicely to

$$\langle \rho \rangle^2 = \frac{2G}{\pi^2 c^3 D_L^2} \int_0^\infty df \frac{dE}{df} \frac{1}{f^2 P_{\rm n}(f)/\mathcal{R}(f)}. \tag{2.56}$$

Finally, noting that $dE/df = dE/dt \times dt/df = \dot{E}/\dot{f}$, we can use the definition of the characteristic strain from Eq. 2.18 to finish up our position, direction, and orientation/polarisation averaged SNR as

$$\langle \rho \rangle^2 = \int_0^\infty df \frac{h_c^2}{f^2 P_{\rm n}(f)/\mathcal{R}(f)} = \int_0^\infty df \frac{h_c^2}{f^2 S_{\rm n}(f)}, \tag{2.57}$$

where we have used that the effective power spectral density of the noise is defined as $S_{\rm n}(f) = P_{\rm n}(f)/\mathcal{R}(f)$.

It is also important to note that this is only the SNR for a circular binary for which we need only consider the $n = 2$ harmonic. In the general case, a binary could be eccentric and requires a sum over *all* harmonics. Thus we can generalise Eq. 2.57 to eccentric binaries with

$$\langle \rho \rangle^2 = \sum_{n=1}^{\infty} \langle \rho_n \rangle^2 = \sum_{n=1}^{\infty} \int_0^\infty df_n \frac{h_{c,n}^2}{f_n^2 S_{\rm n}(f_n)}, \tag{2.58}$$

where $h_{c,n}$ is defined in Eq. 2.26 and $S_{\rm n}$ in Eq. 2.35.

3.5.5 SNR Approximations

Although Eq. 2.58 can be used for every binary, it can be useful to consider different cases in which we can avoid unnecessary sums and integrals. There are four possible cases for binaries in which we can use increasingly simple expressions for the signal-to-noise ratio. Binaries can be circular and stationary in frequency space.

Circular binaries emit only in the $n = 2$ harmonic and so the sum over harmonics can be removed. Stationary binaries have $f_{n,i} \approx f_{n,f}$ and so the small interval allows one to approximate the integral. Note we refer to non-stationary binaries as ‘evolving’ here though many papers also use ‘chirping’.

For an evolving and eccentric binary, no approximation can be made and the SNR is found using Eq. 2.58.

For a evolving and circular binary, the sum can be removed and so the SNR found as

$$\langle \rho \rangle^2_{\rm c,e} = \int_{f_{2,i}}^{f_{2,f}} \frac{h^2_{c,2}}{f_2^2 S_{\rm n}(f_2)} \, {\rm d}f_2 \, . \tag{2.59}$$

For a stationary and eccentric binary we can approximate the integral.

$$\begin{aligned} \langle \rho \rangle^2_{\rm e,s} &= \sum_{n=1}^{\infty} \lim_{\Delta f \to 0} \int_{f_n}^{f_n + \Delta f_n} \frac{h^2_{c,n}}{f_n^2 S_{\rm n}(f_n)} \, {\rm d}f_n \, , &(2.60) \\ &= \sum_{n=1}^{\infty} \frac{\Delta f_n \cdot h^2_{c,n}}{f_n^2 S_{\rm n}(f_n)}, &(2.61) \\ &= \sum_{n=1}^{\infty} \frac{\dot{f}_n \Delta T \cdot h^2_{c,n}}{f_n^2 S_{\rm n}(f_n)}, &(2.62) \\ &= \sum_{n=1}^{\infty} \left(\frac{\dot{f}_n}{f_n^2} h^2_{c,n} \right) \frac{T_{\rm obs}}{S_{\rm n}(f_n)}, &(2.63) \end{aligned}$$

where we have applied Eq. 2.17 to convert between strains. This gives the following expression

$$\langle \rho \rangle^2_{\rm e,s} = \sum_{n=1}^{\infty} \frac{h_n^2 T_{\rm obs}}{S_{\rm n}(f_n)}. \tag{2.64}$$

Finally, for a stationary and circular binary the signal-to-noise ratio is

$$\langle \rho \rangle^2_{\rm c,s} = \frac{h_2^2 T_{\rm obs}}{S_{\rm n}(f_2)}. \tag{2.65}$$

For SNR calculations that take into account the amplitude modulation due to LISA's orbital motion (Section 3.3.3, we apply the calculations as described above but include the modulation in either the strain or characteristic strain as is appropriate.

4 Use cases

In this section, we demonstrate `LEGWORK`'s range of capabilities through a series of example use cases. The plots and results in each subsection are reproduced directly in online demos in the `LEGWORK` documentation, which are each based on individual Jupyter notebooks. These tutorials are linked at the start of each subsection with the 💻 icon.

4.1 Computing the SNR of a binary system -

The most fundamental use case of `LEGWORK` is to compute the SNR of an individual binary system. This can be accomplished in `LEGWORK` with only two lines of code - one line to set up the source and another to compute the SNR. As an example one could consider a binary with the parameters

$$\begin{gathered} m_1 = m_2 = 10\,\mathrm{M}_\odot,\ d = 8\,\mathrm{kpc}, \\ f_{\rm orb} = 10^{-4}\,\mathrm{Hz},\ e = 0.2, \end{gathered}$$

where m_1, m_2 are the primary and secondary masses, d is the distance to the source, $f_{\rm orb}$ is the orbital frequency and e is the eccentricity. As we do not specify a position, polarisation or inclination, `LEGWORK` will calculate the SNR averaged over these quantities. If a user specifies the position, then if either of the polarisation or inclination are not specified then they will be randomly generated. `LEGWORK` does not allow users to specify polarisations or inclinations without positions as we presume that the user must know the position if they know the polarisation or inclination of the source. One can now instantiate a source in `LEGWORK` using these parameters (leaving the mission parameters as the default values, such that we compute the SNR for a 4-year LISA mission). As shown in the linked demo, `LEGWORK` quickly computes that the SNR of this binary is 4.49. In the background, `LEGWORK` decides which SNR approximation is most applicable given the eccentricity of the binary and whether it is stationary in frequency space.

This can be generalised to a population of many binary systems with ease. Instead of inputting single values for each parameter, one can input arrays of values where each entry corresponds to a different binary. As an example, we can take the same parameters as above for 3 different binaries but vary the primary as $m_1 = [5, 10, 15]\,\mathrm{M}_\odot$. Using `LEGWORK` we find that the SNR for each of these cases is $\rho = [2.47, 4.49, 7.85]$. This also need not be limited to a 4-year LISA mission. With `LEGWORK` we can additionally specify various parameters for the detector. For example, using `LEGWORK` we can find that, for a 5-year TianQin mission with no confusion noise, the SNR for each of these cases is $\rho = [1.07, 1.95, 3.41]$.

4.2 Horizon distance -

A common question to consider with stellar mass sources in LISA is how far away a certain source could be detected. In other words, what is the horizon distance beyond which a source no longer has a SNR greater than some chosen threshold. We can explore this question using `LEGWORK`.

Let us compute the horizon distance for a grid of chirp masses and orbital frequencies. First, we can recall that the signal-to-noise ratio of a source is inversely proportional to its distance from LISA and so we can find the horizon distance, $D_{\rm hor}$ as

$$D_{\rm hor} = \frac{\rho(D)}{\rho_{\rm detect}} \cdot D, \tag{2.66}$$

where $\rho(D)$ is the signal-to-noise ratio as some distance D and $\rho_{\rm detect}$ is the threshold above which we consider a source to be detectable. For the purpose of this example we will set

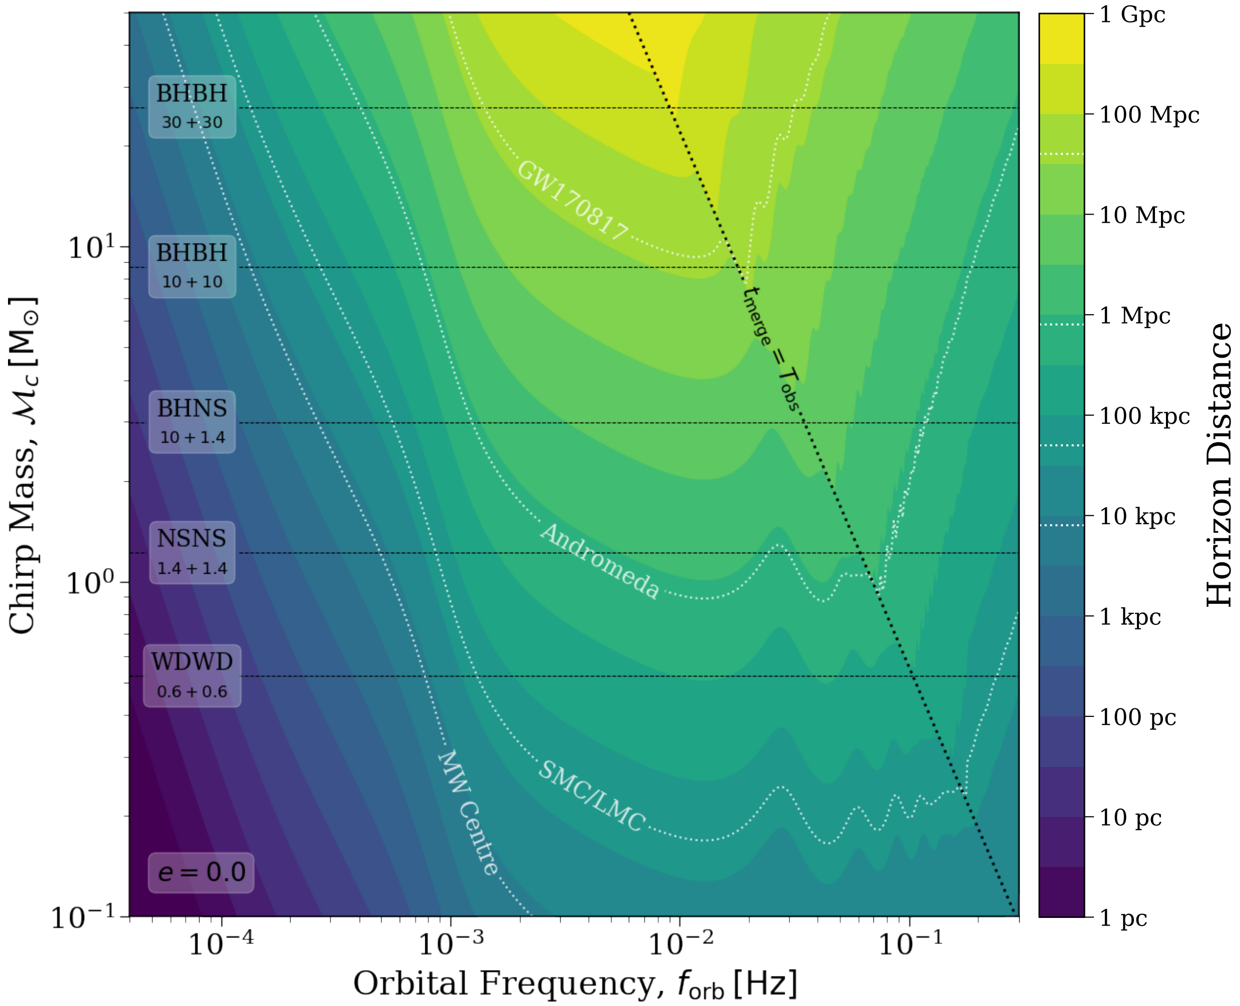

Figure 2.2: The horizon distance for circular, stellar-mass binaries in a 4-year LISA mission. The filled contours indicate the horizon distance for different orbital frequencies and chirp masses. We add white dotted contours at 8, 50, 800 kpc and 40 Mpc to highlight the distances to the centre of the Milky Way, the Magellanic Clouds, the Andromeda galaxy and the nearest ground-based gravitational wave detection (GW170817, Abbott et al., 2017a) respectively. The diagonal black line shows the frequencies and chirp masses at which the merger time is equal to the observation time. This line emphasises the sharp contrast for the horizon distance for binaries that merge before the LISA mission finishes observing (to the right of the line). The horizontal black lines indicate the approximate location of some common double compact object types on this plot, with the assumed masses labelled below in solar masses.

$\rho_{\rm detect} = 7$. We can then, as a function of chirp mass and frequency, determine the maximum distance for which the source of interest is detected based on this SNR threshold.

This can be most efficiently accomplished using `LEGWORK`'s `Source` class, since `LEGWORK` can then compute the merger times and signal-to-noise ratio of each source with only two lines of code. We convert the SNRs to horizon distances using Eq. 2.66 and plot the result in Figure 2.2. The shape of the LISA sensitivity curve is clearly reflected in the horizon distance. This is because circular sources only emit at a single frequency and thus every feature in the sensitivity curve has a strong effect on the SNR (and therefore horizon distance) for circular sources. We also see that once the merger time of a source is shorter than the LISA mission length (shown by the black dotted line), the horizon distance sharply decreases. This is because if a source merges before the LISA mission concludes, it has less time to accumulate signal and thus has a lower SNR and horizon distance.

One can use Figure 2.2 to estimate the horizon distance for any circular source of interest. To illustrate this, we add solid black lines to indicate the typical chirp masses of some possible stellar mass gravitational wave sources as well as white dotted lines to show the distances to nearby galaxies as well as the nearest ground-based gravitational wave detection (GW170817, Abbott et al., 2017a). For example, we see that a circular NSNS with an orbital frequency greater than a mHz is detectable in the Magellanic clouds.

4.3 The role of eccentricity -

The role of eccentricity is important to consider in the detection of gravitational waves with LISA and other space-based detectors, as sources can still have significant eccentricity during their inspiral phase. A high eccentricity has two major effects on the SNR of a gravitational wave source in LISA and we can investigate these effects using `LEGWORK`. Let's consider three hypothetical systems that are identical apart from their eccentricities, e_i.

$$\begin{aligned} m_1 = m_2 = 0.6\,\mathrm{M}_\odot, f_{\rm orb} = 1.5\,\mathrm{mHz}, d = 15\,\mathrm{kpc}, \\ e_i = \{0.0, 0.6, 0.9\}, \end{aligned}$$

Using `LEGWORK` to calculate their signal-to-noise ratios in a 4-year LISA mission, ρ_i, we find

$$\rho_i = \{31.7, 50.2, 38.8\}. \tag{2.67}$$

We see two effects in the signal-to-noise ratio here. First, increasing the eccentricity from essentially circular to $e = 0.6$ results in a higher signal-to-noise ratio ($\rho = 31.7 \to \rho = 50.2$). This is because an eccentric binary has enhanced energy emission via gravitational waves (Peters & Mathews, 1963). This means that an eccentric binary will not only inspiral faster than an otherwise identical circular binary, but also will always produce a stronger gravitational wave strain. We discuss the enhancement factor and its exact dependence on eccentricity in more detail in Section 3 (see specifically Eq. 2.6).

The second effect is more intriguing. We see that increasing the eccentricity from $e = 0.6$ to $e = 0.9$ results in a relative *decrease* in SNR ($\rho = 50.2 \to \rho = 38.8$). The reason for this is that eccentric binaries emit gravitational waves at many harmonic frequencies (unlike circular binaries, which emit predominantly twice the orbital frequency). This leads

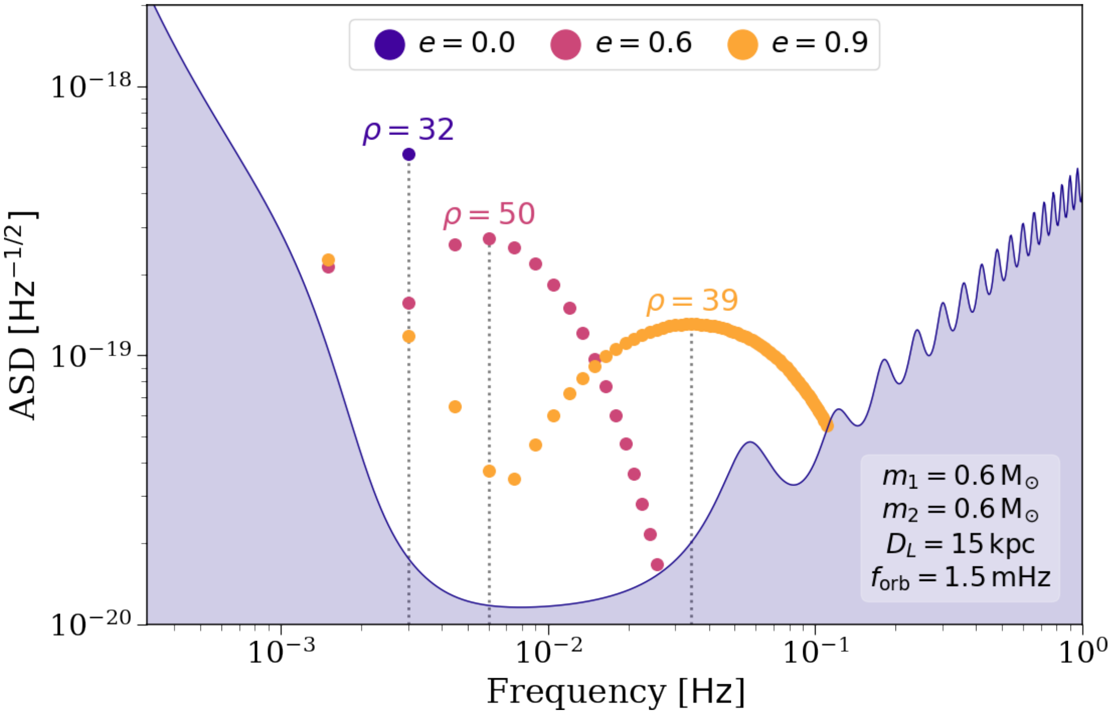

Figure 2.3: An illustration of the effect of eccentricity on the detectability of a LISA source. The three sets of points are coloured by their eccentricity and each individual point corresponds to a harmonic frequency, where its height above the curve gives its SNR. We annotate each set of points with its total SNR and overlay the LISA sensitivity curve. The dotted vertical lines indicate the frequency at which the majority of the gravitational wave signal is concentrated.

to the gravitational wave signal being diluted over many frequencies higher than the orbital frequency, where the higher the eccentricity, the more harmonics are required to capture all of the gravitational luminosity (see Figure 3 of Peters & Mathews, 1963). Therefore, if the eccentricity is too high, the majority of the signal may be emitted at a frequency to which LISA is less sensitive.

We can illustrate this point with `LEGWORK` by calculating the SNR at each individual frequency harmonic using the `snr` module. We plot this distribution of signal over different frequency harmonics with the LISA sensitivity curve overlaid in Figure 2.3. A point is plotted for each harmonic of each source that has an SNR greater than unity and such that its height above the sensitivity curve corresponds to its SNR.

From Figure 2.3, we can better understand why a source with $e = 0.9$ has a lower SNR than the same source with $e = 0.6$. From the dotted lines, we can note that the signal from the $e = 0.9$ source is concentrated at a frequency of around 40 mHz. The LISA sensitivity at this point is much weaker than the 6 mHz at which the $e = 0.6$ source is concentrated. Therefore, although the strain from a more eccentric binary is stronger, the SNR is lower due to the increased noise in the LISA detector.

Overall, we can therefore conclude that for LISA sources of this nature, higher eccentricity will produce more detectable binaries only if the orbital frequency is not already at or above

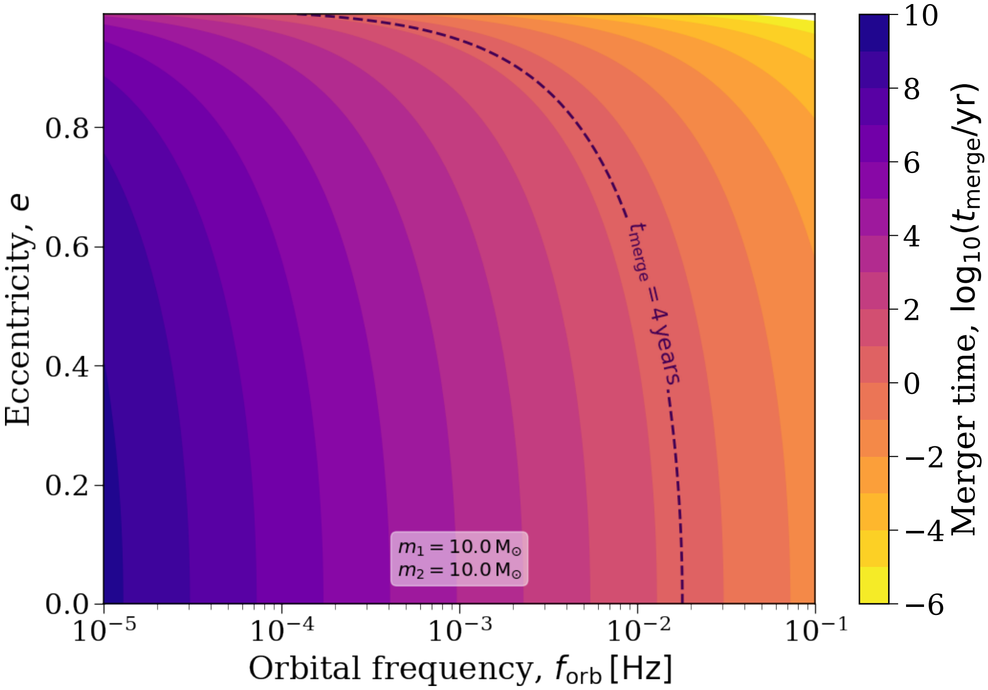

Figure 2.4: The merger time for a binary with both a primary and secondary mass of $10\,{\rm M}_\odot$ over different eccentricities and orbital frequencies. The dashed line indicates a merger time of 4 years, the default LISA mission length.

the minimum of the LISA sensitivity curve.

Another consideration for more massive binaries is whether the increased eccentricity will cause the binary to merge before the mission ends, which would cause a significant decrease in signal-to-noise ratio. We can also use `LEGWORK` to find how the merger time of a source varies with frequency and eccentricity over a grid of sources.

We plot the results of this calculation in Figure 2.4. This plot shows that, for most eccentricities, the merger time is largely determined by the orbital frequency. However, for high eccentricities ($e > 0.8$), the eccentricity leads to a significant reduction in the merger time. Additionally, we can see that any binary that is to the right of the dashed line on this plot at the start of a 4-year LISA mission, would merge before the mission ended. Therefore, if increasing a binary's eccentricity moved it to the right of this line, its SNR will decrease significantly.

4.4 Comparing gravitational wave detectors -

It may also be useful to consider how changing the specifications of the LISA detector, or using a different detector entirely, could affect the SNR of a particular source. `LEGWORK` is capable of adjusting the LISA mission specifications or using a different sensitivity curve and thus we can use it to explore these differences.

As a first step, we can use `LEGWORK` to plot a series of sensitivity curves in Figure 2.5. We show the LISA sensitivity curve for the default 4 year mission length but also illustrate how

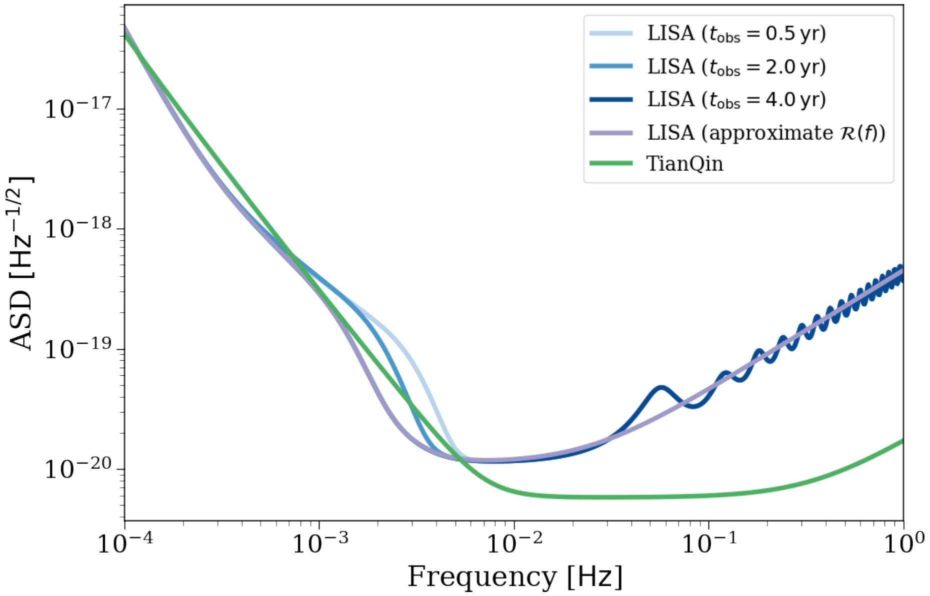

Figure 2.5: The strain spectral density of the LISA detector with different specifications (Robson et al., 2019) and the TianQin detector (Huang et al., 2020). We show the LISA curve for three different mission lengths and once with an approximate response function.

the curve changes for shorter mission lengths. At 0.5 and 2 years, we see a stronger noise level around 3 mHz as a result of the increased Galactic confusion noise. This noise decreases with increasing mission length since more individual foreground sources can be resolved and thus removed from the confusion noise. We also see that using an approximated response function smooths out the sensitivity curve at higher frequencies. Finally, the TianQin curve is higher than the 4 year LISA curve until around 5 mHz, beyond which it has a lower noise level. Although comparing the sensitivity curves would suffice for a stationary and circular source (since it would remain at a single frequency), `LEGWORK` can also be used to see how the relative SNR between two detectors changes over a range of eccentricities and frequencies.

Using `LEGWORK`, we can compute the SNR of a grid of sources (spanning a range of frequencies and eccentricities) for both detectors. In Figure 2.6, we show the ratio of the SNR in LISA to the SNR in TianQin. This plot shows that for circular binaries, the SNR of the source in LISA is stronger up to an orbital frequency of approximately 2.5 mHz, beyond which the SNR of the source is stronger in TianQin. This transition frequency becomes lower with increasing eccentricity as one would expect since eccentric sources emit more at higher harmonics and thus higher frequencies.

4.5 Track SNR of a binary over time -

As a binary inspirals its orbital frequency and eccentricity change and this in turn affects the SNR of the binary. For this use case we will demonstrate how `LEGWORK` can be used to track the evolution of these parameters and pinpoint the moment at which a binary becomes detectable.

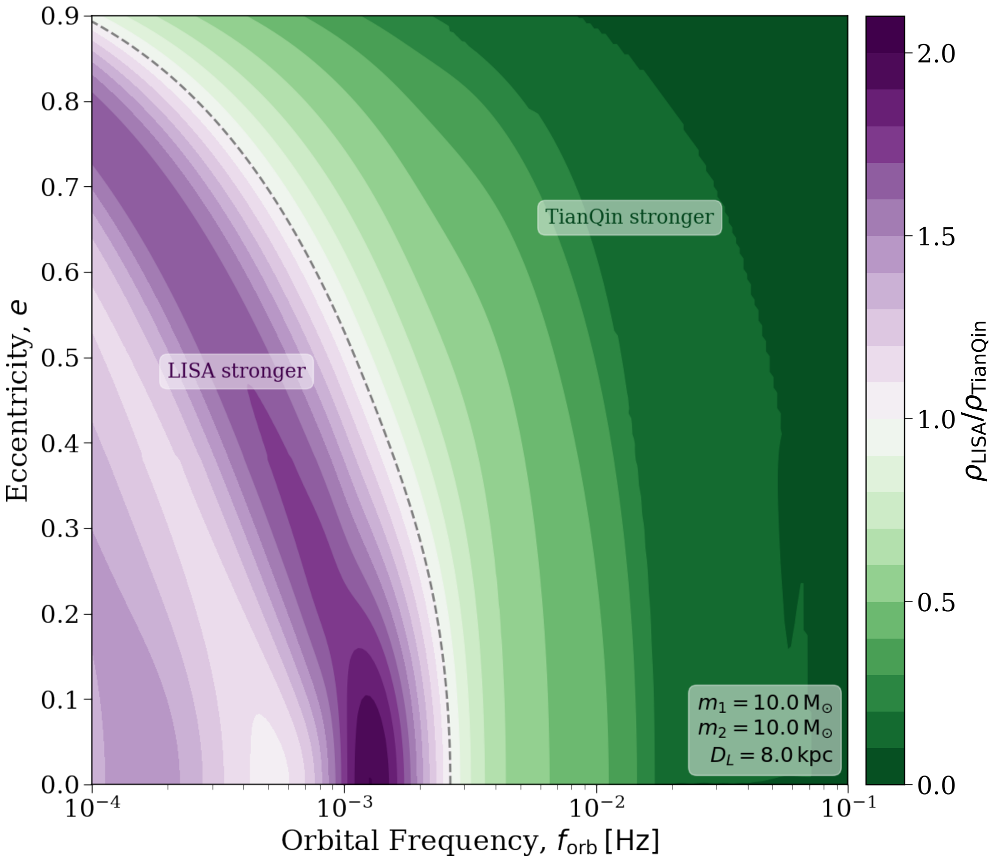

Figure 2.6: The ratio of the SNR in LISA (for a 4-year mission) to the SNR in TianQin. The dashed line indicates the transition at which the SNR is equal in both detectors. We annotate the regions in which either detector has a higher SNR and also annotate the mass and distance of each source in the grid.

Let's consider a binary with the following initial parameters

$$m_1 = m_2 = 15\,{\rm M}_\odot,\ d = 20\,{\rm kpc},$$
$$f_{\rm orb,i} = 3 \times 10^{-5}\,{\rm Hz},\ e_i = 0.5,$$

and use `LEGWORK`'s `evol` module to evolve the system until 100 years before its merger with 1000 linearly spaced timesteps, recording the eccentricity and frequency at each timestep.

We plot this evolution of the eccentricity and frequency in the top two panels of Figure 2.7 as a function of the time before the merger. We see that the binary circularises and increases its orbital frequency as it inspirals as we would expect.

To take this a step further, we can consider the binary at each timestep to be a separate source with the current eccentricity and frequency. It is then trivial to use `LEGWORK` to calculate the SNR for each of these 'sources' and thus attain the SNR evolution, which we plot in the last panel of Figure 2.7.

We see that the SNR increases monotonically over time and sharply increases as the binary approaches its merger. Around 1 Myr before the merger, the SNR reaches the detection

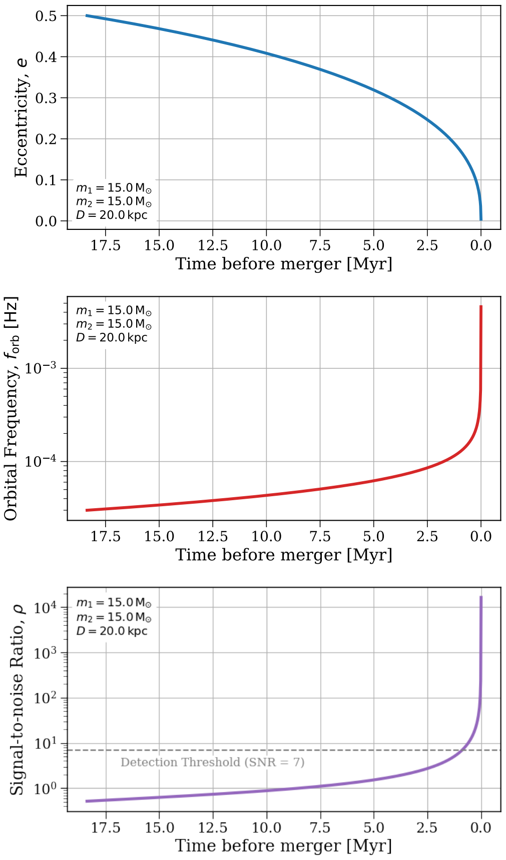

Figure 2.7: The evolution of a binary system's eccentricity (top), orbital frequency (middle) and SNR (bottom). Each panel is annotated with the constant parameters of the system. The SNR is calculated for a 4-year LISA mission. We annotate in a line the bottom panel at SNR = 7 to highlight the moment at which the source becomes detectable.

threshold and thus could then be seen by a 4-year LISA mission.

Note that `LEGWORK` could also be used in this way to find the SNR of any system in which the orbital evolution is known. Thus for a triple system or a binary experiencing gas drag (as long as the evolution of the eccentricity and orbital frequency is known) `LEGWORK` is entirely capable of calculating the SNR evolution.

4.6 LISA verification binaries -

Kupfer et al. (2018) presents the LISA verification binaries, a collection of known binary systems that have gravitational-waves that are strong enough to be detected by LISA. In `LEGWORK`, we provide easy access to this data through the `VerificationBinaries` class, which is a subclass of `Source`. This means that the class works identically to `Source`, but it has the verification binary data (such as their masses and orbital frequencies) pre-loaded into the variables.

In addition to the base variables, this class also includes the designation of each binary and the SNR that Kupfer et al. (2018) computed. We note that this SNR differs from the SNR that `LEGWORK` gives for each binary. This is because Kupfer et al. (2018) run a full, detailed LISA simulation for these binaries, whilst we follow the orbit-averaged approach (see Section 3.3.3), since the former would be intractable for any large number of sources.

As an example of how you could use this data, we use the `VerificationBinaries` class to plot the sources on the LISA sensitivity curve with the SNR calculated by Kupfer et al.

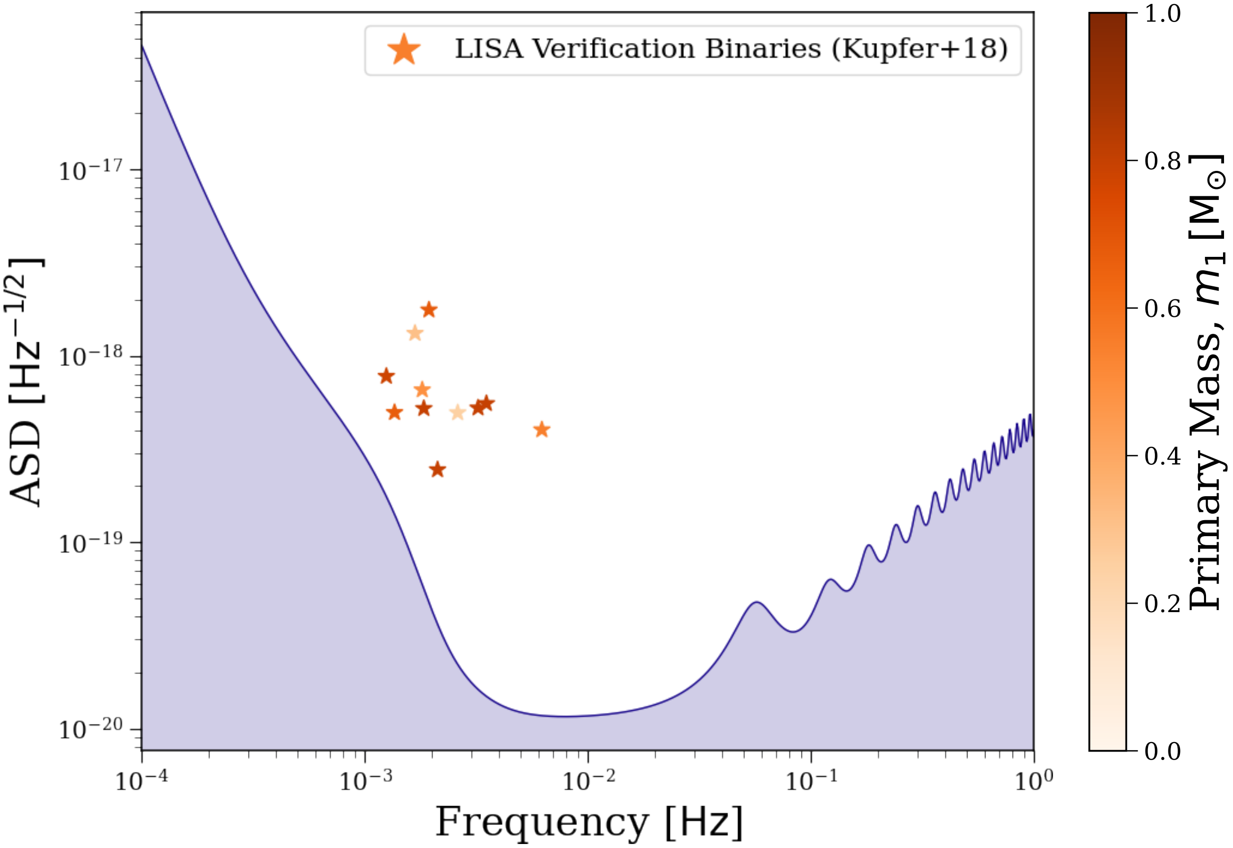

Figure 2.8: The LISA verification binaries from Kupfer et al. (2018) plotted on the LISA sensitivity curve (Robson et al., 2019) for a 4-year LISA mission. Points are coloured by their primary mass.

(2018) for a 4-year LISA mission and colour the binaries by their primary mass. We see that the verification binaries tend to be detected with frequencies between 1-10 mHz and have masses estimated to be between 0.1-1 $\mathrm{M}_\odot$.

5 Conclusion & Summary

We have presented `LEGWORK`, a package designed to aid in calculations for stellar-origin binary sources of mHz gravitational-wave observatories like LISA. We outlined the implementation of orbital evolution due to gravitational-wave emission, gravitational wave strain, SNR, and visualisation modules and provided a detailed derivation for each of the equations required for each module. Finally, we provided several use case examples for how `LEGWORK` can be used to better understand the detectability of compact-object binaries.

CHAPTER 3

GRAVITATIONAL WAVE SOURCES IN OUR GALACTIC BACKYARD

Fitz, home is people. Not a place. If you go back there after the people are gone, then all you can see is what is not there anymore.

– The Fool in *Fool's Fate*

1 Introduction

Since the first direct observation of gravitational waves (Abbott et al., 2016), the number of black hole (BH) and neutron star (NS) binaries observed by ground-based gravitational-wave detectors has rapidly grown (Abbott et al., 2019, 2021a; ?), offering exciting insights into the formation, lives and deaths of massive (binary) stars (e.g. Abbott et al., 2021b).

The Laser Interferometer Space Antenna (LISA, Amaro-Seoane et al., 2017; Colpi et al., 2019) will provide observations in an entirely new regime of gravitational waves. LISA will observe at lower frequencies ($10^{-5} \lesssim f/\,\mathrm{Hz} \lesssim 10^{-1}$) than ground-based detectors and so will enable the study of sources that are imperceptible by ground-based detectors, such as the mergers of supermassive black holes and extreme mass-ratio inspirals (e.g. Begelman et al., 1980; Klein et al., 2016). Moreover, this frequency regime is also of interest for the detection of *local* stellar-mass double compact objects (DCOs) millions of years before their merger. This presents an opportunity for both multi-messenger detections to search for electromagnetic counterparts, as well as multiband gravitational-wave detections that can help to constrain binary characteristics (e.g. Sesana, 2016; Gerosa et al., 2019). In addition, LISA will be able to measure the eccentricities of DCOs, which may yield further constraints on binary evolution, differentiate between formation channels and distinguish between DCO types (e.g. Nelemans et al., 2001; Breivik et al., 2016; Antonini et al., 2017; Rodriguez et al., 2018). Unlike ground-based detectors, the majority of stellar-mass sources detected in LISA will reside in local galaxies, with most situated in the Milky Way. These sources could be used as a probe for the structure of our galaxy (e.g. Korol et al., 2019).

Traditionally, predictions about the detection of stellar-mass sources with LISA focus on double white dwarf (WDWD) binaries, as they are abundantly present in our galaxy and are expected to be the dominant source of stellar-mass binaries that are detectable by LISA (Nelemans et al., 2001; Ruiter et al., 2010; Yu & Jeffery, 2010; Nissanke et al., 2012; Korol et al., 2017; Lamberts et al., 2018). More recently, interest has grown in the detection of NS and BH binaries. Although they are more rare, LISA detections of these sources are potentially valuable for learning more about the evolution and endpoints of massive stars. In this chapter we focus on making LISA predictions for double black hole binaries (BHBHs), black hole neutron star binaries (BHNSs) and double neutron star binaries (NSNSs).

The detection of NSNSs in LISA could improve our understanding of many phenomena. Galactic NSNSs have been observed with electromagnetic signals for several decades (e.g. Hulse & Taylor, 1975; Antoniadis et al., 2016, see also refs. in Tauris et al., 2017) and more recently the mergers of NSNSs have been detected with ground-based gravitational-wave detectors (e.g. Abbott et al., 2017a). A LISA detectable NSNS with a pulsar component close to merger would be ideal for connecting these populations, as the binary could be observed from inspiral to merger. NSNSs (and possibly BHNSs) are useful sources for understanding the origin of r-process elements (e.g. Eichler et al., 1989) as well as the electromagnetic counterparts to gravitational-wave signals, such as kilonovae (e.g. Li & Paczyński, 1998; Metzger, 2017), short gamma-ray bursts (e.g. Berger, 2014), radio emission (e.g. Hotokezaka et al., 2016) and neutrinos (e.g. Kyutoku et al., 2018).

BHBHs in the Milky Way present a greater observational challenge. To date, no BH has been observed in a BHBH binary in the Milky Way, and so LISA could provide the first detection of a Galactic BHBH. The only confirmed BHs in our galaxy have been discovered as components of X-ray binaries with companion stars (e.g. Bolton, 1972; Webster & Murdin, 1972). These detections have observed BHs with masses mainly constrained between 5 and $10\,\mathrm{M}_\odot$ (Corral-Santana et al., 2016), a stark contrast to the more massive BHs observed with LIGO/Virgo that tend to contain at least one BH with a mass greater than $10\,\mathrm{M}_\odot$ (Abbott et al., 2021a). These observations of X-ray binaries suggest the presence of a lower mass gap (from 2-5 $\mathrm{M}_\odot$) in which there are no strong candidates for either black holes or neutron stars (Özel et al., 2010; Farr et al., 2011) but the gap's existence remains an open question (e.g. Kreidberg et al., 2012; Mandel & Müller, 2020; Jayasinghe et al., 2021). Recently there has also been increased discussion over the maximum BH mass in our galaxy, with the claims of a $70\,\mathrm{M}_\odot$ BH (Liu et al., 2019) which has subsequently been challenged (El-Badry & Quataert, 2020; Abdul-Masih et al., 2020; Shenar et al., 2020; Eldridge et al., 2020, see also Liu et al., 2020) and revised measurements of the mass of Cygnus X-1 (Miller-Jones et al., 2021). A sample of BHBHs detected with LISA could possibly help to constrain the BH mass distribution.

One particularly interesting class of potential LISA sources is BHNSs. With the recent detection of two BHNSs by the LIGO scientific collaboration, the existence of these DCOs has been confirmed (Abbott et al., 2021c). However, with only two detections (not including the low-confidence candidate GW190426, Abbott et al. 2021a, or GW190425, GW190814 and GW190917 which have not been ruled out as BHNSs, Abbott et al. 2020a,b, 2024) and no electromagnetic counterparts, the formation rate and properties of BHNSs are still uncertain. A distinctly exciting possibility for BHNSs is the detection of a pulsar–BH system or millisecond pulsar–BH system (Narayan et al., 1991; Pol et al., 2021). These systems could be observed not only by LISA, but also radio telescopes such as MeerKAT and the Square Kilometre Array (SKA, Dewdney et al. 2009), which would help to improve the measurement of individual system parameters and to constrain uncertain binary evolution processes (e.g. Pfahl et al., 2005; Chattopadhyay et al., 2020). Moreover, the pulsar acts as an ultraprecise clock around a BH and thus a pulsar–BH system provides stringent tests of theories of gravity (Keane et al., 2015).

For the purposes of this investigation, we consider the 'classical' isolated binary evolution channel (e.g. Tutukov & Yungelson, 1973, 1993; Smarr & Blandford, 1976; Srinivasan, 1989; Kalogera et al., 2007; Belczynski et al., 2016) in which double compact objects are formed following common-envelope ejection or a phase of highly non-conservative mass transfer (van den Heuvel, 2011; van den Heuvel et al., 2017). We do not, however, account for several alternative proposed formation channels, which could affect the rate and distribution of detectable NS and BH binaries in LISA. These channels include: dynamical formation in dense star clusters (e.g. Sigurdsson & Hernquist, 1993; Portegies Zwart & McMillan, 2000; Miller & Lauburg, 2009; Rodriguez et al., 2015), young/open star clusters (e.g. Ziosi et al., 2014; Di Carlo et al., 2020; Banerjee, 2020; Rastello et al., 2020, 2021) and (active) galactic nuclei discs (e.g. Morris, 1993; Antonini & Rasio, 2016; McKernan et al., 2020), isolated (hierarchical) triple evolution involving Kozai-Lidov oscillations (e.g. Stephan et al., 2016; Silsbee & Tremaine, 2017; Antonini et al., 2017; Toonen et al., 2020), and chemically homogenous evolution through efficient rotational mixing (e.g. de Mink et al., 2009; Mandel & de Mink, 2016; Marchant et al., 2016, 2017; du Buisson et al., 2020).

In this chapter, we present models for the detection rate and distribution of binary properties (masses, frequency, eccentricity, distance, merger time) of BHBHs, BHNSs and NSNSs formed through isolated binary evolution in the Milky Way. We explore the effect of varying physical assumptions in our population synthesis model on our results as well as discuss the effect of extending the LISA mission length and the prospects for distinguishing DCO detections from the WDWD background.

Earlier work on BHBHs, BHNSs and NSNSs in LISA has used a variety of population synthesis codes, Milky Way models and LISA specifications, resulting in a wide range of predictions (Nelemans et al., 2001; Belczynski et al., 2010; Liu & Zhang, 2014; Lamberts et al., 2019; Lau et al., 2020; Breivik et al., 2020b; Sesana et al., 2020; Shao & Li, 2021). We build upon previous efforts but with several important improvements. We explore the effects of varying binary physics assumptions by repeating our analysis for 20 different models (Broekgaarden et al., 2021, 2022) and we compare the effects on the detection rate and distributions of source parameters. We use a model for the Milky Way that accounts for the chemical enrichment history and is calibrated on the APOGEE survey (Majewski et al., 2017; Frankel et al., 2018), whereas most others did not consider the effect of metallicity in detail (see however Lamberts et al., 2019; Sesana et al., 2020). We provide a detailed treatment of the eccentricity of detectable sources, both for the inspiral evolution as well as gravitational wave signal during the LISA mission. Moreover, the grid of binary population synthesis simulations that we use is the most extensive of its kind to date and makes use of the adaptive sampling algorithm STROOPWAFEL (Broekgaarden et al., 2019, 2021). Overall over 2 billion massive binaries were simulated to produce the DCO populations used in this work. We find that this large number of simulations is important to reduce the sampling noise even when using adaptive importance sampling.

All data related to the predictions made in this study are publicly available on Zenodo at Wagg et al. (2021), as are the populations used in our simulations at Broekgaarden (2021a, BHBH) Broekgaarden (2021b, BHNS) and Broekgaarden (2021c, NSNS). We make all code used to produce our results available in a Github repository[1]. In addition, the repository contains step-by-step Jupyter notebooks that explain how to reproduce and change each figure in the chapter. In an earlier chapter, we presented `LEGWORK`[2], the **L**ISA **E**volution and **G**ravitational **W**ave **Or**bit **K**it, a python package designed for making predictions for the detection of sources with LISA, which we use in this chapter.

This chapter is structured as follows. In Section 2, we describe the methods for synthesising a population of binaries, the variations of physical assumptions that we consider, how we simulate the Milky Way distribution of DCOs and our methods for calculating a detection rate for LISA. We present the main results for our fiducial model in Section 3, before exploring the variations in the detectable population when changing our physical assumptions in Section 4. In Section 5 we discuss how and whether these sources can be distinguished from the WDWD population and whether the subpopulations of NSNS, BHNS and BHBH can be distinguished from one other. We further discuss prospects for identifying EM counterparts, the role of different formation channels and overall caveats in the modelling approach. In Section 6, we compare and contrast our methods and findings to previous work.

---

[1] https://github.com/TomWagg/detecting-DCOs-in-LISA

[2] https://legwork.readthedocs.io

2 Method

To produce predictions for the DCOs that are detectable with LISA, we use a synthesised population of DCOs, simulated using the methods described in Section 2.1. In Section 2.2 we describe our model for the Milky Way and how we place DCOs in randomly sampled Milky Way instances. We evolve the orbit of each DCO in a Milky Way instance up to the LISA mission and calculate the detection rate for that instance using the methods presented in Section 2.3.

2.1 Binary population synthesis

We use the grid of 20 binary population synthesis simulations recently presented in Broekgaarden et al. (2021, 2022). This grid of simulations is synthesised using the rapid population synthesis code COMPAS (Stevenson et al., 2017; Vigna-Gómez et al., 2018; Stevenson et al., 2019; Broekgaarden et al., 2019). COMPAS follows the approach of the population synthesis code BSE (Hurley et al., 2000, 2002) and uses fitting formula and rapid algorithms to efficiently predict the final fate of binary systems. The code is open source and documented in the papers listed above, the online documentation[3] and in the methods paper (Riley et al., 2022). The model we refer to as our fiducial model allows for BHs with masses in the heavily debated mass gap, but we also consider 20 physics variations. We summarise the main assumptions and settings relevant for this work in Appendix 8.1. We used these models since Broekgaarden et al. (2022) showed that each physics variation led to a rate prediction that was consistent with the constraints available from ground-based detections in GWTC-3 that were released at that time (**?**). We note that the upper limit on the rates has come down by a factor of two since the original appearance of this catalogue on the arXiv (**?**). We discuss this in further detail in Section 5.3.

The result of the simulations is a sample of binaries, which, for each metallicity Z, have N_{binaries} binaries with parameters

$$\mathbf{b}_{Z,i} = \{m_1, m_2, a_{\text{DCO}}, e_{\text{DCO}}, t_{\text{evolve}}, t_{\text{inspiral}}, w\}, \tag{3.1}$$

for $i = 1, 2, \ldots, N_{\text{binaries}}$, where m_1 and m_2 are the primary and secondary masses, a_{DCO} and e_{DCO} are the semi-major axis and eccentricity at the moment of double compact object (DCO) formation, t_{evolve} is the time between the binary's zero-age main sequence and DCO formation, t_{inspiral} is the time between DCO formation (that is, immediately after the second supernova in the system) and gravitational-wave merger and w is the adaptive importance sampling weight assigned by STROOPWAFEL (Broekgaarden et al., 2019, Eq. 7). We sample from these sets of parameters when creating synthetic galaxies.

2.2 Galaxy synthesis

In order to estimate a detection rate of DCOs with statistical uncertainties, we create a series of random instances of the Milky Way, each populated with a subsample drawn (with replacement) from the synthesised binaries described in Section 2.1.

---

[3] `https://compas.science`

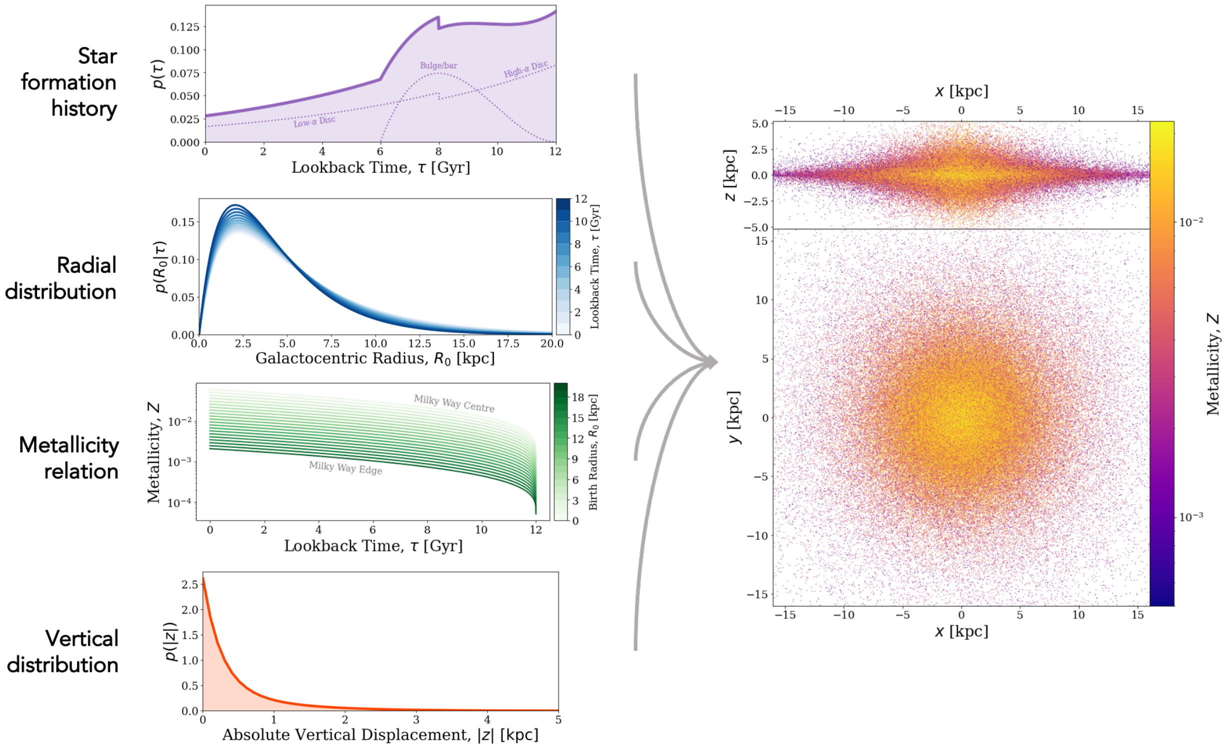

Figure 3.1: A schematic illustrating how we model the Milky Way. The left panel illustrates the different model aspects: star formation history of three galactic components (individually shown in the dotted lines), radial distribution, metallicity-radius-time relation, and height distribution. The right panel shows an example instance of the Milky Way with 250000 binaries shown as points coloured by metallicity. The top panel shows a side-on view and the bottom panel a face-on view. .

Most previous studies that predict a detection rate for LISA place binaries in the Milky Way independently of their age or evolution. We improve upon this as the first study to use an empirically-informed analytical model of the Milky Way that takes into account the galaxy's enrichment history by applying the metallicity-radius-time relation from Frankel et al. (2018). Those authors developed this relation in order to measure the global efficiency of radial migration in the Milky Way and calibrated it using a sample of red clump stars measured with APOGEE (Majewski et al., 2017). We assess the impact of using this improved Milky Way model in Appendix 8.4 and the effect of Galactic models on LISA predictions has been investigated more generally in Storck & Church (2023).

In Section 2.2.1, we outline our model for the Milky Way and in Section 2.2.2 we explain how we combine our population of synthesised DCOs with this Milky Way model.

2.2.1 Milky Way model

Fig. 3.1 shows the distributions and relations outlined in this section and also displays an example random galaxy drawn using this model.

Our model for the Milky Way accounts for the low-[α/Fe][4] disc, high-[α/Fe] disc and a central component approximating a bulge/bar. The low- and high-[α/Fe] discs are often also referred to as the thin and thick discs because the stellar vertical distribution is better fit by a double exponential rather than a single one. However, this doesn't allow one to assign a star to either the thin or thick disk purely based on its height above the Galactic plane. Therefore, we instead use the chemical definition of the two disks (applying the [α/Fe] nomenclature) as there is a clear bimodal distribution in the chemical plane, allowing stars to be assigned to each of the disc components based on their chemical abundances. For each of the three components, we use a separate star formation history and spatial distribution, which we combine into a single model, weighting each component by its present-day stellar mass. Licquia & Newman (2015) gives that the stellar mass of the bulge is $0.9 \times 10^{10}\,\mathrm{M}_\odot$ and the stellar mass of the disc is $5.2 \times 10^{10}\,\mathrm{M}_\odot$, which we split equally between the low- and high-[α/Fe] discs (e.g., Snaith et al., 2014).

Star formation history: We use an exponentially declining star formation history (Frankel et al., 2018) (inspired by the average cosmic star formation history) for the combined low- and high-[α/Fe] discs,

$$p(\tau) \propto \exp\left(-\frac{(\tau_m - \tau)}{\tau_{\mathrm{SFR}}}\right), \tag{3.2}$$

where τ is the lookback time (the amount of time elapsed between the binary's zero-age main sequence and today), $\tau_m = 12\,\mathrm{Gyr}$ is the assumed age of the Milky Way and $\tau_{\mathrm{SFR}} = 6.8\,\mathrm{Gyr}$ is the star formation timescale (Frankel et al., 2018). The two discs form stars in mutually exclusive time periods, such that the high-[α/Fe] disc forms stars in the early history of the galaxy (8–12 Gyr ago) and the low-[α/Fe] disc forms stars more recently (0–8 Gyr ago). Both distributions are normalised so that an equal amount of mass is formed in each of the two components over their respective star forming periods.

The star formation history of the bulge/bar of the Milky Way has many uncertainties due to the (1) sizeable age measurement uncertainties at large ages in observational studies, (2) complex selection processes affecting the observed age distributions, and (3) formation mechanisms that are still under debate. However, the central bar (which we assume to dominate here) was shown to contain stars with an extended age range, with most observed stars between 6 and 12 Gyr with a younger tail of ages that could come from the subsequent secular growth of the Galactic bar (e.g., Bovy et al., 2019). To model the bulge/bar's age distribution more realistically than in previous studies (which assume an old bulge coming from a single starburst), we choose to adopt a more extended star formation history using a $\beta(2,3)$ distribution, shifted and scaled such that stars are only formed in the range $[6, 12]$ Gyr. We show these distributions in the top left panel of Fig. 3.1.

Radial distribution: For each of the three components we employ the same single exponential distribution (but with different scale lengths)

$$p(R) = \exp\left(-\frac{R}{R_d}\right)\frac{R}{R_d^2}, \tag{3.3}$$

where R is the Galactocentric radius and R_d is the scale length of the component. For the

---

[4] Nomenclature used to describe the enhancement of α elements compared to iron in stellar atmospheres

low-[α/Fe] disc, we set $R_d = R_{\rm exp}(\tau)$, where $R_{\rm exp}(\tau)$ is the scale length presented in Frankel et al. (2018, Eq. 6)

$$R_{\rm exp}(\tau) = 4\,{\rm kpc}\left(1 - \alpha_{R_{\rm exp}}\left(\frac{\tau}{8\,{\rm Gyr}}\right)\right), \tag{3.4}$$

where $\alpha_{R_{\rm exp}} = 0.3$ is the inside-out growth parameter[5]

This scale length accounts for the inside-out growth of the low-[α/Fe] disc and hence is age dependent. We assume $R_d = (1/0.43)\,{\rm kpc}$ for the high-[α/Fe] disc (Bovy et al., 2016, Table 1) and $R_d = 1.5\,{\rm kpc}$ for the bulge/bar component (Bovy et al., 2019). Note that in this way we have approximated the bulge/bar component as being axi-symmetric, which is sufficient for our purposes. We show the combination of these distributions in the second panel on the left in Fig. 3.1.

Vertical distribution: Similar to the radial distribution, we use the same single exponential distribution (but with different scale heights) for each component, given by

$$p(|z|) = \frac{1}{z_d}\exp\left(-\frac{z}{z_d}\right), \tag{3.5}$$

where z is the vertical displacement above the Galactic plane and z_d is the scale height. We set $z_d = 0.3\,{\rm kpc}$ for the low-[α/Fe] disc (McMillan, 2011) and $z_d = 0.95\,{\rm kpc}$ for the high-[α/Fe] disc (Bovy et al., 2016). For the bulge/bar, we set $z_d = 0.2\,{\rm kpc}$ (Wegg et al., 2015). We show the combination of these distributions in the bottom left panel of Fig. 3.1.

Metallicity-radius-time relation: To account for the chemical enrichment of star forming gas as the Milky Way evolves, we adopt the relation given by (Frankel et al., 2018, Eq. 7)

$$\begin{aligned}[\mathrm{Fe/H}](R,\tau) = F_m &+ \nabla[\mathrm{Fe/H}]R \\ &- \left(F_m + \nabla[\mathrm{Fe/H}]R^{\rm now}_{[\mathrm{Fe/H}]=0}\right)f(\tau),\end{aligned} \tag{3.6}$$

where

$$f(\tau) = \left(1 - \frac{\tau}{\tau_m}\right)^{\gamma_{[\mathrm{Fe/H}]}}, \tag{3.7}$$

$F_m = -1\,{\rm dex}$ is the metallicity of the gas at the center of the disc at $\tau = \tau_m$, $\nabla[\mathrm{Fe/H}] = -0.075\,{\rm kpc}^{-1}$ is the metallicity gradient, $R^{\rm now}_{[\mathrm{Fe/H}]=0} = 8.7\,{\rm kpc}$ is the radius at which the present day metallicity is solar and $\gamma_{[\mathrm{Fe/H}]} = 0.3$ sets the time dependence of the chemical enrichment. We convert this to the representation of metallicity that we use in this paper by applying (e.g. Bertelli et al., 1994)

$$\log_{10}(Z) = 0.977[\mathrm{Fe/H}] + \log_{10}(Z_\odot). \tag{3.8}$$

Although Frankel et al. (2018) only fit this model for the low-[α/Fe] disc, we also use this metallicity-radius-time relation for the high-α disc and the bar, but focusing on the chemical tracks more representative to the inner disc and large ages. Sharma et al. (2021) showed that using a simple continuous model for both the low- and high-[α/Fe] discs, the Milky Way abundance distributions could be well reproduced. Empirically, the abundance tracks

---

[5] We find that $R_{\rm exp}(\tau) = 4$ kpc fits the data well and adopt this value rather than the 3 kpc quoted in Frankel et al. (2018), which was a fixed parameter (not a fit).

in the [α/Fe]-[Fe/H] plane (and other elements) of the stars in the bulge/bar follow the same track as those of the old stars in the Solar neighbourhood (Griffith et al., 2021; Bovy et al., 2019, Fig. 7,), which motivates our modelling choice to use the same metallicity-radius-time relation.

2.2.2 Combining population and galaxy synthesis

For each Milky Way instance, we randomly sample the following set of parameters

$$\mathbf{g}_j = \{\tau, R, Z, z, \theta\} \tag{3.9}$$

for $j = 1, 2, \ldots, N_{\rm MW}$, where we set $N_{\rm MW} = 2 \times 10^5$, τ, R, Z and z are defined and sampled using the distribution functions specified in Section 2.2.1, θ is the azimuthal angle sampled uniformly on $[0, 2\pi)$ and Z is the metallicity. Fig. 3.1 shows an example of a random Milky Way instance created with these distributions. This shows how these distributions translate to positions and illustrates the gradient in metallicity over radius.

We match each set of galaxy parameters $\mathbf{g}_j$, to a random set of binary parameters $\mathbf{b}_{Z,i}$, by drawing a binary from the closest metallicity bin to the metallicity in $\mathbf{g}_j$. If the metallicity in $\mathbf{g}_j$ is below the minimum COMPAS metallicity bin ($Z = 10^{-4}$), we use this minimum bin. If the metallicity in $\mathbf{g}_j$ is above the maximum COMPAS metallicity bin ($Z = 0.03$), we use a randomly selected bin from the five highest metallicity bins[6].

Each binary is likely to move from its birth orbit. Although all stars in the Galactic disc experience radial migration (Sellwood & Binney, 2002; Frankel et al., 2018), DCOs generally experience stronger dynamical evolution as a result of the effects of both Blaauw kicks (Blaauw, 1961) and natal kicks (e.g. Hobbs et al., 2005).

The magnitude of the systemic kicks are typically small compared to the initial circular velocity of a binary at each Galactocentric radius. Therefore, we expect that kicks will not significantly alter the overall distribution of their positions (see however, e.g., Brandt & Podsiadlowski 1995; Abbott et al. 2017b). Given this, and for the sake of computational efficiency, we do not account for the displacement due to systemic kicks in our analysis.

2.3 Gravitational wave detection

We use the Python package `LEGWORK` (Wagg et al., 2022b) to evolve binaries and calculate their LISA detectability. For a full derivation of the equations given below see (Wagg et al., 2022b, Section 3), or the `LEGWORK` documentation 📖.

2.3.1 Inspiral evolution

Each binary loses orbital energy to gravitational waves throughout its lifetime. This causes the binary to shrink and circularise over time. In order to assess the detectability of a binary, we need to know its eccentricity and frequency at the time of the LISA mission. For each

---

[6] We spread the binaries over the five highest bins, rather than just the highest bin, as we found that using a single bin led to unphysical artifacts in our results. These artifacts arose because the small population of binaries in the highest bin were oversampled.

binary in our simulated Milky Way, we know that the time from DCO formation to today is $\tau - t_{\rm evolve}$ and that the initial eccentricity and semi-major axis are $e_{\rm DCO}$ and $a_{\rm DCO}$. We find the eccentricity of the binary at the start of the LISA mission, $e_{\rm LISA}$, by numerically integrating its time derivative (Peters, 1964, Eq. 5.13) given the initial conditions. This can be converted to the semi-major axis at the start of LISA, $a_{\rm LISA}$(Peters, 1964, Eq. 5.11), which in turn gives the orbital frequency, $f_{\rm orb,LISA}$, by Kepler's third law since we know the component masses.

2.3.2 Binary detectability

We define a binary as detectable if its gravitational wave signal has a signal-to-noise ratio (SNR) of greater than 7 by the end of the LISA mission (e.g. Breivik et al., 2020b; Korol et al., 2020). The sky-, polarisation- and orientation-averaged signal-to-noise ratio, ρ, of an inspiraling binary can be calculated with the following (e.g. Finn & Thorne, 2000)

$$\rho^2 = \sum_{n=1}^{\infty} \int_{f_{n,i}}^{f_{n,f}} \frac{h_{c,n}^2}{f_n^2 S_{\rm n}\left(f_n\right)} \, {\rm d}f_n \,, \tag{3.10}$$

where n is a harmonic of the gravitational wave signal, $f_n = n \cdot f_{\rm orb}$ is the frequency of the $n^{\rm th}$ harmonic of the gravitational wave signal, $f_{\rm orb}$ is the orbital frequency, $S_{\rm n}(f_n)$ is the LISA sensitivity curve at frequency f_n (e.g. Robson et al., 2019) and $h_{c,n}$ is the characteristic strain of the $n^{\rm th}$ harmonic, given by (e.g. Barack & Cutler, 2004)

$$h_{c,n}^2 = \frac{2^{5/3}}{3\pi^{4/3}} \frac{(G\mathcal{M}_c)^{5/3}}{c^3 D_L^2} \frac{1}{f_{\rm orb}^{1/3}} \frac{g(n,e)}{nF(e)}, \tag{3.11}$$

where D_L is the luminosity distance to the source, $f_{\rm orb}$ is the orbital frequency, $g(n,e)$ and $F(e)$ are given in Peters & Mathews (1963) and $\mathcal{M}_c$ is the chirp mass, defined as

$$\mathcal{M}_c = \frac{(m_1 m_2)^{3/5}}{(m_1+m_2)^{1/5}}. \tag{3.12}$$

Note that increasing the length of the LISA mission allows more time for a DCO to evolve over the mission. Therefore the frequency limits in Eq. 3.10 are dictated by the LISA mission length. The SNR generally scales as $\sqrt{T_{\rm obs}}$ (with exceptions for sources very close to merging) and thus the SNR of a typical source in a 10-year LISA mission is approximately 1.58 ($=\sqrt{10/4}$) times stronger than in a 4-year mission.

We use `LEGWORK` (Wagg et al., 2022b) to calculate the signal-to-noise ratio for each binary and the package ensures that enough harmonics are computed for each binary such that the error on the gravitational-wave luminosity remains below 1%.

2.3.3 Detection rate calculation

For each physics variation model and DCO type, we first convert the COMPAS simulation results into a total number of DCOs in the Milky Way, $N_{\rm DCO}$. We do this by integrating the full mass and period distributions and stars and normalising to the total Milky Way mass. For more details see Appendix 8.2.

We then determine the fraction of binaries that are detectable in each Milky Way instance by summing the adaptive importance sampling weights of the binaries that have an SNR greater than 7, and dividing by the total weights in the simulation. We multiply this fraction by $N_{\rm DCO}$ to find a detection rate (which we write as a total number of detections per LISA mission)

$$N_{\rm detect} = \frac{\sum_{i=0}^{N_{\rm MW}} w_i \phi(i)}{\sum_{i=0}^{N_{\rm MW}} w_i} N_{\rm DCO}, \tag{3.13}$$

where $\phi(i) = 1$ if a binary is detectable and 0 otherwise. We calculate the detection rate by Monte Carlo sampling 2500 Milky Way instances (each containing 200,000 DCOs) for each DCO type and every physics variation in order to obtain values for the uncertainty on the expected detection rate.

3 Results I - Predictions for LISA sources

In this section we present our predictions for the population of detectable LISA sources for our fiducial model. In total we expect, on average, 124 detections in a 4-year LISA mission, of which 74, 42 and 8 are BHBHs, BHNSs and NSNSs respectively, based on our fiducial simulations. In the remainder of this section, we discuss where the sources are expected relative to LISA's sensitivity curve (Sec. 3.1), their properties (Sec. 3.2), their locations in the Milky Way (Sec. 3.3), their formation channels (Sec. 3.4) and finally we discuss the expected SNR and how accurately we expect that the parameters can be measured (Sec. 3.5). Note that all results shown in this section are based on our fiducial simulations. A discussion of the impact of variations in the physics assumptions is provided in Sec. 4.

3.1 The LISA sensitivity curve and the population of detectable sources

We show the expected LISA sensitivity curve (Robson et al., 2019) in Fig. 3.2, which includes the confusion noise arising from the Galactic WDWD population, and overplot our predictions for the distribution of detectable sources. Eccentric systems emit gravitational waves in multiple harmonic frequencies ($n f_{\rm orb}$, with $n = 2, 3, \dots$). We choose to plot them at the x-coordinate that corresponds to the frequency of the harmonic that individually accumulates the largest SNR. For circular systems, the x-coordinate simply corresponds to $2 f_{\rm orb}$. The y-coordinate indicates the strength of the signal (or to be more precise, the amplitude spectral density, ASD), including the contribution from *all* harmonics.

For reference, we show dotted lines to indicate where a hypothetical binary system would reside assuming a given distance from earth (diagonal lines) and a fixed remaining inspiral time (vertical lines). For each line we assume the binary is circular and has a chirp mass equal to the average of the sample ($\langle \mathcal{M}_c \rangle$, annotated in each panel). We also overplot the LISA verification binaries (star symbols, Kupfer et al. 2018).

We observe several features in Fig. 3.2 that are worth discussing and explaining. We note

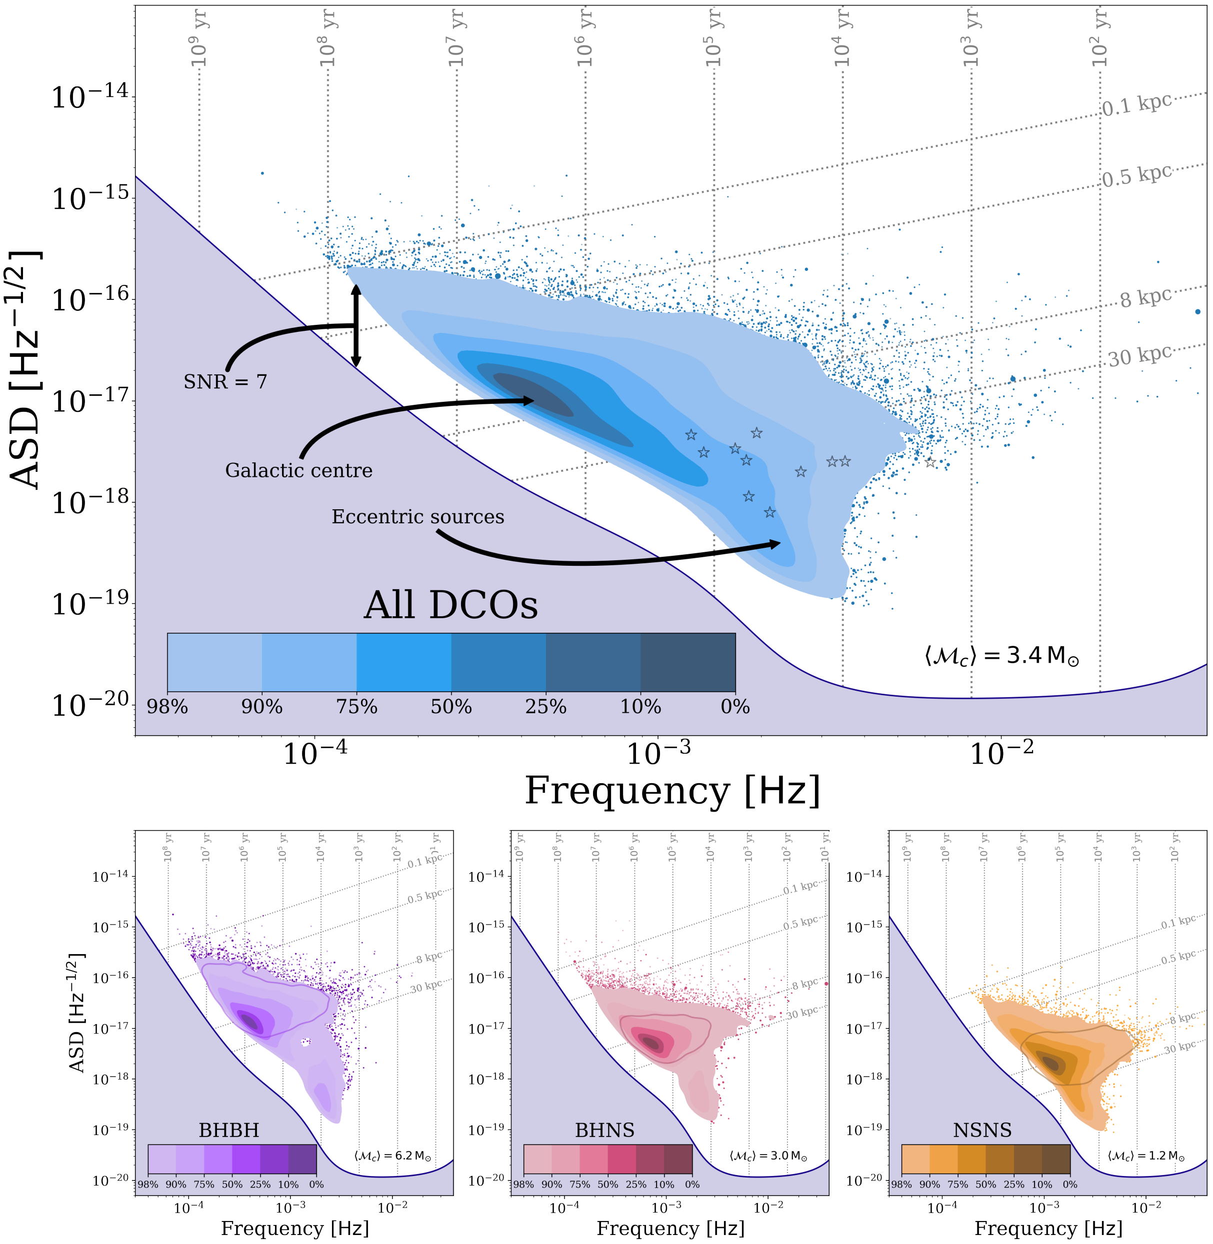

Figure 3.2: The LISA sensitivity curve is shown together with the density distribution of the characteristic strain and the dominant frequency for all detectable sources in our simulations (top) and separated by type (bottom). Contours show the percentage of the population enclosed. The remaining 2% of the population is shown as dots with a size that scales with the sampling weight. To guide the interpretation, we show reference lines marking where a circular binary would reside for a given distance (diagonal line) and remaining inspiral time (vertical lines), assuming an average chirp mass $\langle \mathcal{M}_c \rangle$, orientation and sky location. The coloured lines in the bottom panels show a contour that encloses 90% of the population that is circular. LISA verification binaries are overplotted in the top panel (star symbols). See also Fig. 3.15 and Sec. 3.1 for a discussion.

that some of these have also been described in earlier studies (see Sec. 6). Firstly, we note the empty band that separates the LISA sensitivity curve and the detectable population. This is the result of the criteria for detection where we require SNR > 7.

We further note that the detectable population is concentrated on the left side of LISA's sensitivity window. The peak is located at a frequency of about 0.4 mHz, which is much lower than the frequency where LISA will be most sensitive (about 10 mHz). This can be understood from the acceleration of inspiraling DCOs as they evolve towards higher frequencies. DCOs typically form with wide orbits (low frequencies) that would not be detectable yet. Their orbits shrink as they lose angular momentum in the form of gravitational waves leading to an increase of their orbital frequency until they become detectable. These systems are increasingly rare because they evolve faster and faster towards higher frequency as the inspiral accelerates, even though the signal emitted by a more compact binary is stronger (Peters, 1964). The vertical grid lines show these rapidly decreasing inspiral times at increasing frequencies. Most of the population is thus expected to reside at low frequencies.

In the lower three panels, we show the contributions of the different types of sources. Comparing them, one can observe the shift in the frequency at which the peak is located, at 0.3 mHz, 0.7 mHz and 1 mHz for BHBH, BHNS and NSNS systems respectively. This is a result of the difference in chirp mass. A higher mass system can emit at lower frequency and still produce the minimum SNR needed for detection. We note that this effect can be used to distinguish the heavier DCOs that we discuss in this work from WDWD systems, at least probabilistically (see Sec. 5.1.1). In the same way, this also explains the offset in frequency between the population of sources we predict and the LISA verification binaries.

Inspecting the dotted reference lines, we note that the peak of the density distribution of observable sources coincides with the location expected for circular systems at 8 kpc, which is the distance to the centre of the Milky Way. As can be seen best in the lower panels, the reference lines for 0.1 and 30 kpc enclose the majority of systems, as expected given the dimensions of the Milky Way.

There is a distinct subpopulation of binaries, most clearly visible in the lower panels as an offshoot that extends downwards to ASDs of $10^{-19}\,\mathrm{Hz}^{-1/2}$, especially around 2 mHz. This offshoot is almost uniquely composed of eccentric binaries, as can be seen in Fig. 3.15, which shows a similar figure but colouring individual systems by their eccentricity. This can also be seen, albeit indirectly, from the contour lines shown in the bottom panels of Fig. 3.2, which encompass 90% of the *circular* sources in each population. This contour does not include the offshoot. We conclude that eccentric sources occupy a very different region in this diagram.

3.2 Properties of the detectable systems

In Fig. 3.3, we show the 1D distribution of several individual parameters of the population of detectable binaries together with the 1- and 2-σ uncertainties obtained via bootstrapping. These uncertainties represent the fluctuations in our results over different random instances of the Milky Way. The distributions shown here are approximated by kernel density estimators, corrected for edge effects by mirroring the sample (Schuster, 1985).

Orbital Frequency — The orbital frequency distributions for BHBHs, BHNSs and NSNSs (Fig. 3.3a) peak at progressively increasing frequencies. As mentioned in Sec. 3.1, this is

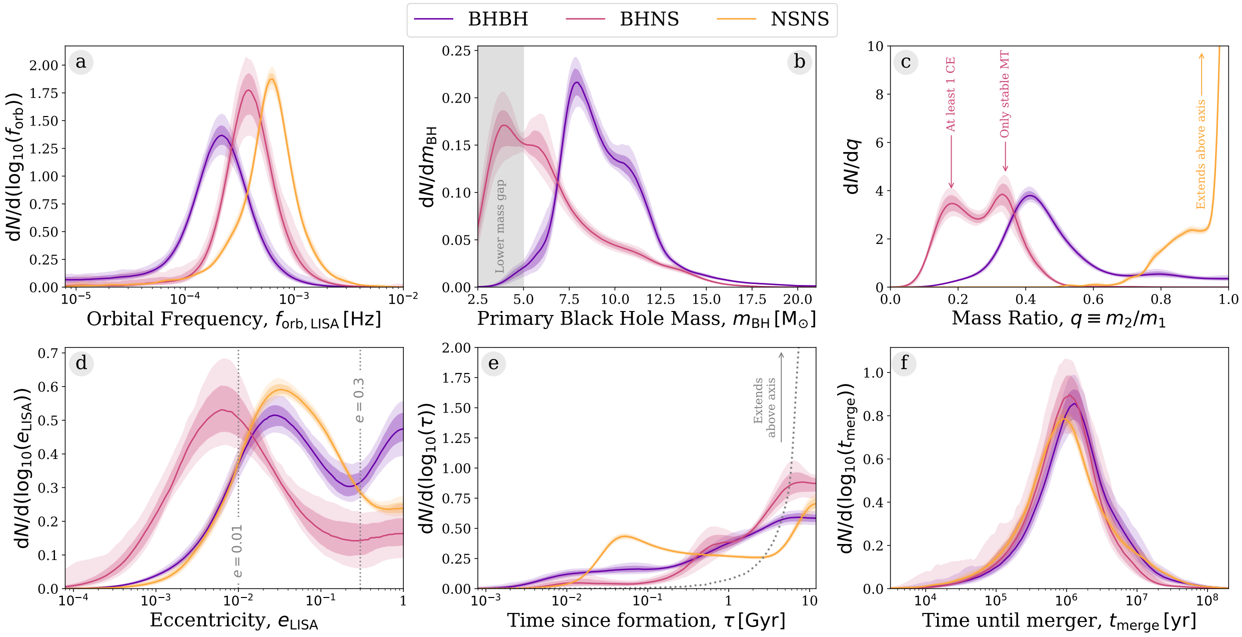

Figure 3.3: Properties of detectable systems for a 4-year LISA mission in our fiducial model. Each panel shows a kernel density estimator for a single property, coloured by DCO type. Each curve has been individually normalised. The shaded areas show the 1- and 2-σ sampling uncertainties (obtained via bootstrapping). The dotted line in panel e shows the star formation history we have assumed in our Milky Way model. See Sec. 3.2 for a discussion.

because a higher mass DCO at the same distance and eccentricity requires a lower frequency to produce the same signal-to-noise ratio and thus be detected. The distributions appear nearly symmetric, but the left tail of the distribution is more populated, especially for the BHBHs. This is due to the contribution of highly eccentric binaries, which are most abundant in the BHBH population. These systems are still detectable by LISA, despite their low orbital frequency, as the high eccentricity means that the majority of the GW signal is emitted at higher harmonics, where LISA is more sensitive.

Black Hole Mass — In Fig. 3.3b, we show the distribution of primary black hole masses for BHNS and BHBH systems. The BHNS distribution shows a bimodality. This is due to the bimodal mass ratio distribution (see Fig. 3.3c) since most NS components have approximately the same mass.

For both the BHBHs and BHNSs, we see that the black hole mass distribution favours low masses. About 90% of BHs have masses below $12\,\mathrm{M}_\odot$, in our fiducial simulations shown here. This is in contrast with observations from ground-based GW detectors, where heavy BHs with masses of about $30\,\mathrm{M}_\odot$ and higher have been common. There are two main reasons for this discrepancy. First, the population of BHs in the Milky Way (and, in particular, those detectable by LISA due to their recent formation times, see Fig. 3.3e) primarily come from progenitors that formed from high metallicity gas according to our simulations. Stellar winds are stronger at high metallicity leading to increased mass loss. This affects the mass of the most massive black holes that can be formed (Belczynski et al., 2010). Secondly, ground-based detectors are strongly biased towards high mass systems, since they can be detected out to larger distances and thus a greater volume is probed. In contrast, LISA has no such bias and is, in principle, sensitive to inspiraling DCOs throughout the entire Milky Way regardless of their mass, as long as we catch them when they are emitting in the LISA band. For this reason, LISA is more likely to detect the more common lower mass DCOs.

We also note that our BH mass distribution extends down below $5\,\mathrm{M}_\odot$ to $2.5\,\mathrm{M}_\odot$ which is our fiducial maximum neutron star mass. The BHs in this simulation fill the so-called “lower mass gap” marked as a grey band (Özel et al., 2010; Farr et al., 2011), see also Shao & Li (2021) who recently also pointed this out. This prediction is sensitive to adopted model for fallback during the SN explosion as we discuss this further in Section 4.2.1.

Mass Ratio — The mass ratio distributions for each DCO type are very distinct from one another, as can be seen in Fig. 3.3c. The majority of NSNSs have a mass ratio close to unity, with 90% of systems having $q > 0.8$, where $q \equiv m_2/m_1$. The reason for the concentration around equal masses is that most NSs are formed either through electron-capture supernovae (ECSN) or from low mass stars in our simulations. We assume a remnant mass for any NS formed through ECSN of $1.26\,\mathrm{M}_\odot$ (see Sec. 8.1.2). The remnant mass prescription that we use assumes a fixed fallback mass for any star with a CO core mass less than $2.5\,\mathrm{M}_\odot$, such that many NSs end up with an identical mass of $1.278\,\mathrm{M}_\odot$ (see Fryer et al., 2012, Eq. 19). This means that many NSs are formed with equal masses and hence we see a mass ratio distribution peaked around unity.

In contrast, only 8% of detectable BHBHs are formed with $q > 0.8$ and the distribution peaks around $q = 0.4$, we find that the strong stellar winds in our (typically high-metallicity) progenitors are the reason behind this, as we will discuss below. However, we note that the preference for unequal mass ratios is possibly at tension with the current inferred distribution from GWTC-3, which prefers equal masses (**?**, Fig. 10). The population detectable by ground-based detectors and the population to which LISA is sensitive may differ, for example because of difference in the underlying metallicity-specific star formation history or due to the lower masses of the LISA detectable systems, but it is unclear if this is sufficient to explain the discrepancy.

BHBHs with unequal masses typically come from progenitors that also had more extreme mass ratios at birth (90 and $30\,\mathrm{M}_\odot$ are typical for the progenitors of detectable BHBH systems in our simulations). The primary in such systems experiences strong mass loss by winds before filling its Roche lobe. This mostly happens during its early hydrogen-shell burning phase. The wind mainly reduces the mass of the envelope, but does not have a very strong effect on the core. By the time the primary fills its Roche lobe, it has become less massive and the mass ratio is closer to one. This favours stable mass transfer and most detectable BHBHs form through stable mass transfer (see Section. 3.4). The massive core of the primary star typically becomes the more massive BH. Accretion on the secondary star is limited and the secondary eventually provides the less massive black hole.

At the same time, stellar wind mass loss disfavours the formation of black holes with comparable masses. Such systems would have originated from progenitors that also started with comparable masses. The rather massive secondaries in these systems (especially after they accreted from the primary) experience very strong stellar wind mass loss due to LBV-like eruptions. This limits the amount by which they can expand (e.g. van Son et al., 2022a). At the same time, wind mass loss (and also stable and more conservative mas transfer in these systems) lead to widening of the orbit. Both effects, the reduced expansion and increased widening of the orbit, tend to prevent the secondaries from being able to fill their Roche lobe. This thus limits the number of systems that experience the reverse mass transfer or common-envelope phase needed shrink the orbits and make them tight enough to be detected as gravitational-wave sources.

We find that detectable BHNSs have even more unequal mass ratios. Moreover, the mass ratio distribution is bimodal, where the two peaks arise from two distinct formation scenarios. Around two thirds of detectable BHNSs experience at least one common-envelope event, whilst the last third are formed through only stable mass transfer. The first peak at $q = 0.18$ is from systems that experience at least one common-envelope phase and occurs at the expected mass ratio, which approximately follows the mean BH mass ($\sim 6.5\,\mathrm{M}_\odot$) and NS mass ($\sim 1.2\,\mathrm{M}_\odot$). Yet we also see a second peak at higher mass ratios around $q = 0.34$, which arises from the fraction of the population that underwent only stable mass transfer phases. The stability of mass transfer depends on the mass ratio, preferentially forming systems with more equal masses, i.e. at higher q.

Eccentricity — In Fig. 3.3d we show the eccentricity distributions. We find that most systems (73%) will have eccentricities larger than 0.01 during the LISA mission, which should in principle be detectable according to Nishizawa et al. (2016). This means that we will potentially be able to use eccentricity to distinguish these sources from WDWDs, which are expected to have little to no eccentricity (see Sec. 5.1.1). We note that several previous studies assumed all systems to be circular when calculating the detection rates (e.g. Liu & Zhang, 2014; Lamberts et al., 2018; Sesana et al., 2020). We discuss the impact of this assumption in Section 6.

For systems with eccentricities higher than $e \gtrsim 0.3$, most gravitational wave energy is emitted in higher harmonics. Such systems are more rare, but we find them to be significant among the BHBH population, where they account for 21% of systems. Detectable BHBHs in our simulation (and, in particular, those that are eccentric) are primarily systems formed through the stable mass transfer channel (see Fig. 3.17). These systems are still relatively wide (compared to those formed through the CE channel) immediately prior to formation of the second BH, which makes them more easily affected by kicks. If the kick is oriented roughly in opposite direction to the orbital motion and has a velocity that is of similar magnitude as the orbital velocity, it will lead to the formation of a highly eccentric system.

It is rare to get such a “lucky kick”, but there are a few effects that favour this for BHBHs. The kicks of BHs are reduced by fallback and they are thus less likely to disrupt the system. Moreover, because of their higher masses, BHBHs can be observed already at lower orbital frequencies. This means that they have not had as much time to circularise and so still have significant eccentricity by the time of the LISA mission. Finally, LISA favours the detection of eccentric systems, if all other properties are held fixed. This is because the gravitational-wave emission is stronger (Eq. 3.28) and the energy is emitted at higher frequencies (Peters & Mathews, 1963, Eq. 20) where LISA is more sensitive.

The lower abundance of highly eccentric systems among the NSNS and BHNS systems may seem counter-intuitive since neutron stars are lower mass and would be more strongly affected by natal kicks, which one may expect to lead to more eccentric systems. However, the majority of NSs in our simulations are formed through ECSN and USSN and for these types of supernovae we draw from a Maxwellian with $\sigma^{\rm 1D}_{\rm rms} = 30\,{\rm km\,s^{-1}}$. Thus the kicks received by NSs in our simulations are often much smaller than for BHs.

Time since formation — In Fig. 3.3e we show how long ago the LISA detectable DCOs formed. Star formation was highest at early times $6-12$ Gyr ago, after which it declined. In contrast, the LISA detectable DCOs primarily formed in the relatively *recent* history of our Milky Way, about 2 Gyr ago. This reflects the fact that binaries in our simulation typically take about a Gyr to merge.

When comparing the distribution of formation times for the three different DCO types we see that NSNSs are most strongly concentrated at recent times, followed by BHBHs and then BHNSs. To understand this it is helpful to consider how the inspiral time scales with various parameters (Peters, 1964)

$$t_{\rm inspiral} \propto \frac{a^4}{(m_1+m_2)^3} \cdot \frac{q}{(1+q)^2} \cdot (1-e^2)^{-7/2}. \tag{3.14}$$

The inspiral time depends most strongly on the separation at DCO formation, a, and this is where the three types also differ most strongly (see Fig. 3.16). The detectable NSNS systems have the tightest orbits at DCO formation. The median of the distribution of separations at DCO formation, $\langle a_{\rm DCO}\rangle_{\rm med}$, relate as 8:3:1 for detectable BHBH:BHNS:NSNS in our simulations. This results in increase of the inspiral time by a factor of about 4000:80:1.

The distribution of separations at DCO formation is also the reason that the time since formation for NSNSs shows a bimodality. For NSNSs, the separation at DCO formation peaks around 2×10^{-2} AU but the distribution has a second population that extends to much lower separations relative to the the BHNS and BHBH distributions. It is this population of shorter separations that corresponds to the NSNSs that were formed more recently (within the last 100 Myr).

The total masses affect the inspiral time to the third power and this where the heavier BHBH systems are favoured. The median total masses differ by ratios of 6:4:1 for detectable BHBH:BHNS:NSNS in our simulations, impacting the inspiral times such that they are a factor of 200:60:1 shorter, partially counteracting the effect of the separations. The term depending on the mass ratio q only varies by about 30% for the mass ratio ranges considered here and so is not of interest. The eccentricity term is not of importance for mildly eccentric systems, $f(e_{\rm DCO} \leq 0.3) \leq 1.4$ but of large importance for the very eccentric $f(e_{\rm DCO} \geq 0.9) \geq 300$. The fraction of highly eccentric systems with $e_{\rm DCO} > 0.9$ is 33%, 16%, and 8% of for BHBH, BHNS and NSNS respectively, see also Fig. 3.16.

We conclude that the shorter median separations at DCO formation are the main reason why NSNS are most strongly peaked at short delay times. They are followed by BHBHs rather than BHNSs due to the high masses and substantial eccentricities of BHBHs.

Time until merger — Fig. 3.3f shows the remaining time until merger for each of the DCO types at the start of LISA mission. The distributions are strikingly similar and peak with merger times of around a Myr.

The merger time is a function of the mass, frequency and eccentricity of the sources, such

Component	Formation	Detectable			
	All	All	BHBH	BHNS	NSNS
Low-α disc	42.5%	86%	89%	82%	85%
High-α disc	42.5%	6%	5%	8%	12%
Bulge/bar	15.0%	8%	6%	10%	3%

Table 3.1: Percentage of systems in each Galactic component.

that more massive, higher frequency and more eccentric sources merge faster (Peters, 1964, Eq. 5.14). So, despite the fact that each DCO type often has higher values in any one of these properties, the convolution of all three tends to negate the differences. For example, NSNSs have the highest orbital frequencies and are mildly eccentric whilst BHNSs have moderate orbital frequencies and are more circular. However, BHNSs are more massive in general and so the overall merger times are distributed very similarly for both DCO types.

3.3 Distribution in the Milky Way

For our Milky way model we consider three different components, a low-[α/Fe] ("thin") disc, an older high-[α/Fe] ("thick") disc and a bulge/bar (see Sec. 2.2). In Table 3.1 we summarise the number of detections originating from each of these components. Despite the fact that only 42.5% of systems are formed in the low-[α/Fe] disc, we find that 86% of the detections originate from this component. This is because most detectable systems were formed relatively recently (see Fig. 3.3e) and so the high-[α/Fe] disc and bulge are effectively too old to contribute many detectable systems. Nevertheless, we do find a significant fraction of detectable systems originate in the high-[α/Fe] disc and bulge, indicating that it is still important to include these components, as was ignored in some earlier works (see Sec. 6).

In the top panel of Fig. 3.4, we show the density distribution for detectable DCOs in the galaxy. We see that most detectable sources are concentrated towards the Galactic centre, with a strong bias towards sources that are on our side of the Milky Way in the vicinity of the solar system (indicated with the $\odot$ symbol). In principle, systems are detectable out to large distances of about 20 kpc and more, although they become increasingly rare, as can be seen from the 95% contour.

The differences between different DCO types can be seen more clearly in the bottom panel of Fig. 3.4 where we show the distribution of the distances, D, from earth to the detectable systems. Each distribution peaks around 8 kpc, which is the distance to the centre of the Milky Way. The distribution for BHBH and BHNS systems follow a very similar shape, favouring the detection of systems with distance < 8 kpc, but with a tail extending out to about 20 kpc.

The distribution for NSNS stands out by being flatter, making them more common nearby and, surprisingly, also at larger distances compared to the stellar density. This may seem counter intuitive as one might naively expect the less massive NSNS systems would not be observable out to larger distances than the more massive BHNS and NSNS systems. To understand the differences we need to consider not only the mass distribution of binaries, but also their eccentricity and frequency distributions. Since, each parameter contributes to the calculation of the SNR (and thus affects the maximum distance at which systems can be

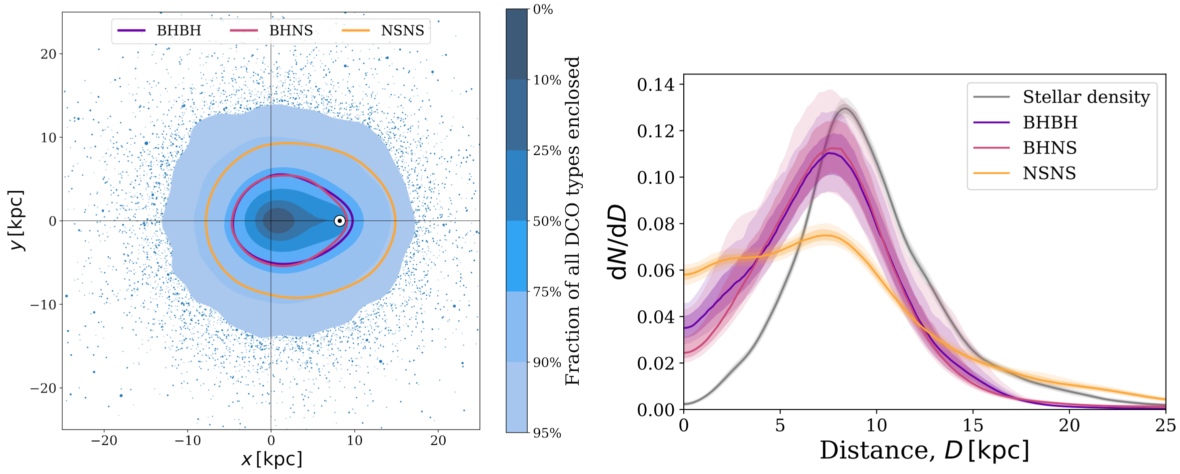

Figure 3.4: Left: A face-on view of the Galactic density distribution for detectable DCOs. We show the density distribution for the top 95% of the sources, the rest are indicated by scatter points whose sizes correspond to their sampling weights. The coloured lines show the 75% contour for each of the individual DCO types. The large cross passes through (0, 0) and helps to highlight the bias towards the position of the sun, which is indicated by the $\odot$. **Right:** As Fig. 3.3, but for the luminosity distance. .

detected).

The reason that the NSNSs are favoured at higher distances is that the NSNS population has the highest fraction of "mildly" eccentric systems ($0.01 < e < 0.03$). In contrast, the BHNS population has a much higher fraction of effectively circular systems ($e < 0.01$), which emit weaker gravitational waves compared to equivalent eccentric systems. Therefore, despite their typically higher masses, the distance at which a BHNS source is detectable is generally lower than for the mildly eccentric NSNS.

Conversely, the BHBH population has the highest fraction of highly eccentric systems ($e > 0.3$). Although one may naively expect that this would result in stronger signals (and so further distances), for a system to have these high eccentricities in LISA, it must still be early in its evolution (otherwise it would have circularised) and thus have a low orbital frequency. The result of this is that highest eccentricity systems tend to have lower SNRs and so cannot be detected at large distances.

Overall we see that the eccentricity distribution of NSNSs occupies a "sweet spot" where the gravitational wave power is increased compared to circular systems, but it isn't too high that the frequency is significantly impacted. This means that NSNSs can be seen out to the largest distances of the three DCO types.

3.4 Formation channels

In our fiducial model, approximately two thirds of detectable BHBHs are formed through the 'only stable mass transfer' channel, whilst the remaining third are primarily formed through the 'classic' CE channel , in which the first phase of mass transfer is stable whilst the second is unstable and results in a common-envelope event. Detectable BHNSs follow an inverse

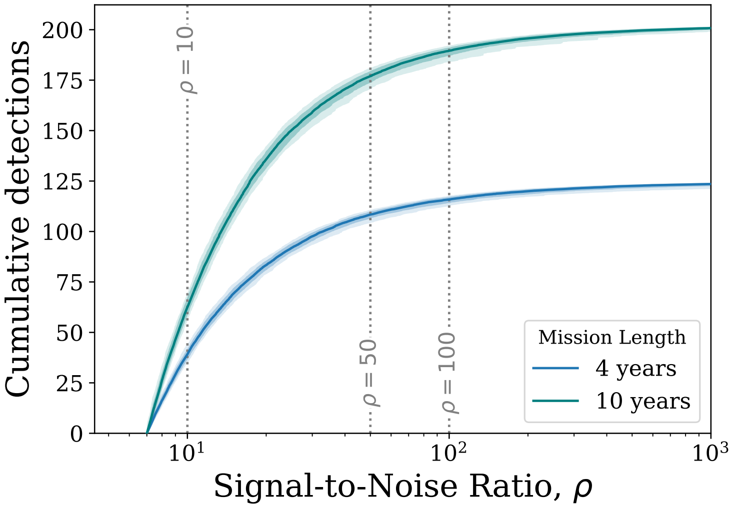

Figure 3.5: Cumulative number of LISA detections with a given signal-to-noise ratio. Colours indicate LISA mission length and shading shows 1- and 2-σ uncertainties (obtained via bootstrapping).

pattern, such that around two thirds are formed through the classic channel and the rest are mainly formed through only stable mass transfer.

In contrast, detectable NSNSs are very rarely formed through only stable mass transfer. Approximately half of systems are formed through the ‘classic’ channel and the rest are formed through a double-core common-envelope event (Brown, 1995) where both progenitors evolve on a similar timescale and initiate a double-core common-envelope event whilst they are on the giant branch. All detectable DCOs show a small fraction of systems are formed through a channel that does not fit into the other categories and hence are labelled ‘other’. These systems tend to be formed through ‘lucky’ supernova kicks that happen to shrink the binary significantly by chance.

The fraction of detectable DCOs that are formed through different formation channels in the other model variations are shown Fig. 3.17, where the first column in the plot corresponds to the fiducial model that we described above.

3.5 How accurately will we be able to infer the parameters of detected systems?

So far we have discussed the properties of the all detectable sources. However, only for a subset of systems we expect to get high enough SNR to obtain accurate and useful measurements of these parameters. Below we discuss the expected SNR distribution and the typical uncertainties expected for the most relevant parameters, namely, the chirp mass and sky localisation.

Signal-to-noise ratio — In Fig. 3.5 we show the cumulative number of detections with a given SNR. Although many a large fraction of sources have SNRs around our assumed detection threshold of 7, many systems are detected with very high SNRs. We find that on average for a 4(10)-year LISA mission, of the 124(202) detections, 85(138), 16(27) and 9(14) systems have SNRs greater than 10, 50 and 100, respectively. These high SNR systems are typically, but not only, the more massive BHBH systems.

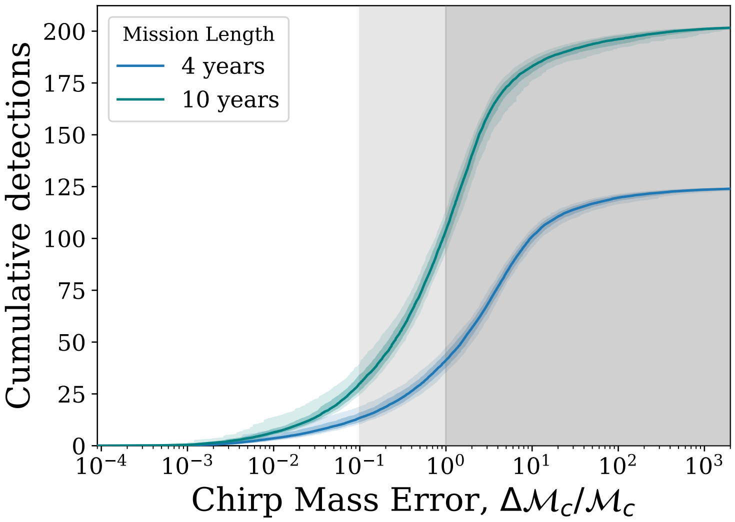

Figure 3.6: As Fig. 3.5, but for the chirp mass uncertainty. The shaded areas indicate regions with more than 10% and 100% uncertainty. 🖼 📖.

Chirp mass — The chirp mass is important for identifying the type of the source of a detected GW signal. The uncertainty of the chirp mass depends on the uncertainty in the measured orbital frequency, the time derivative of the orbital frequency and the eccentricity as detailed in Appendix 8.3.

We find that for a 4(10)-year LISA mission, approximately 41(105) detections have measurable chirp masses ($\Delta\mathcal{M}_c/\mathcal{M}_c < 1$, indicated by the dark shaded region in Fig. 3.6) whilst 13(31) have chirp mass uncertainties below 10% ($\Delta\mathcal{M}_c/\mathcal{M}_c < 0.1$, indicated by the light shaded region in Fig. 3.6). This uncertainty is generally dominated by the uncertainty on the time derivative of the frequency, since most of the binaries are too early in their inspiral for LISA to measure a strong chirp. Note from Fig. 3.6 that increasing the mission length significantly increases the number of detections for which the chirp mass uncertainty is below 100%. The total number of detections only scales as $\sqrt{T_{\rm obs}}$, yet we find that the the total number of detections with a chirp mass uncertainty is below 100% and 10% both scale approximately as $T_{\rm obs}$.

Sky localisation — An accurate sky localisation will be essential to possibly identify electromagnetic counterparts or distinguish sources that come from different components of our Milky Way.

We quantify the sky localisation of a source by estimating the angular resolution for the detectable sources. Since all potential sources are effectively stationary on the timescale of the LISA mission, we can follow Mandel et al. (2018) and use the timing accuracy of LISA and the effective detector baseline to calculate the angular resolution, σ_θ, as

$$\sigma_\theta = 16.6^\circ \left(\frac{7}{\rho}\right)\left(\frac{5\times10^{-4}\,\mathrm{Hz}}{f_{\rm dom}}\right)\left(\frac{2\,\mathrm{AU}}{L}\right), \tag{3.15}$$

where L is the effective detector baseline, which for LISA is 2 AU in the frame of the solar system, since it will orbit the Sun.

We plot the distribution of expected angular resolutions in Fig. 3.7. We see that, for a 4(10)-year LISA mission, approximately 82(123) sources can be resolved to an angular

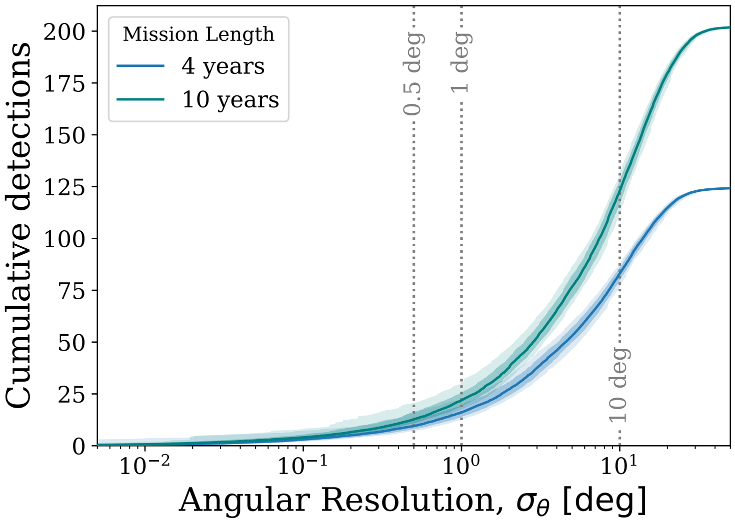

Figure 3.7: As Fig. 3.5, but for the angular resolution. .

resolution better than 10 degrees and only 16(23) better than 1 degree.

For comparison, the size of a pencil beam for a 15 m diameter SKA dish observing at 1.4 GHz is roughly 0.67 square degrees (Smits et al., 2009), corresponding to an angular resolution of $\sigma_\theta = \sqrt{(0.67/\pi)} = 0.46°$ (similar to the angular size of the moon). We will further discuss the prospects of matching LISA detections to radio pulsars with SKA in Sec. 5.2.

4 Results II - Impact of physics assumptions

In this section we explore the effect of varying the uncertain assumptions governing the evolution of binary system and the formation of compact objects. For this we use the population synthesis simulations and model variations presented in Broekgaarden et al. (2021, 2022).

We first discuss the robustness of our predictions for the number of detectable systems (Section 4.1). We then discuss examples of how the observable properties of the detectable systems, such as the distribution of masses and eccentricities, are affected and can potentially be used to probe the physics of double compact object formation (Section 4.2).

4.1 Detection rates

We predict approximately 30-370 detections in a 4-year LISA mission, across all our simulations for varying physics assumptions. This increases to about 50-550 for a 10-year LISA mission. Although the number of detections per type can vary by about 2 orders of magnitude, we find that the total detection rate is fairly robust, among the variations we have considered.

In Fig. 3.8, we show the expected number of LISA detections based on our simulations considering variations in the physical assumptions. We show the expected number of detections for BHBH, BHNS and NSNS systems in the top, middle and bottom panel respectively. In the sections that follow, we briefly explain the variations considered and discuss the most prominent trends.

4.1.1 Efficiency of mass transfer

The efficiency of mass transfer, i.e. the fraction of mass lost by the donor through Roche-lobe overflow that is accreted by the companion, is poorly constrained and is considered as one of the main uncertainties in binary evolution (e.g. de Mink et al., 2007). In our fiducial model A, we use a prescription in which the accretion rate onto stellar companions is regulated by their thermal timescale, i.e. the timescale on which a star can react to changes and restore thermal equilibrium (see e.g. **?**).

In models B-D, we instead adopt a fixed value for the mass transfer efficiency, β, from $\beta = 0.25$ up to 0.75, in cases of stable mass transfer onto a stellar companion. For accretion onto NS and BH we still assume that their accretion is limited to the Eddington rate.

Since nearly all systems that can be detected form through channels where the very first interaction is stable mass transfer, changing β can affect the majority of systems. Generally, higher mass transfer efficiencies lead to higher masses for the accreting stars, but also lead to wider orbits (Soberman et al., 1997; van Son et al., 2020). Changing β thus already affects the masses and orbital separation after the first interaction phase, which in turn changes the starting conditions and outcome of all subsequent interactions phases. This makes it complicated to fully understand the impact of varying β in simple terms, but we can distinguish two main patterns for higher and lower mass systems respectively.

For the most massive progenitors, increasing β leads to secondary stars that are so massive and luminous that they experience strong wind mass loss. This leads to further widening of the orbit. In addition the more massive secondaries may not be able to fully expand as mass loss may prematurely remove their hydrogen envelope. Both of these effects tend to prevent the most massive secondaries from filling their Roche lobe. This means that they cannot initiate the reverse interaction needed to shrink the binary system and eventually produce a detectable double compact object. We indeed see that increasing β leads to a decrease of the expected number of detections for BHBHs and BHNSs, which originate from the most massive progenitors.

For lower mass progenitors, which primarily produce NSNS systems, we find the opposite: increasing β leads to an *increase* in the number of detectable NSNS systems. This is in part because the changes in secondary mass and the orbital widening, affect the number of systems for which the reverse interaction successfully ejects the envelope and shrinks the orbit. Furthermore, the increased mass of the secondary stars allows stars that would have otherwise ended their life as a WD to instead become massive enough to form a NS (e.g. Zapartas et al., 2017). The same effect also allows some NS progenitors to become massive enough to become BH progenitors, which partially cancels the extra progenitors that would have originally been destined to become WDs.

4.1.2 Common-envelope evolution

The common-envelope phase constitutes a highly uncertain phase in the evolution of interacting binary systems (e.g. Ivanova et al., 2013). The uncertainties concerns the conditions required for the onset of a common-envelope phase and, if a common-envelope phase occurs, what the outcome is. Rapid population synthesis simulations such as ours approximate both questions in a crude way. We therefore consider several model variations.

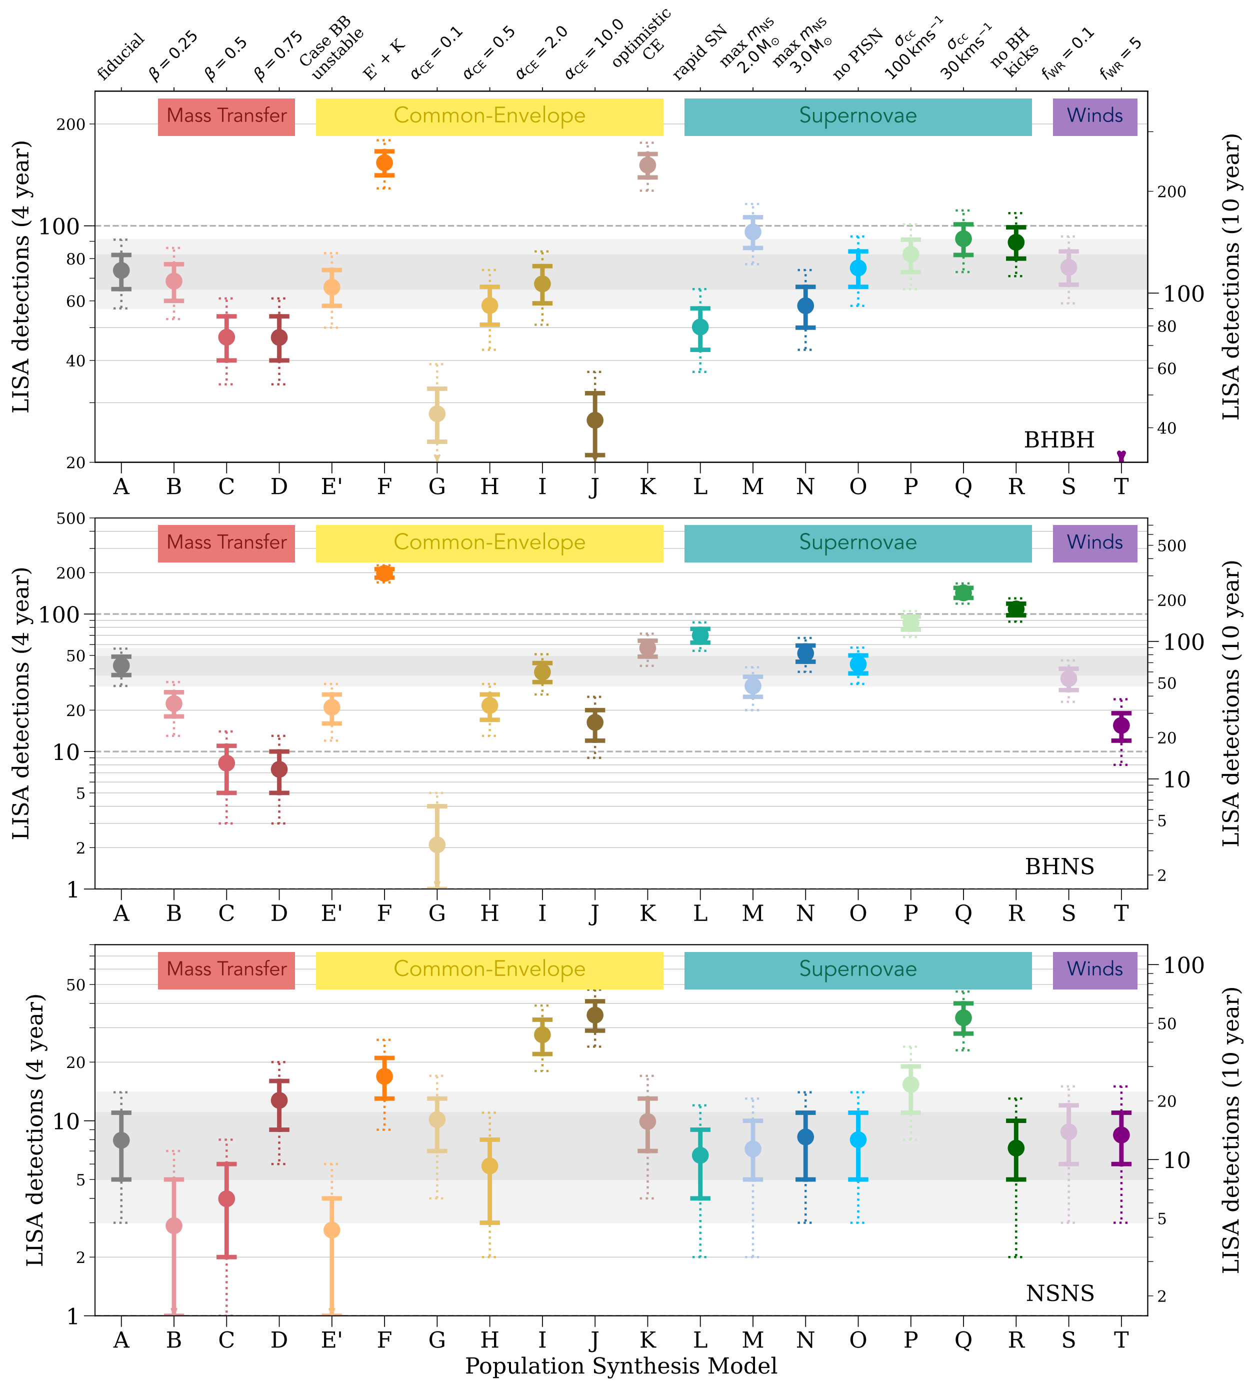

Figure 3.8: The number of expected detections in the LISA mission for different DCO types and model variations. Error bars show the 1- (solid) and 2-σ (dotted) Poisson uncertainties. An arrow indicates that the error bar extends to zero. The left axis and grid lines show the number of detections in a 4-year LISA mission and the right axis shows an approximation of the number of detections in a 10-year mission (we scale the axis by $\sqrt{T_{\rm obs}}$). The grey shaded regions indicate the 1- and 2-σ region for the fiducial model. Each model is described in further detail in Table 3.2 and details of the fiducial assumptions are in Section 8.1.2. See Sec. 4.1 for a discussion.

The efficiency parameter $\alpha_{\rm CE}$ — To estimate the outcome of a CE phase, we use a simple consideration of the binding energy and orbital energy (Webbink, 1984; de Kool, 1990). Our fiducial model assumes a common-envelope efficiency parameter $\alpha_{\rm CE} = 1$ which can be interpreted as the case where all the energy liberated by shrinking the orbit is used in an optimal way to unbind the envelope. There have been many attempts to constrain this parameter using observations and more recently also using 3D simulations (e.g. De Marco et al., 2011; Law-Smith et al., 2020; Lau et al., 2022), but no consistent picture has emerged. Therefore we consider large variations in this parameter.

In models G-J we alter the common-envelope efficiency parameter $\alpha_{\rm CE}$ to $0.1, 0.5, 2.0$ and 10.0 respectively. Values smaller than 1 may represent cases where not all energy is used efficiently to unbind the envelope, for example when part of the energy escapes in the form of radiation or if part is used to impart additional kinetic energy in the ejecta (e.g. Ivanova et al., 2013; Nandez & Ivanova, 2016). Values larger than 1 may represent cases where additional energy sources can be tapped into, such as for example jets powered by accretion (e.g. Schreier et al., 2021). The variations can also be seen as a way to cover uncertainties in estimates for the binding energy itself.

Increasing $\alpha_{\rm CE}$ makes common-envelope ejection more efficient or, in other words, less orbital shrinkage is needed to successfully eject the envelope. This has two consequences. (1) A larger fraction of systems avoids merging during the CE phase. This increases the overall number of DCOs. (2) The systems that survive the CE phase are wider, possibly too wide to become detectable as gravitational wave sources (e.g. Klencki et al., 2021). So, while the first consequence favours the formation of DCOs, the second consequence disfavours the formation of DCOs that are tight enough to be detected.

These two opposing effect result in a fine tuning situation. Only a very small subset of progenitor systems have the right orbital parameters prior to the CE phase to successfully produce detectable systems. Changing $\alpha_{\rm CE}$ moves and changes this window in the parameter space that successfully leads to the formation of detectable systems. How the number of the detectable systems changes depends on whether the relevant parameter space grows or shrinks and on how well the relevant part of the parameters space is populated. Fully unravelling these effects and how they interplay with the assumed star formation history is beyond our scope (and possibly not even of large relevance given the severe simplifications). We will limit the further discussions to simply stating the trends we observe.

We find that the BHBH rate peaks for $\alpha_{\rm CE} = 1$ (model A) and reduces whether we increase *or* decrease $\alpha_{\rm CE}$. The BHNS rate follows the same pattern as BHBHs although the value of $\alpha_{\rm CE}$ which maximises BHNSs seems to be between 1 and 2. In contrast, for NSNSs we find that increasing $\alpha_{\rm CE}$ (models I and J) results in significantly higher rates.

The "optimistic" CE treatment — We further explore a model variation introduced by Belczynski et al. (2007) often referred to as the "optimistic" CE scenario. This variation (model K) relaxes our restriction that donor stars that are on the Hertzsprung cannot survive common-envelope events.

In agreement with other studies, we find that this treatment leads to a significant increase in the formation rate of BHBHs, by a factor of two. This is because the progenitors expand significantly during the Hertzsprung gap phase in our simulations. In our fiducial simulation, all progenitors that initiate unstable interaction during this phase would end as stellar mergers, while in this variation they will survive. The progenitors of BHNSs and NSNSs are less strongly affected with an increase about 30% increase.

Case BB mass transfer — In models E′ and F we consider uncertainties in case BB mass transfer. This is a phase of mass transfer where the donor star has already lost its hydrogen envelope in a prior interaction, but fills its Roche lobe again as it expands during helium shell burning phase (e.g. Dewi et al., 2002; Tauris et al., 2015, 2017). This is of particular interest for the formation of NSs, as their lower mass progenitors are swelling the most during this phase (e.g. Laplace et al., 2020, and references therein). Population synthesis studies find that nearly all NSNS systems form through a phase of case BB mass transfer (Vigna-Gómez et al., 2018).

In model E′ we enforce that case BB mass transfer is always unstable, such that it always leads to a CE. Note that this is slightly different from model E described in Broekgaarden et al. (2021, 2022). In their work the pessimistic approach to CE evolution is implemented such that all HG stars are excluded including helium stars in the helium shell burning phase. This, in combination with the assumption that case BB mass transfer is always unstable, effectively leads to the exclusion of all systems that originate through this channel. We are interested in NSNS systems, which are frequently formed through this channel in our simulations. Therefore, we adapted this model to only exclude H-rich HG donors, but allow systems with donors that are helium stars in the helium shell burning phase to survive a CE phase.

As expected, in model E′, we find that case BB systems form through a common-envelope phase rather than only stable mass transfer (see Fig. 3.17). This model gives the lowest NSNS rate of all of our variations, which is a factor of 3 lower than our fiducial rate. We also find a reduction of the BHNS systems by a factor 2. The BHBH systems are not significantly affected, as expected, since case BB mass transfer does not play a role for high mass progenitors.

Finally, in model F, we again enforce that case BB mass transfer is always unstable, but in combination with the optimistic treatment for CE (essentially combining models E′ and K). This allows the systems that have HG donors for common-envelopes (as well as those formed through case BB mass transfer) to survive the CE phase. We find that this model leads to the highest predictions for the detections among all variations that we have considered.

4.1.3 Supernovae and compact remnants

The formation of compact remnants and their associated natal kicks also constitute important uncertainties.

In model L we consider the so-called “rapid” remnant mass function by Fryer et al. (2015) as an alternative to the “delayed” description used in our fiducial simulations. This affects the mass distribution (see Sec. 4.2.1), but the effects on the number of detectable systems is modest for BHBH and BHNS, and negligible for NSNS.

The same is true for the impact of changing the assumed maximum neutron star mass, $m_{\mathrm{NS,max}}$ (M and N). Lowering $m_{\mathrm{NS,max}}$ increases the number of detections involving BHs (since more stars form BHs instead of NSs) and vice versa, but has no significant effect on the number of NSNS detections since the vast majority are formed from low mass NSs.

We find that not implementing pair-instability supernovae (PISN) or pulsational pair-instability supernovae (PPISN) in model O has no effect on the number of detections with LISA. This is because the average metallicity of the Milky Way is high enough such that no progenitor retains enough mass to initiate a PISN or PPISN.

Decreasing the natal kicks for all core-collapse supernovae (models P-Q) increases the detection rates for each DCO type, since lower kicks result in fewer disrupted binaries and hence a more numerous detectable population. The BHNS and NSNS systems are strongly affected whilst the impact on BHBH systems is insignificant. The reason that BHBHs are relatively unaffected is that, in our models, the natal kicks for BHs are scaled down with the amount of mass that falls back. In the case of BHBHs, the black holes have very massive cores and thus low kicks.

In model R we assume BHs form without any kick, while using our fiducial assumption for the natal kicks of neutron stars. This increases the predictions for BHNS by a factor 3 but the impact on BHBH systems is much smaller for the same reason as for models P-Q. As expected, the NSNS population is not affected.

4.1.4 Stellar winds

Mass loss in the form of stellar winds or eruptions is also a main uncertainty. It affects by how much stars can grow in size and it affects the final core masses. We consider variations in the mass loss by naked helium (Wolf-Rayet like) stars and choose to vary the efficiency of these winds between 0.1 and 5, to account for uncertainties in the derived rates (e.g. Vink, 2017; Shenar et al., 2019; Hamann et al., 2019; Sander & Vink, 2020).

We find that a reduction of the wind mass loss has very little effect on our predictions. This is the consequence of several effects that cancel each other. Firstly, the decreased Wolf-Rayet like winds mean that the DCOs (particularly those containing BHs) are generally more massive and so more detectable in LISA. Secondly, one may expect that LISA sources in model S would be higher frequency than in our fiducial model as decreased winds generally result in tighter binaries. However, though this is the case at DCO formation, we find that by the time the sources have evolved until they are observable by the LISA mission, they have lower orbital frequencies than in our fiducial model. This is because the reduced winds allow DCOs to be formed at higher metallicity and, therefore, at more recent times. This means that most DCOs do not evolve for as long before the LISA mission and so remain at lower frequencies (wider separations) thus making them less detectable.

In addition, we find that NSNSs are more eccentric and BHBHs are less eccentric than our fiducial model (with BHNSs relatively unchanged). The increase in eccentricity for NSNSs comes from the same reason as the lower frequency, more recent birth times mean that binaries have less time to circularise. The same is not true for BHBHs as the more massive systems are less affected by supernova kicks and so fewer high eccentricity systems are formed. Overall, despite the large differences in the system properties, these three effects in combination leave the detection rates relatively unchanged.

In model T we instead *increase* the efficiency of Wolf-Rayet winds by a factor of 5. In this model the detection rate of BHBHs decreases by over a factor of 10 and BHNSs by over a factor of 2 whilst the NSNS rate is relatively unchanged. Increasing the efficiency of WR winds widens the orbit and decreases the final masses of DCOs. This means that some progenitors that would have formed LISA sources under our fiducial assumptions would not have enough mass to produce a DCO, or produce a NS instead of a BH. The effect is strongest for the progenitors of black holes, since these are most strongly affected by Wolf-Rayet winds. This effect is less pronounced in NSs since the rate of mass loss for the progenitors of these systems is low enough that changing by a factor of 5 still does not impact the final fate. Moreover, DCOs that are formed tend to be less massive and therefore less detectable.

4.2 Properties of detectable systems

In this section, we consider how varying underlying physics assumptions changes the properties of detectable systems. We focus on several key differences across physics variations rather than showing the differences in every model and thus this section is by no means exhaustive.

4.2.1 Effect of remnant mass prescription on lower mass gap

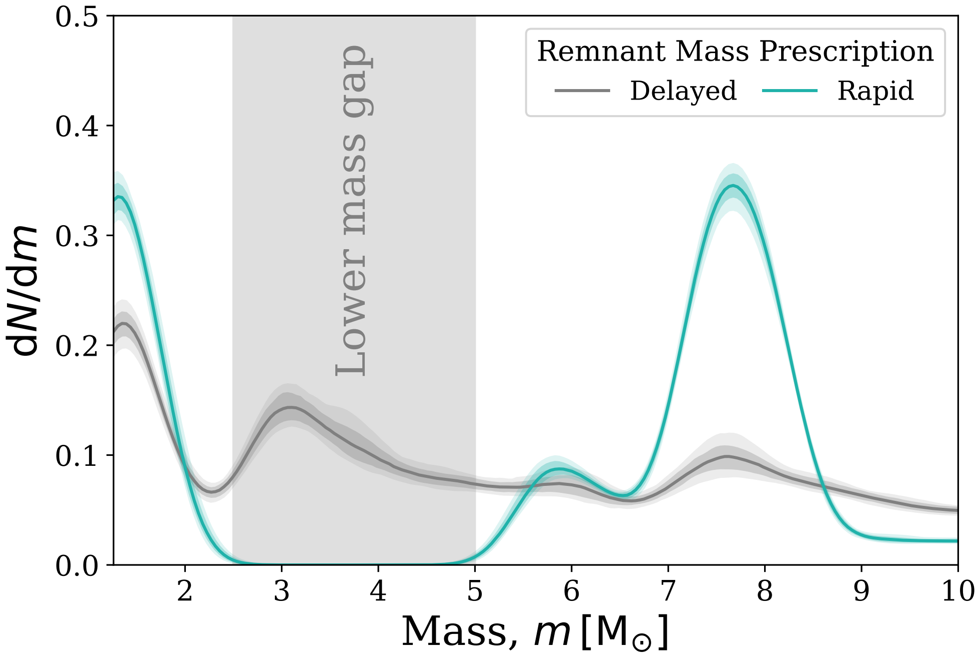

Figure 3.9: Comparison of the component mass distribution of LISA detectable DCOs when using the Fryer delayed (model A) and rapid (model L) remnant mass prescriptions. Distributions are plotted in the same way as Fig. 3.3, except all DCOs types are shown in one curve and each type is weighted by its detection rate in the respective model.

In Fig. 3.9, we show the component mass distribution for all LISA detectable DCOs (BHBHs, BHNSs and NSNSs) for two different remnant mass prescriptions. The grey distribution uses the Fryer *delayed* remnant mass prescription (Fryer et al., 2012), which is our fiducial assumption (model A). This prescription produces compact objects in the lower mass gap ($2.5\,\mathrm{M}_\odot \leq m \leq 5\,\mathrm{M}_\odot$) and indeed we find that, of the LISA detectable DCOs, approximately 69% of BHBHs, 39% of BHNSs and 0% of NSNSs have at least one component in the lower mass gap. Overall, weighting by the relative detection rates, this gives that, in our fiducial model, 55% of our predicted LISA DCO detections would have at least one component in the lower mass gap when using this remnant mass prescription. This equates to approximately 69 systems being detected in the lower mass gap. Alternatively, the blue curve in Fig. 3.9 shows the same distribution but for the *rapid* remnant mass prescription (Fryer et al., 2012), which we use in model L. In this case, no compact objects are formed (and therefore, detected) in the lower mass gap.

From the stark difference between these models, it is clear that it is difficult at this point to say with any certainty what fraction of systems LISA will detect in the lower mass gap given the highly uncertain formation rate of systems in this mass range.

We find that the percentage of detectable systems with at least one component in the lower mass gap varies between approximately 30-70% (or, in terms of detections, from 15 to 156) for the different model variations that we consider, except for model L, for which it is 0%. We note however that the masses of detectable systems are not always well constrained (see Section 3.5). It may therefore be difficult to assess whether a detected system truly has components in the mass gap.

Lastly, we highlight that recent investigations of the GW detections so far suggest a significant reduction of BHs in the lower mass gap (e.g. Farah et al., 2022; Ye & Fishbach, 2022, van Son et al. in prep.). All our models except model L allow for systems in the mass gap, which may be at tension with the observational constraints so far.

4.2.2 Effect of natal kicks on eccentricity distribution

In Fig. 3.10, we investigate how decreasing the magnitude of natal kicks from core-collapse supernovae affects the eccentricity distribution of LISA detectable DCOs. For reference, we show the mean fiducial distributions (model A) as dashed lines (see Fig. 3.3d for full comparison). In the main curves, we reduce the velocity dispersion for core-collapse supernovae from $265\,\mathrm{km\,s^{-1}}$ to $30\,\mathrm{km\,s^{-1}}$ (model Q).

We find that the LISA detectable BHBHs are significantly less eccentric with weaker kicks, such that the population above $e = 0.2$ is nearly completely eliminated. This is because BHBHs are often massive enough to withstand strong natal kicks without disrupting and these kicks tend to impart significant eccentricity. In model Q, very few systems are ever given such strong kicks and thus very few BHBHs are detected with significant eccentricity.

Since BHNSs are less massive than BHBHs and have more unequal mass ratios, they are more vulnerable to disruption during supernova kicks. BHNSs can only withstand strong kicks when they are aimed in the correct direction and so only a small ‘lucky’ fraction of the fiducial population is highly eccentric. Therefore in model Q, although we see that the population of highly eccentric BHNS systems is eliminated (similar to BHBHs), the peak of the distribution actually shifts to *higher* eccentricity. This is because a larger fraction of systems are given weaker kicks that BHNSs can withstand and these impart much more moderate eccentricities.

Finally, we find that the NSNS distribution is relatively unchanged between model A and Q. This is not surprising however since the majority of NSNSs are formed through electron-capture supernovae and ultra-stripped supernovae and for these types of supernovae we use

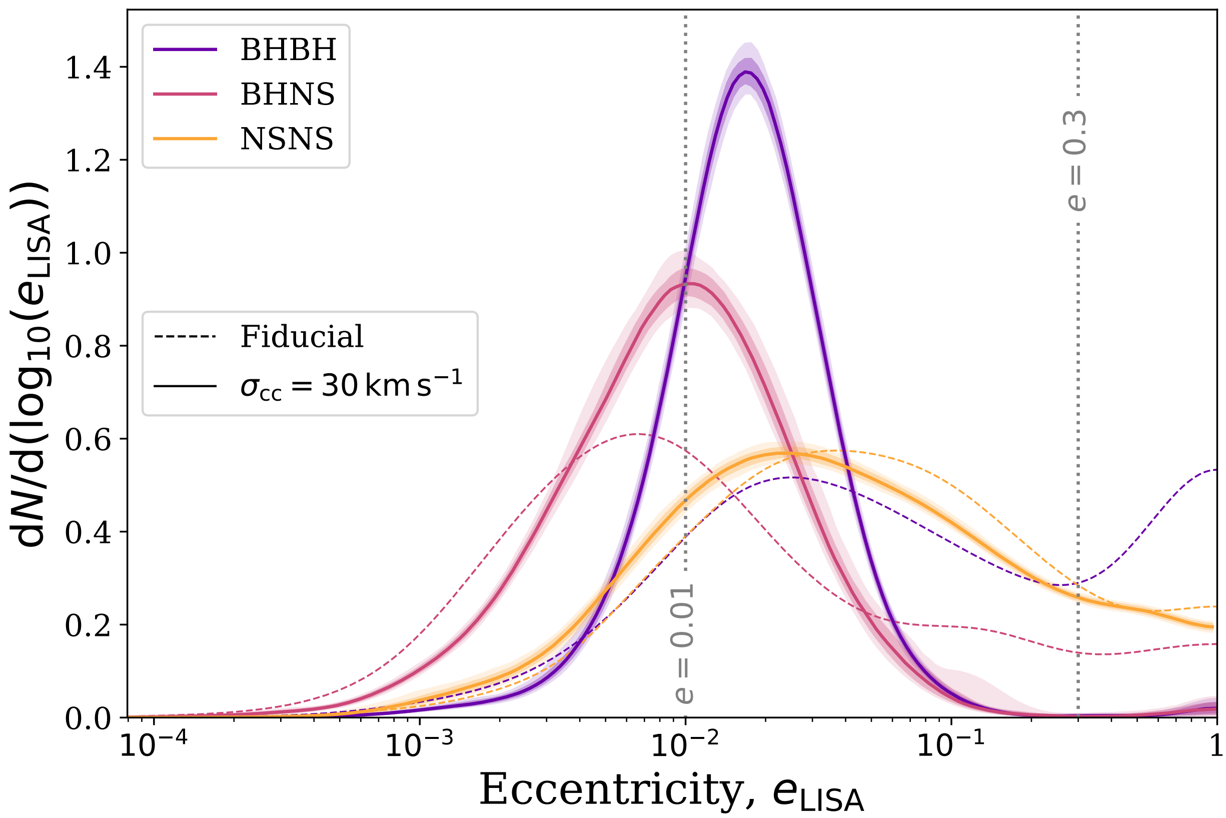

Figure 3.10: As Fig. 3.3d, but for model Q. For comparison, we show the mean distribution for the fiducial model (model A) as dashed lines. .

$\sigma^{\rm 1D}_{\rm rms} = 30\,{\rm km\,s^{-1}}$ already (see App. 8.1.2) and thus there is little difference between the models.

Overall, compared to our fiducial model, we find that decreasing supernova natal kicks, though it strongly increases the number of detections (see Fig. 3.8), strongly *decreases* the fraction of highly eccentric systems that are detected.

4.2.3 Effect of Wolf-Rayet winds on mass distribution

In Fig. 3.11 we show the effect that changing the efficiency of Wolf-Rayet winds has on the individual component mass distribution. Decreasing the Wolf-Rayet wind efficiency allows the formation of more massive DCOs in the Milky Way and, indeed, we see that the distribution extends to $25\,{\rm M}_{\odot}$ and relatively fewer detectable systems are formed at low masses. By contrast, increasing the Wolf-Rayet efficiency by a factor of 5 strongly disfavours the formation of systems at high masses and approximately 85% of detectable systems have masses below $5\,{\rm M}_{\odot}$. These three distributions are very distinct and so it is possible that the mass distribution of LISA could help to constrain the efficiency of Wolf-Rayet winds.

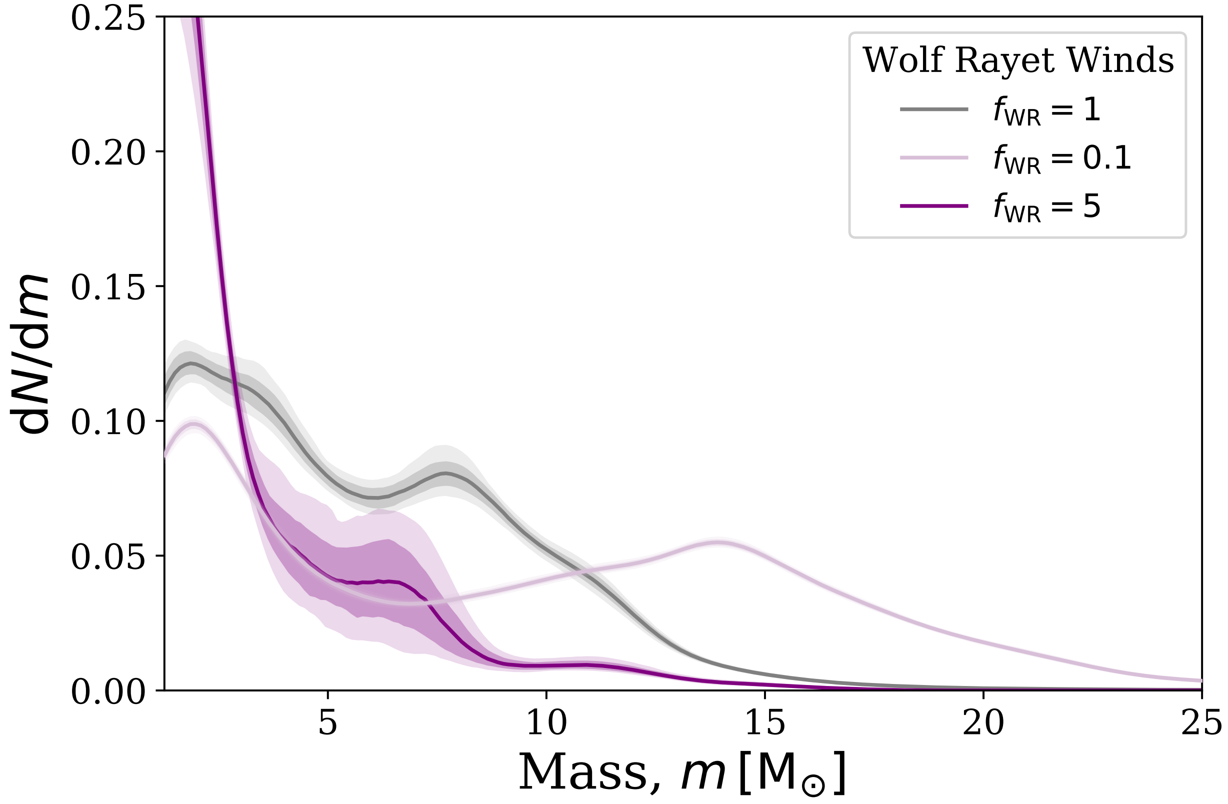

Figure 3.11: As Fig. 3.9, but instead varying Wolf-Rayet wind efficiency (models S and T). The curve for $f_{\rm WR} = 5$ has much higher uncertainties as there are many fewer DCO systems formed in this model (the inverse reasoning also explains the lower uncertainties for $f_{\rm WR} = 0.1$). 🖻 📘.

5 Discussion

In this section we discuss the prospects of (and methods for) identifying LISA sources (Sec. 5.1), the possibility of matching LISA signals to SKA detections (Sec. 5.2), the main caveats for this study (Sec. 5.3) and the possible contribution from other formation channels (Sec. 5.4). All predictions quoted in each subsection are derived for the fiducial model (model A).

5.1 Identification of GW sources

It is important to note that, though we present predictions for the detection rates of specific DCO types, the nature of the source may not be immediately (or ever) apparent from the gravitational wave signal. LISA can detect a variety of sources, from exoplanets (e.g. Tamanini & Danielski, 2019) to common-envelopes (e.g. Ginat et al., 2020; Renzo & Götberg, 2021) that may cause confusion. However, by far the most prominent will be the population of Galactic WDWDs detectable with LISA, which will be several orders of magnitude larger than the population of the more massive DCOs that we focus on in this paper (e.g. Korol et al., 2017). It is therefore imperative that we consider how to distinguish NS and BH binaries from this much more numerous population of sources. In addition to distinguishing them from WDWDs, we must consider how to discriminate between BHBHs, BHNSs and NSNSs themselves.

5.1.1 Distinguishing from WDWD population

The simplest way to check whether a source is a WDWD is to evaluate its chirp mass. The mass of a non-rotating white dwarf cannot be larger than the Chandrasekhar limit of $1.4\,\mathrm{M}_\odot$ (Chandrasekhar, 1931; Hamada & Salpeter, 1961), so we can take the maximum chirp mass of a WDWD to be $\sim 1.2\,\mathrm{M}_\odot$. Therefore, any DCO with a chirp mass that satisfies $\mathcal{M}_c > 1.2\,\mathrm{M}_\odot + \Delta\mathcal{M}_c$ must not be a WDWD (where $\Delta\mathcal{M}_c$ is the error on the chirp mass, estimated using Eq. 3.29). We find that for the detectable population of a 4(10)-year LISA mission, 24(38)% of BHBHs, 28(41)% of BHNSs and 4(5)% of NSNSs satisfy this condition. This method is not particularly effective for NSNSs since their average chirp mass, $1.17\,\mathrm{M}_\odot$, is below the Chandrasekhar limit.

Another discriminator between WDWDs and other DCOs is eccentricity. WDWDs formed in the disc are thought to be formed mainly through isolated binary formation and have little to no eccentricity (e.g. Nelemans et al. 2001, see however Dosopoulou & Kalogera 2016a,b; Gosnell et al. 2019). This is because WDWDs formed through isolated binary evolution all experience a phase of mass transfer or a common envelope, which typically circularises the binary (e.g. Marsh et al., 2004). However, in contrast to the more massive DCOs that we study, WDWDs do not experience strong natal kicks which we find to be the main source of eccentricity. Therefore, if any system is detected with anything other than one detectable harmonic, this suggests that the system is unlikely to be a WDWD. We find that for a 4(10)-year LISA mission, 55(61)% of BHBHs, 27(29)% of BHNSs and 66(68)% of NSNSs are detected with multiple harmonics (see also Sec. 3.2). Both the absolute percentage and the relative improvement with an extended LISA mission is lower for the BHNSs with respect to other DCOs as we find that these BHNSs are less eccentric on average (see Fig. 3.3d and discussion in Sec. 3.2).

However, we should also consider that eccentric WDWDs could be formed through dynamical formation in Milky Way globular clusters (e.g. Willems et al., 2007; Kremer et al., 2018), or with third companions (e.g. Antonini et al., 2017). This means that we cannot assume that eccentric binaries are not WDWDs unless they are detected in the Galactic plane (though even then there is a chance they were formed dynamically). We can use the sky localisation, scale height of the disc and distance to the source to estimate what fraction of eccentric sources can be localised to the Galactic plane. This condition can be written as $\sigma_\theta < \arcsin(z_{\mathrm{plane}}/D_L)$ or $D_L < z_{\mathrm{plane}}$, where we set the height of the Galactic plane, $z_{\mathrm{plane}} = 0.95\,\mathrm{kpc}$, to the scale height of the high-α disc. We apply this condition to find that the fraction of sources that are eccentric and localised within the disc for a 4(10)-year LISA mission are 40(40)% for BHBHs, 23(23)% for BHNSs and 59(59)% for NSNSs. Note that although the fractions are the same for the 10-year mission, the absolute number of detections is still greater.

Overall, combining these methods (chirp mass, eccentricity and sky localisation) we find that for a 4(10)-year mission, LISA will detect at least 37(70) BHBHs, 18(38) BHNSs and

5(8) NSNSs that are distinguishable from the WDWD population. Thus we will be able to confidently distinguish approximately half of all detected sources from WDWDs. This increases to roughly 60% for a 10-year mission. We highlight that, though the overall number of LISA detections in an extended mission only increases by a factor of $\sqrt{T_{\rm obs}}$, the number of distinguishable detections increases by a greater factor since each of the more numerous sources are better measured. This further underlines the benefits of extending the LISA mission to 10 years.

5.1.2 Discriminating between BHBHs, BHNSs and NSNSs

The problem of discriminating between the BHBH, BHNS and NSNS populations can be more difficult than distinguishing them from WDWDs. For NSNSs, we can follow a similar method to the WDWDs (see Sec. 5.1.1) by applying our knowledge of the maximum mass of a neutron star. Following our fiducial assumption, we can take the maximum mass of a neutron star as $2.5\,\mathrm{M}_\odot$ and thus the maximum chirp mass that a system can attain without one of the components being a black hole is $\mathcal{M}_c = 2.2\,\mathrm{M}_\odot$. For a 4(10)-year LISA mission, the fraction of systems that are above or below this limit (and thus *must* respectively contain or not contain a BH component) by more than $\Delta\mathcal{M}_c$ is 21(33)% for BHBHs, 18(24)% for BHNSs and 47(62)% of NSNSs, which in terms of absolute detections is 16(39) for BHBHs, 7(17) for BHNSs and 4(8) for NSNSs.

For separating the BHBH and BHNS population one could do so probabilistically given the properties that are measured, particularly the orbital frequency and eccentricity, since these distributions are fairly different for the two DCO types (see Fig. 3.3). This method would pose a challenge, however, as it would likely only indicate which type was more likely rather than discriminate between them with strong evidence.

Another possible solution would be the existence of electromagnetic counterparts to the gravitational wave signal. In Section 5.2 we consider the possibility of detecting a pulsar within a BHNS or NSNS system. This could be used to identify the type of the source.

5.2 Matching LISA detections to pulsars with the SKA

Since the vast majority of the LISA detectable population of DCOs will not merge for many years, the main type of electromagnetic counterpart for this population is pulsars. Therefore, for this section we focus only on BHNSs and NSNSs since no BHBH system will contain a pulsar. The joint detection of a binary pulsar with LISA and the Square Kilometre Array (SKA, Dewdney et al. 2009) would not only help to constrain the parameters of the binary, but also enable investigation of other compact object physics. A pulsar(PSR)+BH can provide stringent tests of theories of gravity, in particular the "No-hair theorem" (Keane et al., 2015). Alternatively, an ultrarelativstic PSR+NS system could be used to measure the neutron star equation of state up to an order of magnitude more accurately than other proposed observational constraints (Kyutoku et al., 2019; Thrane et al., 2020).

We estimate that on average, given the number of detectable pulsars and the SKA sky area, each pulsar in the SKA occupies a region with an angular resolution of $\sigma_\theta < 1.3^\circ$ or 0.7° for SKA-1 and SKA-2 respectively (see Appendix 8.5). Therefore, any DCOs containing NSs localised by LISA with an angular resolution lower than these values can be unambiguously matched to the radio signal in the SKA. By considering Fig. 3.7, approximately 11 and 6 (for SKA-1 and SKA-2) DCOs will satisfy this constraint.

If there is more than one pulsar in the region given by the LISA sky localisation, one can compare the measured parameters of the system in LISA and the SKA. Both the SKA and LISA will measure the orbital frequency to high precision, as well as the time derivative of the frequency and chirp mass to a lesser precision, of each of these systems. Therefore, one could perform a targeted search with the SKA that checks the sky location given by LISA, only looking for binary pulsars with orbital frequencies within the uncertainties. If there was *still* more than one possible pulsar one could also check against the chirp mass. In this way, we expect it will be possible to get a joint detection between the SKA and LISA even when the sky area implied by the LISA detection contains more than one pulsar.

In order to assess the efficacy of this method, we would need to know the probability that two random binary pulsars would have orbital frequencies and chirp masses close enough that one could not tell which pulsar matches the LISA detection. This would require simulating the SKA population of pulsars with a code such as PSRPOPPy (Bates et al., 2014) to find the frequency and chirp mass distribution, which is beyond the scope of this paper. However, the uncertainty in the orbital frequency of a binary on the detection threshold (SNR = 7) for a 4-year LISA mission is 2.5×10^{-9} Hz and 1.0×10^{-9} Hz for a 10-year mission (calculated using Eq. 3.30). Therefore, we expect that the SKA could likely isolate the correct binary pulsar to match to a LISA detection even when several are present in the sky localisation region.

5.3 Caveats

Our predictions are subject to various uncertainties which can be broadly divided into two different categories: those arising from the progenitor models for the population of DCOs and those arising from the choices we have made when placing these DCOs in our model for the Milky-Way. Although we are unable, at present, to evaluate the impact of all these uncertainties, the reader should nevertheless keep in mind that they are likely very substantial. Most of these concerns are not unique to these study, but apply to most of the predictions available in present literature. We highlight a few main concerns.

Progenitor models — Our binary-star progenitors models have been computed with a rapid population synthesis code (see Sec. 2.1). This code relies on approximate parametric prescriptions for the stellar evolutionary tracks of single tracks and simple algorithms to mimic the effects of evolutionary and binary interaction processes. Even though we explicitly consider the impact of some of the main physics uncertainties (see Sect. 8.1.3) the list of variations that we considered is far from exhaustive. Moreover, it is by no means guaranteed that the parametric prescriptions used in this code lead to realistic results, even when varying the values of the parameters to their extremes. We stress in particular the uncertainties affecting our most massive progenitor models. Observational constraints are scarce for high mass stars and practically non-existent for the rapid evolutionary phases (e.g. Langer, 2012; Mapelli, 2021). This is even more true for the evolution of massive stars at low metallicity. In addition to our limited understanding of massive stars, we note that the rapid population synthesis code, such as the one employed to compute the models used in this study, rely on extrapolations of the original fitting formulae to approximate the evolutionary tracks for these higher mass progenitors (Hurley et al., 2000, 2002). A further caveat is whether the population synthesis predictions used in our study realistically describe the population of compact objects. Broekgaarden et al. (2022) shows that none of the physics variations that we use can be excluded at present based on the overall constraints on the GW rates, but this does not mean that they are accurate. For example, at present it is unclear whether these models reproduce the features in the inferred mass distribution (van Son et al. in prep. Farah et al., 2022; Abbott et al., 2021c; Li et al., 2021; Veske et al., 2021; Tiwari, 2022; Edelman et al., 2022).

Populating the Milky Way — Our Milky Way model is semi-empirical and has been calibrated based on observations. Unfortunately, the early evolution of the (metallicity dependence of the) star formation history is poorly constrained. We do not expect this to be a major concern, as most of the double compact objects have relatively short delay times of less than 2 Gyr (see Fig. 3.3e), but this is a caveat that should be kept in mind. Furthermore, to estimate the rate of detectable systems, we rely on normalisation choices (e.g. how many detectable double compacts are formed per unit of star formation). This depends heavily on the initial mass function, as low mass stars account for most of the mass while high mass stars are the progenitors of double compact objects. Further choices, such as the binary fraction and the initial distributions of binary parameters also play a lesser but probably still significant role de Mink & Belczynski (e.g. 2015); Chruslinska et al. (e.g. 2017); Klencki et al. (e.g. 2018).

We also note that, for reasons of computational efficiency, we have not accounted for the spatial velocities resulting from the Blaauw-Boersma kick (Blaauw, 1961; **?**). In test simulations we find that accounting for this spreads out the population (increasing the typical height above the Galactic plane and Galactocentric radius), but we find that the impact on the rate is limited. In light of the other much larger uncertainties, we felt that this was justified (see however, e.g., Brandt & Podsiadlowski 1995; Abbott et al. 2017b). We have further ignored a possible contribution coming from the Galactic halo, as Sesana et al. (2020) estimates this is not significant compared to the contributions from the Galactic bulge and discs. However, this may not be true for other formation channels other than those we have considered here.

Correspondence with ground-based detections — The models used in this paper were shown to be consistent with the inferred merger rates from the GWTC-3 catalogue available at the time of submission (for example the BHBH rate was quoted as $16 - 130\,\mathrm{Gpc}^{-3}\,\mathrm{yr}^{-1}$ in section IV.A of version 1 and 2 on the arXiv in Nov 2021 **?**). We note, however, that the upper limit on the rate has come down by a factor two since, after the completion and submission of our work. The inferred BHBH merger rate is now quoted as $16 - 61\,\mathrm{Gpc}^{-3}\,\mathrm{yr}^{-1}$ in the final version of the GWTC-3 paper that appeared in Feb 2022 (**?**). This means that some model variations may over predict the current rates.

5.4 Other formation channels

In this paper we considered the formation of NS and BH binaries formed via isolated binary evolution, through the classical CE channel, the stable mass transfer channel and variations on these (see Fig. 3.17). We did not consider further possible contributions from other formation channels, which may play a role.

We highlight the possible role of dynamical formation in globular clusters. Kremer et al. (2018) predict, for a nominal 4-year LISA mission, that 21 sources will have SNR > 7, of which 7 are BHBHs, 0 are BHNSs and 1 is a NSNS (see Table 1 Kremer et al., 2018). This is significantly lower than the rates we predict for nearly every model variation. If true, this would mean that formation through isolated binary formation will dominate the LISA detections.

Banerjee (2020) investigates formation of LISA detectable BHBHs in young massive and open stellar clusters and estimates approximately 128 BHBHs with SNR > 5 in a 5-year LISA mission (see Table 1, Column 9 Banerjee, 2020). Although this is similar to the number we predict for our fiducial model, we note these authors adopt a threshold SNR required for a detection that is lower and a mission length is slightly longer than what is typically assumed (i.e. SNR > 7 and 4 years, as we have also adopted in our work). We expect that, after correcting for this and making a fair comparison, our fiducial model predicts about twice as many detections.

The contribution of triples systems (e.g. Antonini et al., 2017), or even higher order multiple systems (e.g. Vynatheya & Hamers, 2022) will likely also be of interest, in particular for the formation of eccentric sources. We are, however, not aware of specific predictions for the detection rates that we can compare to directly.

6 Comparison with previous studies

In Fig. 3.12, we compare our results to similar previous studies that investigate the population of stellar-mass BHBHs, BHNSs and NSNSs that are detectable with LISA. Fig. 3.12 details the expected detection rates predicted by each paper as well as their assumptions regarding their Milky Way galaxy model, binary population synthesis simulation and LISA mission specifications. We only include papers that are similar to our work, such that they use population synthesis and simulate sources in the Galactic plane. Moreover, Fig. 3.12 does not include the numerous papers on the LISA WDWD population as we do not make predictions for WDWDs.

Nelemans et al. (2001)—were the first to investigate the population of LISA detectable stellar-mass double compact objects. We find a significantly higher detection rate for BHBHs and BHNSs, as well as a slightly lower rate for NSNSs. We can understand this difference from changes both to the specifications of LISA (such as the mission length and SNR threshold for detection) and our understanding of massive star evolution since the publication of their paper, which both strongly affect the expected detections rates.

Belczynski et al. (2010)—built upon the work of Nelemans et al. (2001), by using a different population synthesis code with two model variations and a multi-component model for the Milky Way. They find a much lower detection rate for BHNSs and NSNSs (and agreed on zero BHBHs) when compared to Nelemans et al. (2001). They state that this discrepancy from Nelemans et al. (2001) comes from differences in their population synthesis and an overall lower formation rate rather than any changes to LISA detectability. The low total

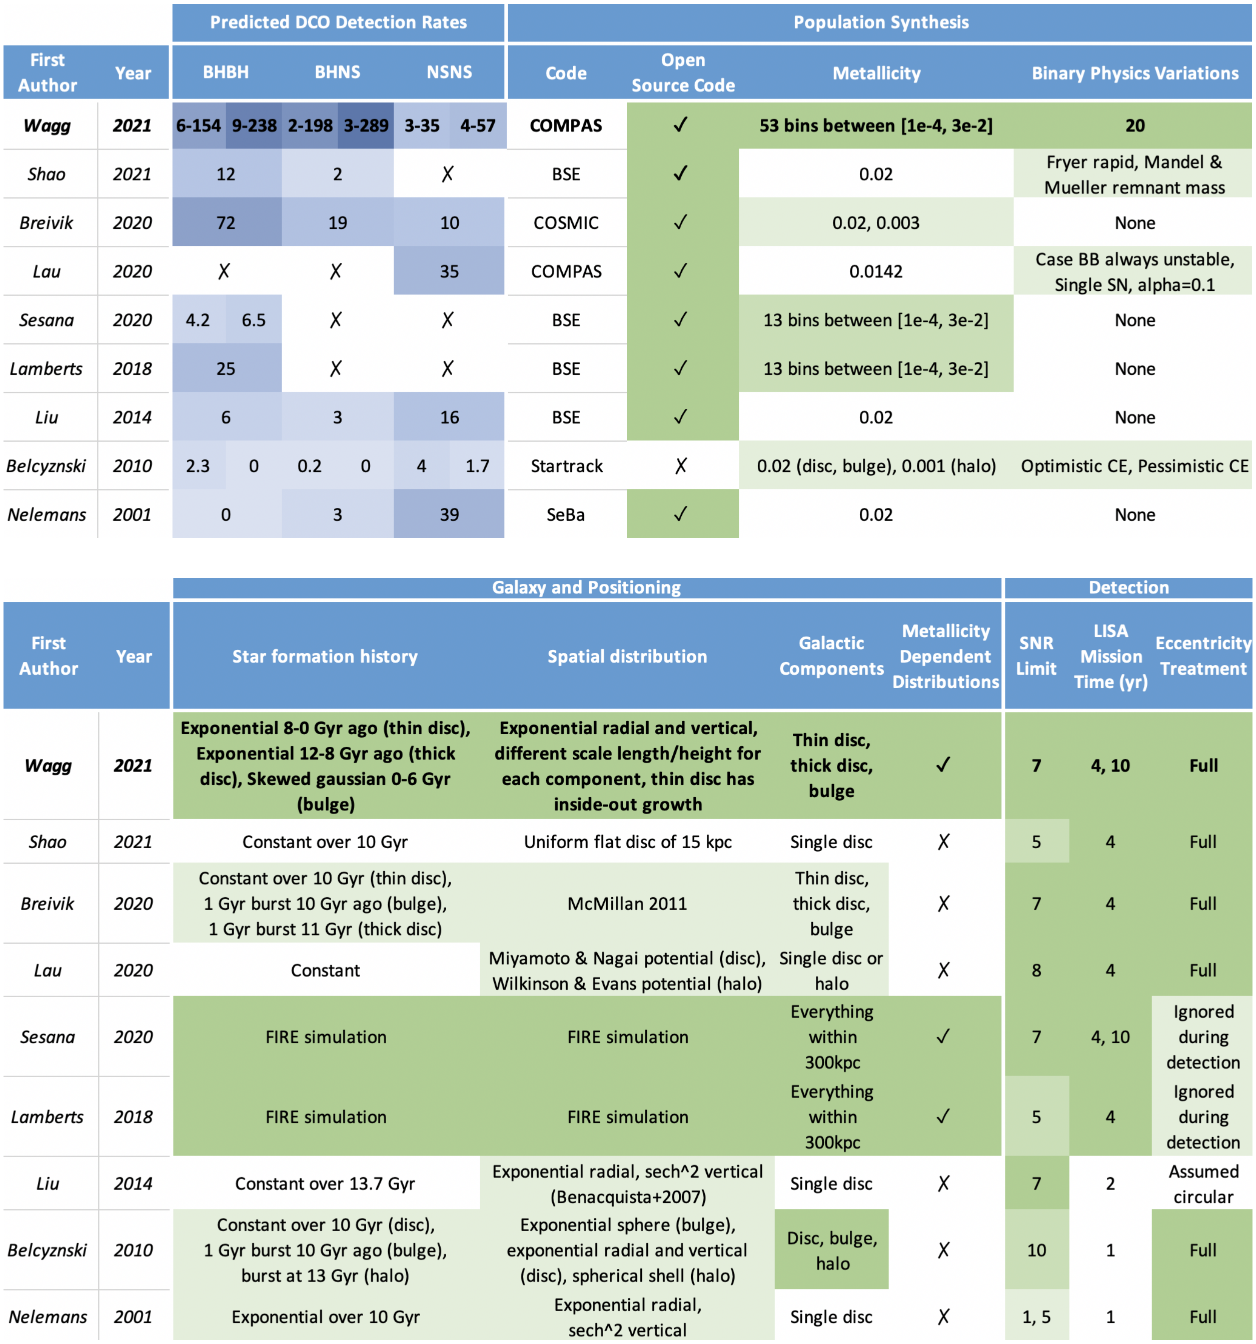

		Predicted DCO Detection Rates						Population Synthesis			
First Author	Year	BHBH		BHNS		NSNS		Code	Open Source Code	Metallicity	Binary Physics Variations
Wagg	***2021***	**6-154**	**9-238**	**2-198**	**3-289**	**3-35**	**4-57**	**COMPAS**	✓	**53 bins between [1e-4, 3e-2]**	**20**
Shao	*2021*	12		2		✗		BSE	✓	0.02	Fryer rapid, Mandel & Mueller remnant mass
Breivik	*2020*	72		19		10		COSMIC	✓	0.02, 0.003	None
Lau	*2020*	✗		✗		35		COMPAS	✓	0.0142	Case BB always unstable, Single SN, alpha=0.1
Sesana	*2020*	4.2	6.5	✗		✗		BSE	✓	13 bins between [1e-4, 3e-2]	None
Lamberts	*2018*	25		✗		✗		BSE	✓	13 bins between [1e-4, 3e-2]	None
Liu	*2014*	6		3		16		BSE	✓	0.02	None
Belcyznski	*2010*	2.3	0	0.2	0	4	1.7	Startrack	✗	0.02 (disc, bulge), 0.001 (halo)	Optimistic CE, Pessimistic CE
Nelemans	*2001*	0		3		39		SeBa	✓	0.02	None

		Galaxy and Positioning				Detection		
First Author	Year	Star formation history	Spatial distribution	Galactic Components	Metallicity Dependent Distributions	SNR Limit	LISA Mission Time (yr)	Eccentricity Treatment
Wagg	***2021***	**Exponential 8-0 Gyr ago (thin disc), Exponential 12-8 Gyr ago (thick disc), Skewed gaussian 0-6 Gyr (bulge)**	**Exponential radial and vertical, different scale length/height for each component, thin disc has inside-out growth**	**Thin disc, thick disc, bulge**	✓	**7**	**4, 10**	**Full**
Shao	*2021*	Constant over 10 Gyr	Uniform flat disc of 15 kpc	Single disc	✗	5	4	Full
Breivik	*2020*	Constant over 10 Gyr (thin disc), 1 Gyr burst 10 Gyr ago (bulge), 1 Gyr burst 11 Gyr (thick disc)	McMillan 2011	Thin disc, thick disc, bulge	✗	7	4	Full
Lau	*2020*	Constant	Miyamoto & Nagai potential (disc), Wilkinson & Evans potential (halo)	Single disc or halo	✗	8	4	Full
Sesana	*2020*	FIRE simulation	FIRE simulation	Everything within 300kpc	✓	7	4, 10	Ignored during detection
Lamberts	*2018*	FIRE simulation	FIRE simulation	Everything within 300kpc	✓	5	4	Ignored during detection
Liu	*2014*	Constant over 13.7 Gyr	Exponential radial, sech^2 vertical (Benacquista+2007)	Single disc	✗	7	2	Assumed circular
Belcyznski	*2010*	Constant over 10 Gyr (disc), 1 Gyr burst 10 Gyr ago (bulge), burst at 13 Gyr (halo)	Exponential sphere (bulge), exponential radial and vertical (disc), spherical shell (halo)	Disc, bulge, halo	✗	10	1	Full
Nelemans	*2001*	Exponential over 10 Gyr	Exponential radial, sech^2 vertical	Single disc	✗	1, 5	1	Full

Figure 3.12: A table comparing previous studies of a similar nature to this work. In the detection rate section darker shades of blue indicate higher detection rates. In other sections, darker shades of green indicate more detailed models or more accurate choices. Entries in detection rates with two values listed correspond to the different mission lengths (penultimate column). The works listed in the table are Nelemans et al. (2001), Belczynski et al. (2010), Liu & Zhang (2014), Lamberts et al. (2018), Sesana et al. (2020), Lau et al. (2020), Breivik et al. (2020b) and Shao & Li (2021).

detection rate for all DCOs in this paper compared to our work is unsurprising given the relatively high SNR threshold of 10 and short mission length of 1 year. The reduced mission length means that the source signal has much less time to accumulate, whilst also fewer WDWDs can be resolved in this time, leading to a weaker signal and an increased Galactic confusion noise relative to our work.

Liu & Zhang (2014) — performed a similar investigation using a different population synthesis code and find higher rates than earlier works. Their lower detection threshold and longer mission length compared to Belczynski et al. (2010) likely explains the relatively increased rates. Yet their rates are still significantly below what we find. This could be for several reasons; they assume all binaries are circular both in their evolution and for detection. This means that systems may not have inspiralled as far before the LISA mission or may appear to have weaker gravitational waves when eccentricity is not accounted for. They also use a simplified model for the Milky Way with a single disc of one metallicity and constant star formation, whilst also using a mission length half what we assume. Each of these factors likely contributes to the lower overall detection rates.

Lamberts et al. (2018) — presented a new approach to the problem by using the FIRE simulation (Hopkins et al., 2014) to distribute their sources rather than an analytical model of the Milky Way and were the first paper in this area to incorporate metallicity dependence into their Milky Way model. Sesana et al. (2020) followed up on this paper using the same simulated BHBH population and presented updated results for the number of expected BHBH detections. They find significantly fewer BHBHs than our fiducial model despite using the same SNR threshold and LISA mission length. The discrepancy between the results of Sesana et al. (2020) and those presented in this work could be caused by different treatments of eccentricity. Unlike our work, Sesana et al. (2020) assume that all binaries are circular for the purpose of detection in LISA, which could result in a lower number of detections by missing eccentric binaries that appear as weaker signals when assumed to be circular. This is especially relevant as we find that around 87% of LISA detectable BHBHs are not circular and around 21% have significant eccentricity (see Fig. 3.3d). We also improve upon this work by using a larger number of metallicity bins compared to Sesana et al. (2020), since a low number of metallicity bins can produce artificial features in the mass distribution of DCOs and possibly affect the detection rate (see Dominik et al., 2015; Neijssel et al., 2019; Kummer, 2020, and also appendix 8.4 for further discussion). Finally, it could be that different implicit assumptions in their population synthesis code lead to differences in our results (Toonen et al., 2014).

Lau et al. (2020) — focussed on the number of Galactic NSNS binaries that could be detected by LISA. Their study uses the same population synthesis code, COMPAS, as this work, though an earlier version. Despite this, their study finds a much larger number of detections. They make several different physical assumptions in their population synthesis, using the Fryer et al. (2012) *rapid* remnant mass prescription, assuming the optimistic CE scenario, limiting the maximum neutron star mass to $2\,\mathrm{M}_{\odot}$ and not implementing PISN. However, we note that none of these assumptions strongly affect the NSNS LISA detection rate (see bottom panel of Fig. 3.8, models K, L, M and O) and so this is unlikely to entirely account for the differences. It is also important to highlight that COMPAS has received several improvements and bug fixes since Vigna-Gómez et al. (2018) (which contains the simulations used by Lau et al. 2020) and these could possibly have affected the formation rate of NSNSs.

Yet it is most likely that the remaining difference between our results is due to way in which we simulate the Milky Way. Lau et al. (2020) use a model for the Milky Way similar to that of Breivik et al. (2020b), which we use to estimate the impact of the choice of MW model in Appendix 8.4. The Milky Way model by Breivik et al. (2020b) applies only two metallicity bins, while we consider a range of metallicities between 10^{-4} and 0.03. When applying the simpler model for the Milky Way, we find that the NSNS detection rate is increased by at least a factor of two. Lau et al. (2020) only uses a single metallicity, and so assuming that all star formation happens at a single high metallicity (which has a high efficiency of producing NSNS), could lead to an even greater overestimate of the detection rate. Hence, we expect the low number of metallicity bins in their Milky Way model to be the main driver behind the discrepancy between our results.

Breivik et al. (2020b) — introduced the population synthesis code COSMIC and presented detections for many different DCO types in LISA using this code. They find that LISA will detect 72 BHBHs, 19 BHNSs and 10 NSNSs in the Milky Way over a 4 year mission. Breivik et al. (2020b) make many physical assumptions that differ from our fiducial model, the most notable being that they assume the optimistic CE scenario and that case BB mass transfer is always unstable, whilst also using a simpler model for the Milky Way (see Appendix 8.4). Thus for better comparison we ran our simulation using model F and the Milky Way model from Breivik et al. (2020b). This results in 97, 101 and 43 detections for BHBHs, BHNSs and NSNSs respectively. Therefore, though our BHBH rates agree reasonably well, we predict much higher rates for BHNSs and NSNSs. These differences are likely due to using a different population synthesis code (COSMIC), which has different underlying physics assumptions from COMPAS. Given our strong agreement for BHBHs, it is possible that COSMIC and COMPAS handle NSs differently and so lead to different detection rates for DCOs containing NSs. However checking this would require a more in-depth study of the intrinsic formation rate of DCOs containing NSs in the two codes.

Shao & Li (2021) — most recently investigated the detectability binaries containing BHs in LISA using BSE and a relatively simple model for the Milky Way (assuming a uniform flat disc, constant star formation and a single metallicity). They assume that kicks for NSs formed through ECSN are slightly higher than our work ($50\,\mathrm{km\,s^{-1}}$ instead of $30\,\mathrm{km\,s^{-1}}$). This may account for their particularly low BHNS rate (as the binaries would be more likely to disrupt), which is a factor of 20 lower than ours, but we expect their assumption of the optimistic CE scenario, reduced Wolf-Rayet winds and lower SNR detection threshold could partially offset this. As we show in Appendix 8.4, their use of a simpler Milky Way model, especially with only a single metallicity, would lead to an underestimate of the BHBH and BHNS rates, which may explain the discrepancy in our results.

Overall, since the work of Nelemans et al. (2001), in addition to the LISA mission specifications, the methods that we use to simulate binaries and the Milky Way have all changed significantly. We now predict that LISA detections of these massive DCOs are dominated by BHBHs, rather than NSNSs, whilst the absolute detection rates for BHBHs and BHNSs are much higher. Further studies in this area could improve on this work by including the effects of systemic kicks on the position of systems in the Milky Way and accounting for contributions from other formation channels.

7 Conclusion & Summary

We provide predictions for the detection rate and population properties of LISA detectable BHBH, BHNS and NSNS. We use a novel empirically-informed analytical model for the metallicity dependent star formation history of the Milky Way, calibrated against the APOGEE stellar spectroscopic survey. We use this to model Monte-Carlo realisations of the present-day BHBH, BHNS and NSNS populations in our Milky Way. For the binary population, we use the results of a large grid of simulations performed with the rapid population synthesis code COMPAS (Broekgaarden et al., 2021, 2022). Our suite of simulations consider 20 physics variations, which represent the most common uncertainties in binary physics. We investigate the detectability of sources in the planned space-based GW detector LISA using the `LEGWORK` package (Wagg et al., 2022b), that we specifically developed for this purpose and is publicly available. We investigate the results expected for a 4- and 10-year LISA mission. Our main conclusions can be summarised as:

1. **Total detections:** We predict 30-370 detections in a 4-year LISA mission, across all our simulations for varying physics assumptions. This increases to about 50-550 for a 10-year LISA mission. Although the number of detections per type can vary by about 2 orders of magnitude, we find that the total detection rate is fairly robust, among the variations we have considered.

2. **Detections by type:** For our fiducial model, we predict a total of 124 ± 11 detections and out of these we find about 74 ± 9 BHBHs, 42 ± 6 BHNSs and 8 ± 3 NSNSs. The errors quoted here are the 1-σ Poisson uncertainties resulting from the random initialisation of the Milky Way.

3. **Physics variations:** Among the model variations we consider, we find that our predictions for the rates for the different DCO types are robust within a factor of 2 of the fiducial rate, with the following exceptions. For BHBHs, the rate is most sensitive to the treatment of common envelope phases or an increase of the WR wind mass loss. For BHNSs and NSNSs, the assumptions regarding the assumed common envelope ejection efficiency, treatment of case BB mass transfer and the kicks are most important. In addition, the assumed mass transfer efficiency impacts the BHNS (see Fig. 3.8).

4. **Probing the black hole mass distribution and the lower mass gap:** We expect LISA to predominantly detect lower mass BHs (with 90% of BHBH and BHNSs having BH masses lower than $11\,\mathrm{M}_\odot$ in our fiducial simulations). This is in stark contrast to current ground-based detectors which are heavily biased towards high mass systems, since they can be detected out to larger distances. We predict that between 15-156 systems with a component with a mass between 2.5-$5\,\mathrm{M}_\odot$ would be detected by a 4-year LISA, for our different physics variations. An exception is model using the Fryer Rapid remnant mass prescription, which produces none. However, it will likely be hard for

LISA data to make a significant impact on this debate on the existence of the lower mass gap given the uncertainty of the individual masses (see Fig. 3.9).

5. **Eccentricity distribution:** We find that for all DCO types a large fraction of detectable systems still have nonzero eccentricities ($e > 0.01$) when entering the LISA band, which can be used to distinguish them from the more numerous WDWD binaries, which are largely expected to be circular. In particular, for our fiducial model, we find that this is the case for around three quarters of detectable binaries. Furthermore, around 16% of detectable binaries have eccentricities that are so high ($e > 0.3$) that the emission at frequencies corresponding to higher order harmonics start to dominate (see Fig. 3.3).

6. **Distinguishing from WDWD sources:** For about half of all detections we expect that we will be able to confidently determine the type of compact objects involved and this increases to 60% for a 10-year LISA mission (see Sec. 5.1.1).

7. **Chirp mass determinations:** For about 10% of systems we expect to be able to determine the chirp mass better than 10% and this increases to 15% for a 10-year LISA mission (see Fig. 3.6).

8. **Prospects for finding EM counterparts:** We expect about 13% of detections with a sky localisation better than 1 degree for our fiducial model (though the fraction remains roughly constant among model variations). This fraction remains the same for a 10-year LISA mission, meaning that the number increases proportionally. This will be of interest for electromagnetic searches for counterparts, in particular for radio pulsar searches with SKA (see Fig. 5.2).

9. **Benefits of extending the LISA mission:** The number of detections scale approximately as $\sqrt{T_{\rm obs}}$, where $T_{\rm obs}$ is the mission length. Therefore, extending the LISA mission from 4- to 10-years increases the number of detections by about 60% for each model variation. A further important benefit is the improvement of the characterisation of the sources, since the relative error on the frequency derivative (which dominates the relative error in the chirp mass) scales as $T_{\rm obs}^{-2.5}$ for stationary sources (Eq. 3.31). We find that the number of systems with chirp masses that can be measured better than 10% increases by a factor of 2.4 for each model variation. In addition, the number of systems with a sky localisation better than one degree increases by a factor of 1.5. Overall, the number of sources that can be unambiguously distinguished from WDWDs increases by almost a factor of 2 (see Section 5.1.1).

8 Appendices

8.1 Population Synthesis

In this section we summarise the main assumptions and settings from the population synthesis simulation from Broekgaarden et al. (2021, 2022).

8.1.1 Initial conditions

Broekgaarden et al. (2021, 2022) simulate between 1 and 100 million massive binaries for each of 50 metallicities equally spaced in log space between $Z \in [0.0001, 0.022]$, where Z is the mass fraction of heavy elements. They simulate more binaries for higher metallicities so that large enough sample of DCOs at each metallicity (since DCOs are formed at a lower rate at higher metallicities). These metallicities span the allowed metallicity range for the original fitting formulae on which COMPAS is based (Hurley et al., 2000). This is repeated for 19 physics variations (see Section 8.1.3) and so in total over two billion binaries were simulated.

Each binary is sampled from initial distributions for the primary and secondary masses as well as the separation. The primary mass, that is the mass of the initially more massive star, is restricted to $m_1 \in [5, 150]\,\mathrm{M}_\odot$, which spans the range of interest for NS and BH formation in binary systems, and drawn from the Kroupa (2001) initial mass function (IMF), $p(m_1) \propto m_1^{-2.3}$. The secondary mass, m_2, is drawn using the initial mass ratio of the binary, $q \equiv m_2/m_1$, which Broekgaarden et al. (2021, 2022) assume to be uniform on $[0, 1]$, therefore $p(q) = 1$ (e.g. consistent with Sana et al., 2012). They additionally restrict the secondary masses $m_2 \geq 0.1\,\mathrm{M}_\odot$, which is approximately the minimum mass for a main sequence star. They assume that the initial separation follows a flat in the log distribution with $p(a_i) \propto 1/a_i$ and $a_i \in [0.01, 1000]\,\mathrm{AU}$ (**?**Abt, 1983). They assume that all binary orbits are circular at birth to reduce the dimensions of initial parameters. Since they focus on post-interaction binaries which will have circularised after mass transfer they argue this is an reasonable assumption (as many studies have in the past) and is likely not critical for predicting detection rates (Hurley et al., 2002; de Mink & Belczynski, 2015).

Broekgaarden et al. (2021, 2022) apply the adaptive importance sampling algorithm STROOPWAFEL (Broekgaarden et al., 2019) to improve the yield of their sample. This algorithm increases the prevalence of target DCOs (BHBHs, BHNSs and NSNSs in this case) in the sample and assigns each a weight, w, which represents the probability of drawing the DCO without STROOPWAFEL in effect.

8.1.2 Physical assumptions in the fiducial model

Stellar Evolution: To follow the evolution of massive stars, COMPAS relies on fitting formulae by Hurley et al. (2000) to detailed single star models by Pols et al. (1998). COMPAS models the evolution of stars that lose or gain mass closely following the algorithms originally described in Tout et al. (1996) and Hurley et al. (2002).

Wind mass loss: Broekgaarden et al. (2021, 2022) follow the wind prescription from Belczynski et al. (2008), which was based on results from Monte Carlo radiative transfer simulation of Vink et al. (2000, 2001). They use the wind mass loss rates from Vink et al. (2001) for stars above 12500 K and the rates from Hurley et al. (2000) for cooler stars. Additionally, they use a separate, higher wind mass loss rate for luminous blue variable (LBV) stars, following Belczynski et al. (2008), to mimic observed LBV eruptions for stars with luminosities and effective temperatures above the Humphreys-Davidson limit. They use the Wolf-Rayet-like mass loss rate from Hamann & Koesterke (1998) with an additional metallicity scaling from Vink & de Koter (2005) for helium stars, and set $f_{\mathrm{WR}} = 1$. See Riley et al. (2022), Section 3 for the explicit equations.

Mass Transfer: In determining the stability of mass transfer Broekgaarden et al. (2021, 2022) use the ζ-prescription, which compares the radial response of the star with the response of the Roche lobe radius to the mass transfer (e.g. Hjellming & Webbink, 1987). The mass transfer efficiency, $\beta \equiv \Delta M_{\rm acc}/\Delta M_{\rm don}$, is the fraction of the mass transferred by the donor that is actually accreted by the accretor. They limit the maximum accretion rate for stars to $\Delta M_{\rm acc}/\Delta t \leq 10 M_{\rm acc}/\tau_{\rm KH}$, where $\tau_{\rm KH}$ is the Kelvin-Helmholtz timescale of the star (Paczyński & Sienkiewicz, 1972; Hurley et al., 2002). The maximum accretion rate for compact objects is limited to the Eddington accretion rate. If more mass than these rates is accreted then they assume that the excess is lost through isotropic re-emission in the vicinity of the accreting star (e.g. ?Soberman et al., 1997). They assume that all mass transfer from a stripped post-helium-burning-star (case BB) onto a neutron star or black hole is unstable (Tauris et al., 2015).

Common-Envelope: A common-envelope phase follows dynamically unstable mass transfer and Broekgaarden et al. (2021, 2022) parameterise this using the α-λ prescription from Webbink (1984) and de Kool (1990). They assume $\alpha = 1$, such that all of the gravitational binding energy is available for the ejection of the envelope. For λ they use the fitting formulae from Xu & Li (2010a,b). They assume that any Hertzsprung gap donor stars that initiate a common-envelope phase will not survive this phase due to a lack of a steep density gradient between the core and envelope (Taam & Sandquist, 2000; Ivanova & Taam, 2004; Klencki et al., 2021). This follows the 'pessimistic' common-envelope scenario (c.f. Belczynski et al., 2007). They remove any binaries where the secondary immediately fills its Roche lobe upon the conclusion of the common-envelope phase as they treat these as failed common-envelope ejections, likely leading to a stellar merger.

Supernovae: Broekgaarden et al. (2021, 2022) draw the remnant masses and natal kick magnitudes from different distributions depending on the type of supernova that occurs. For stars undergoing a general core-collapse supernova, they use the *delayed* supernova remnant mass prescription from Fryer et al. (2012). The *delayed* prescription does not reproduce a neutron star black hole mass gap and they use this as their default as it has been shown to provide a better fit for observed populations of DCOs (e.g. Vigna-Gómez et al., 2018). They draw the natal kick magnitudes from a Maxwellian velocity distribution with a one-dimensional root-mean-square velocity dispersion of $\sigma_{\rm rms}^{\rm 1D} = 265\,{\rm km\,s^{-1}}$ (Lyne & Lorimer, 1994; Hobbs et al., 2005). They assume that stars with helium core masses between 1.6–$2.25\,{\rm M_\odot}$ (Hurley et al., 2002) experience electron-capture supernovae (ECSN) (Nomoto, 1984, 1987; Ivanova et al., 2008). They set all remnant masses to $1.26\,{\rm M_\odot}$ in this case as an approximation of the solution to Equation 8 of Timmes et al. (1996). For these supernovae, they set $\sigma_{\rm rms}^{\rm 1D} = 30\,{\rm km\,s^{-1}}$ (e.g. Pfahl et al., 2002; Podsiadlowski et al., 2004). They assume that stars that undergo case BB mass transfer (Dewi et al., 2002) experience extreme stripping which leads to an ultra-stripped supernova (Tauris et al., 2013, 2015). For these supernovae they calculate the remnant mass using the Fryer et al. (2012) prescription and use $\sigma_{\rm rms}^{\rm 1D} = 30\,{\rm km\,s^{-1}}$ (as with ECSN). Stars with final helium core masses between 35-135 ${\rm M_\odot}$ are presumed to undergo a pair-instability, or pulsational pair-instability supernova (e.g. Woosley et al., 2007; Farmer et al., 2019). They follow the prescription from Marchant et al. (2019) as implemented in (Stevenson et al., 2019) for these supernovae. They assume that kicks are isotropic in the frame of the collapsing star. They adopt a maximum neutron star mass of $2.5\,\mathrm{M}_\odot$ (e.g. Kalogera & Baym, 1996; Fryer et al., 2015; Margalit & Metzger, 2017) for the fiducial model and change the Fryer et al. (2012) prescription accordingly.

8.1.3 Model variations

In addition to their fiducial model for the formation of DCOs, Broekgaarden et al. (2021, 2022) explore 19 other models in which they change various aspects of the mass transfer, common-envelope, supernova and wind mass loss physics assumptions in order to assess the effect of their uncertainties on the overall double compact object detection rates and distributions. Each of the models varies a single physics assumption (fiducial assumptions are outlined in Section 8.1.2) and these models are outlined in Table 3.2.

Their fiducial model is labelled model A. Models B-D focus on changes to the mass transfer physics assumptions. They explore the effect of fixing the mass transfer efficiency β to a constant value, rather than allowing it to vary based on the maximum accretion rate. In models B, C, D, in which they set the value of β to 0.25, 0.5 and 0.75 respectively.

Models E′-K focus on altering the common-envelope physics. In model E′ we modify model E from Broekgaarden et al. (2021, 2022) to investigate the consequence of assuming that case BB mass transfer is always unstable, whilst allowing Helium HG donors to survive CE events. They change the common-envelope efficiency parameter to $\alpha_{\mathrm{CE}} = 0.1, 0.5, 2.0, 10.0$ in models G, H, I and J respectively. In model K, they relax their restriction that Hertzsprung gap donor stars cannot survive common-envelope events, thereby following the 'optimistic' common-envelope scenario. They combine this with model E′ in model F.

In models L-R they consider changes related to their assumptions about supernova physics. Model L uses the alternate *rapid* remnant mass prescription from Fryer et al. (2012) instead of the *delayed* prescription. They change the maximum neutron star mass in models M and N to 2 and $3\,\mathrm{M}_\odot$ respectively to account for the range of predicted maximum neutron star masses. Model O removes the implementation of pair-instability and pulsational pair-instability supernovae. In models P and Q they decrease the root-mean-square velocity dispersion for core-collapse supernovae to explore the effect of lower kicks. Model R removes the natal kick for all black holes.

Finally, in models S-T Broekgaarden et al. (2021, 2022) investigate the effect of changing their assumption about wind mass loss rates, specifically for Wolf-Rayet winds. They vary f_{WR} to 0.1 and 5.0 in models S and T respectively. These values approximately span the current range of possible Wolf-Rayet wind efficiencies suggested from observations (e.g. Vink, 2017; Hamann et al., 2019; Shenar et al., 2019; Miller-Jones et al., 2021; van Son et al., 2022a).

8.2 Detection Rate Normalisation

In this section we explain the normalisation process that we refer to in Section 2.3. From each simulated instance of the Milky Way we extract the fraction of targets that are detectable, where we define a target as one of BHBH, BHNS or NSNS that merges in a Hubble time. To convert the detectable fraction to a detection rate for the Milky Way, we write that the *number* of detectable targets in the Milky Way is

$$N_{\mathrm{detect}} = f_{\mathrm{detect}} \cdot N_{\mathrm{target,MW}}, \tag{3.16}$$

Model	Physics Variation
A	Fiducial (see Section 8.1.2)
B	Fixed mass transfer efficiency of $\beta = 0.25$
C	Fixed mass transfer efficiency of $\beta = 0.5$
D	Fixed mass transfer efficiency of $\beta = 0.75$
E′	Case BB mass transfer always unstable
F	Case BB always unstable + Optimistic CE
G	CE efficiency parameter $\alpha = 0.1$
H	CE efficiency parameter $\alpha = 0.5$
I	CE efficiency parameter $\alpha = 2$
J	CE efficiency parameter $\alpha = 10$
K	HG donor stars initiating a CE survive CE
L	Fryer rapid SN remnant mass prescription
M	Maximum NS mass is fixed to $2\,\mathrm{M}_\odot$
N	Maximum NS mass is fixed to $3\,\mathrm{M}_\odot$
O	PISN and pulsational-PISN not implemented
P	$\sigma_{\rm rms}^{\rm 1D} = 100\,\mathrm{km\,s^{-1}}$ for core-collapse supernova
Q	$\sigma_{\rm rms}^{\rm 1D} = 30\,\mathrm{km\,s^{-1}}$ for core-collapse supernova
R	Black holes receive no natal kick
S	Wolf-Rayet wind factor $f_{\rm WR} = 0.1$
T	Wolf-Rayet wind factor $f_{\rm WR} = 5.0$

Table 3.2: A description of the 20 binary population synthesis models used in this study. A is the fiducial model, B-D change mass transfer physics, E′-K change common-envelope physics , L-R change supernova physics and S-T change wind mass loss (c.f. Broekgaarden et al., 2021, Table 2).

where $f_{\rm detect}$ is the fraction of targets in the instance that were detectable and $N_{\rm target,MW}$ is the total number of targets that have been formed in the Milky Way's history. We can further break this total down into

$$N_{\rm target,MW} = \langle \mathcal{R}_{\rm target} \rangle \cdot M_{\rm SF,MW}, \tag{3.17}$$

where $\langle \mathcal{R}_{\rm target} \rangle$ is the average number of targets formed per star forming mass and $M_{\rm SF,MW}$ is the star forming mass of the Milky Way, meaning the total mass of every star ever formed in the Milky Way.

8.2.1 Average target formation rate

Double compact object formation is metallicity dependent, so we find the average rate as the integral over metallicity, which is given by

$$\langle \mathcal{R}_{\rm target} \rangle = \int_{Z_{\rm min}}^{Z_{\rm max}} p_Z \mathcal{R}_{\rm target,Z}\, \mathrm{d}Z\,, \tag{3.18}$$

where $Z_{\rm min}, Z_{\rm max}$ are the minimum and maximum sampled metallicities, p_Z is the probability of forming a star at the metallicity Z (which can be found using the distribution in Frankel et al., 2018) and $\mathcal{R}_{\rm target,Z}$ is the number of targets formed per star forming mass,

$$\mathcal{R}_{\rm target,Z} = \frac{N_{\rm target,Z}}{M_{\rm SF,Z}}. \tag{3.19}$$

In practice, this integral is instead approximated as a sum over the metallicity bins that we use in our simulation. The number of targets in our sample at a metallicity Z, $N_{\rm target,Z}$, can be written simply as the sum of the targets' weights:

$$N_{\rm target,Z} = \sum_{i=1}^{N_{\rm binaries,Z}} w_i \theta_{\rm target,i}, \tag{3.20}$$

where w_i is the binary's adaptive importance sampling weight assigned, $N_{\rm binaries,Z}$ is the number of binaries at metallicity Z in our sample and $\theta_{\rm target,i}$ is only 1 when the binary is a target and otherwise 0.

The total star forming mass at a metallicity Z, $M_{\rm SF,Z}$, can be written as

$$M_{\rm SF,Z} = \frac{\langle m \rangle_{\rm COMPAS,Z}}{f_{\rm trunc}} N_{\rm binaries,Z}, \tag{3.21}$$

where $\langle m \rangle_{\rm COMPAS}$ is the average star forming mass of a binary in a simulation using our cutoffs (discussed in Section 2.1) and $f_{\rm trunc}$ is the fraction of the total stellar mass from which our COMPAS simulations sample, given our truncated mass and separation ranges (see Section 2.1). These truncations mean that only $f_{\rm trunc} \approx 0.17$ of the stellar mass in the galaxy is sampled from.

8.2.2 Total star forming mass in the Milky Way

It is important to distinguish between the *total* mass of every star formed over the entire history of the Milky Way and the *current* stellar mass in the Milky Way. Many stars born in the Milky Way are no longer living and have lost much of their mass to stellar winds and supernovae, thus the current stellar mass in the Milky Way is an underestimate of the total star forming mass.

Licquia & Newman (2015) find that the total stellar mass today in the Milky Way is $6.08 \pm 1.14 \times 10^{10}\,\mathrm{M}_\odot$. This total includes all stars and stellar remnants (white dwarfs, neutrons stars and black holes) but *excludes* brown dwarfs. We can write that the total mass of every star every formed in the Milky Way is

$$M_{\rm SF,MW} = (6.08 \pm 1.14) \times 10^{10}\,\mathrm{M}_\odot \cdot \frac{\langle m \rangle_{\rm SF,total}}{\langle m \rangle_{\rm SF,today}}, \tag{3.22}$$

where $\langle m \rangle_{\rm SF,total}$ is the average mass of a star over the history of the Milky Way and is defined as

$$\langle m \rangle_{\rm SF,total} = \int_0^{t_{\rm MW}} p_{\rm birth}(\tau) \int_{0.01}^{200} \zeta(m)\ m \, {\rm d}m \, {\rm d}\tau \, , \tag{3.23}$$

where $t_{\rm MW}$ is the age of the Milky Way, $\zeta(m)$ is the Kroupa (2001) IMF function and $p_{\rm birth}(\tau)$ is the probability of a star being formed at a lookback time τ (Eq. 3.2). $\langle m \rangle_{\rm SF,today}$ is the average mass of all stars and stellar remnants (excluding brown dwarfs) present in the Milky Way today is defined as follows (note that we integrate from 0.08 not 0.01 since observations of today's Milky Way mass exclude brown dwarfs)

$$\langle m \rangle_{\rm SF,today} = \int_0^{t_{\rm MW}} p_{\rm birth}(\tau) \int_{0.08}^{200} \zeta(m)\ m_{\rm today} \, {\rm d}m \, {\rm d}\tau \, , \tag{3.24}$$

where $m_{\rm today}(m, Z, \tau)$ is the current mass of a star that was formed τ years ago at a metallicity Z. We calculate $m_{\rm today}(m, Z, \tau)$ by interpolating the final masses given by COMPAS for a grid of single stars over different masses and metallicities using the Fryer et al. (2012) delayed prescription and default wind mass loss settings. For Z, we use the average star forming metallicity in the Milky Way at a lookback time τ using our galaxy model. Evaluating Equation 3.22, we find that the total mass of every star that has ever formed in the Milky Way is

$$\begin{aligned} M_{\rm SF,MW} &= (6.1 \pm 1.1) \times 10^{10}\,{\rm M_\odot} \cdot \frac{0.378\,{\rm M_\odot}}{0.221\,{\rm M_\odot}}, \\ &= (10.4 \pm 1.1) \times 10^{10}\,{\rm M_\odot}, \end{aligned} \tag{3.25}$$

an increase of approximately 70% from the value still in stars today!

8.2.3 Normalisation summary

Finally, we can substitute Equations 3.18 and 3.22 into 3.16 and write that the overall normalisation of the detection rate is calculated as

$$N_{\rm detect} = f_{\rm detect} \cdot 10.4 \times 10^{10}\,{\rm M_\odot} \times \sum_{Z=Z_{\rm min}}^{Z_{\rm max}} p_Z \frac{\sum_{i=1}^{N_{\rm binaries,Z}} w_i \theta_{\rm target,i}}{\frac{\langle m \rangle_{\rm COMPAS,Z}}{f_{\rm trunc}} \sum_{i=1}^{N_{\rm binaries,Z}} w_i}. \tag{3.26}$$

8.3 Calculation of the uncertainties in the chirp mass for detectable sources

How accurately the chirp mass of a detected binary can be determined depends on the signal to noise ratio, duration of the mission, its orbital frequency and the time derivative of the orbital frequency.

Here we describe how we estimate the uncertainty of the chirp mass. First, consider the chirp mass, which can be expressed as

$$\mathcal{M}_c = \frac{c^3}{G}\left(\frac{5\pi}{48n}\frac{\dot{f}_n}{F(e)}\right)^{3/5}\frac{1}{(2\pi f_{\rm orb})^{11/5}}, \tag{3.27}$$

where f_n is the frequency of the n-th harmonic, $f_{\rm orb}$ is the orbital frequency, $\mathcal{M}_c$ is the chirp mass (defined in Eq. 3.12), e is the eccentricity and

$$F(e) = \frac{1+\frac{73}{24}e^2+\frac{37}{96}e^4}{(1-e^2)^{7/2}}, \tag{3.28}$$

is the enhancement factor of gravitational wave emission for an eccentric binary over an otherwise identical circular binary (Peters & Mathews, 1963, Eq. 17). In practice, we will use the dominating harmonic, with $n = n_{\rm dom}$ and $f_n = n_{\rm dom} f_{\rm orb} = f_{\rm dom}$. The dominating harmonic for circular binaries is $n_{\rm dom} = 2$ and so the dominating frequency is twice the orbital frequency.

Therefore the chirp mass uncertainty can be estimated as

$$\frac{\Delta\mathcal{M}_c}{\mathcal{M}_c} = \frac{11}{5}\frac{\Delta f_{\rm orb}}{f_{\rm orb}} + \frac{3}{5}\frac{\Delta \dot{f}_{\rm dom}}{\dot{f}_{\rm dom}} + \frac{3}{5}\frac{\Delta F(e)}{F(e)}, \tag{3.29}$$

We estimate the frequency uncertainties using Takahashi & Seto (2002), such that

$$\frac{\Delta f_{\rm orb}}{f_{\rm orb}} = 4\sqrt{3}\cdot\frac{1}{\rho}\frac{1}{T_{\rm obs}}\frac{1}{f_{\rm orb}}, \tag{3.30}$$

$$\frac{\Delta \dot{f}_{\rm dom}}{\dot{f}_{\rm dom}} = 6\sqrt{5}\cdot\frac{1}{\rho}\left(\frac{1}{T_{\rm obs}}\right)^2\frac{1}{\dot{f}_{\rm dom}}, \tag{3.31}$$

where ρ is the signal-to-noise ratio and $T_{\rm obs}$ is the LISA mission length. We estimate the eccentricity uncertainty, Δe, following the methods of Lau et al. (2020) and Korol & Safarzadeh (2021), which use the relative SNRs of different harmonics to work out the eccentricity. We propagate this uncertainty such that

$$\frac{\Delta F(e)}{F(e)} = \Delta e\cdot\frac{(1256+1608e^2+111e^4)e}{96+196e^2-255e^4-37e^6}. \tag{3.32}$$

We use Eq. 3.29 to calculate the chirp mass uncertainty for each DCO type in our sample and plot it in Fig. 3.6.

8.4 Assessing the impact of Milky Way model choices

The model that we use for the Milky Way adds several layers of complexity, accounting for the inside-out growth of the thin disc, using empirically informed star formation histories that are a function of time and assigning metallicities based on the position and age of binaries. In this section, we repeat our main analysis but instead apply a simpler model for the Milky Way in order to assess the effect of these added features. For this purpose, we use model for the Milky Way used in Breivik et al. (2020b) as this is representative of the models used in most previous works.

Their model can be summarised as follows: the Milky Way is assumed to comprise of three components, a thin disc, a thick disc and a bulge. The spatial distributions and relative masses for these components are given in McMillan (2011). Breivik et al. (2020b) assume constant star formation over 10 Gyr for the thin disc, a 1 Gyr burst of star formation 11 Gyr ago for the thick disc and a 1 Gyr burst of star formation 10 Gyr ago for the bulge. A major difference is that only two metallicities are used and they are assigned to binaries independent of age or position. Binaries formed in the thin disc and bulge are assumed to have a metallicity of $Z = 0.02$ and those formed in the thick disc are assumed to have $Z = 0.003$.

We show the spatial metallicity distribution for this model in Fig. 3.13 in the same form as Fig. 3.1 for ease of comparison between our models. The two main differences we can see between Fig. 3.1 and 3.13 are that the Breivik et al. (2020b) model is more centrally concentrated and only has two fixed metallicity populations.

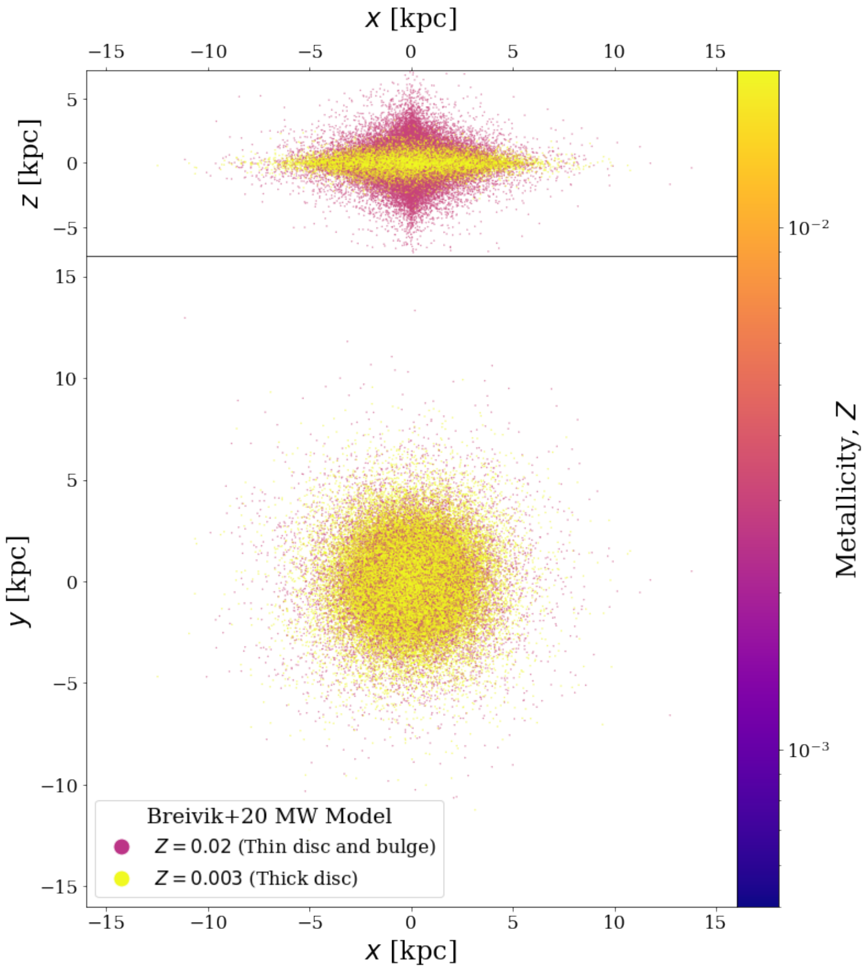

Figure 3.13: As Fig. 3.1 (right panel), but for the Milky Way model used in Breivik et al. (2020b).

When applying this simpler Milky Way model in combination with our fiducial binary physics assumptions (model A), we find that the expected number of detections for BHBHs, BHNSs and NSNSs for a 4-year LISA mission is 52, 25 and 17 respectively. Thus the BHBH detection has decreased slightly compared to our main findings, whilst for BHNSs and NSNSs the rate has approximately halved and doubled respectively.

Moreover, the distribution of parameters within the population, particularly the mass distributions, are notably disparate. By using only two fixed metallicity populations, unphysical artifacts are introduced into distribution of DCO masses (e.g. Dominik et al., 2015; Neijssel et al., 2019; Kummer, 2020). For example, in Fig. 3.14, we show the black hole mass distribution for BHNSs produced by the simulation using the simple Milky Way model. Despite the fact that these KDEs use the same bandwidth as Fig. 3.3, the distributions show many more sharp transitions, which is a result of pileups occurring at specific masses for specific metallicities. Moreover, the lack of lower metallicities systems means that higher mass systems are not formed and so we see the distributions do not include a high mass tail such as in our fiducial results.

The unphysical artifacts present in the mass distributions can have far-reaching effects since the masses of DCOs affect most other parameters. The inspiral time and SNR are directly dependent on the mass, whilst the uncertainty estimates depend on the SNR. This means that the artifacts can affect the predictions for most distributions of LISA detectable populations.

Overall, we find that previous studies that use Milky Way models analogous to this simpler model may significantly underestimate the LISA BHNS rate whilst overestimating the NSNS detection rate. They may also miss higher mass systems (particular for BHNSs) and contain unphysical artifacts in their parameter distributions.

8.5 Estimating the number of pulsars for a given sky area in SKA

In this section, we perform some back-of-the-envelope calculations in order to estimate the number of pulsars that SKA will observe within a given sky area.

First, we consider how many pulsars SKA is likely to detect. Keane et al. (2015) uses PSRPOPPy (Bates et al., 2014) to simulate the Milky Way pulsar population. They find that for SKA-1, approximately 10000 pulsars will be discovered. The second phase of SKA, which should be in operation by the time of the LISA mission, would yield a total of 35000-41000 pulsars (Keane et al., 2015). We use the average, 38000, in further estimates below. Moreover, we are only interested in pulsars that are part of a binary system. We estimate this pulsar binary fraction as the fraction of known pulsars that are in binaries using the ATNF Pulsar Catalogue[7] (Manchester et al., 2005). 290 of the 2872 currently known pulsars are in binary systems and thus we estimate the binary fraction of pulsars as 10%. Therefore, we expect that SKA-1 and SKA-2 will detect approximately 1000 and 3800 binary pulsars respectively.

Next, we can find the total number of pulsars SKA will detect in a patch on the sky. The total sky area that the SKA covers is approximately $5700\,\mathrm{deg}^2$, which is calculated by integrating over the sky for all Galactic longitudes and Galactic latitudes limited to $|b| < 10^\circ$

---

[7] https://www.atnf.csiro.au/research/pulsar/psrcat

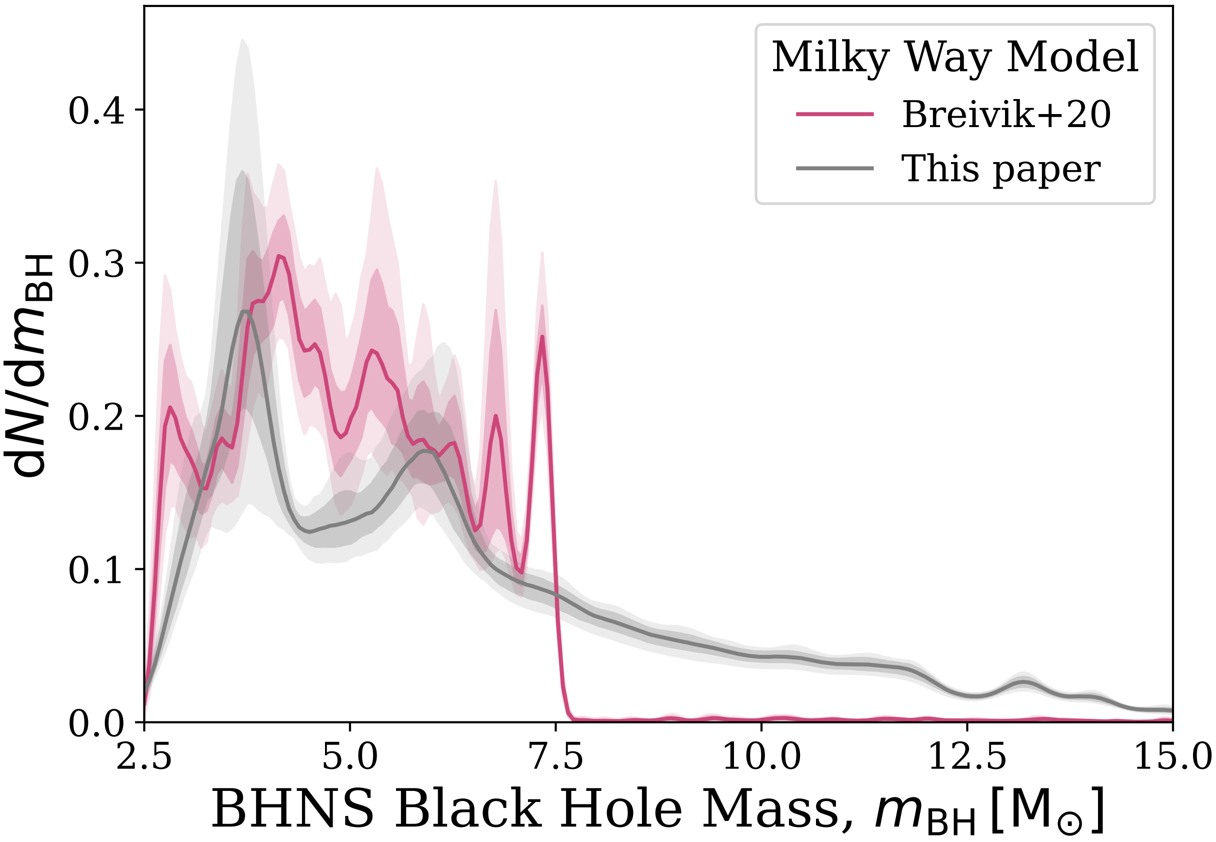

Figure 3.14: As Fig. 3.3b, but for the Milky Way model used in Breivik et al. (2020b). Dotted lines show the distribution from Fig. 3.3b for comparison. .

and $\delta < 45°$, which are the limits on SKA-mid (Keane et al., 2015). If we assume that the pulsars are found uniformly across the sky, this means that roughly 0.2 and 0.7 binary pulsars are expected per square degree for SKA-1 and SKA-2 respectively. Note that the assumption of a uniform distribution is not realistic as pulsars will tend to be far more concentrated in the Galactic centre but we use it to provide a slightly optimistic estimate.

Overall, we therefore expect a single pulsar per $5.7\,\mathrm{deg}^2$ and $1.5\,\mathrm{deg}^2$ for SKA-1 and SKA-2 respectively, which correspond to angular resolutions of $\sigma_\theta = 1.3°$ and $\sigma_\theta = 0.7°$.

8.6 Supplementary material

This section contains various supplementary plots for this work. Figure 3.15 highlights the location of eccentric sources on the LISA sensitivity curve plot, whilst Figure 3.16 shows the properties of detectable sources at the moment of double compact object formation. Figure 3.17 shows the relative rates of different formation channels that result in detectable LISA sources.

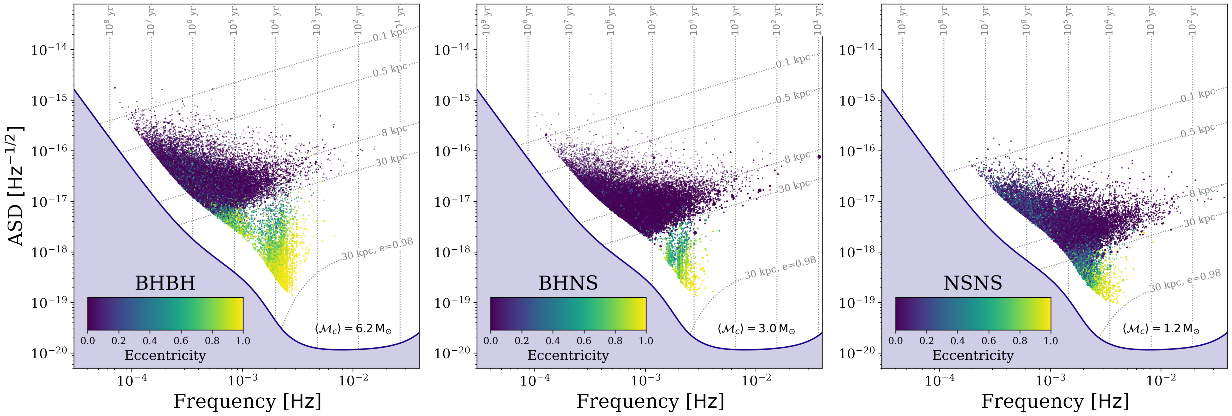

Figure 3.15: As the bottom panels of Fig. 3.2, but without the density distributions and scatter points are coloured by their eccentricity. We show eccentric sources are located in an offshoot below the 30 kpc around 2 mHz.

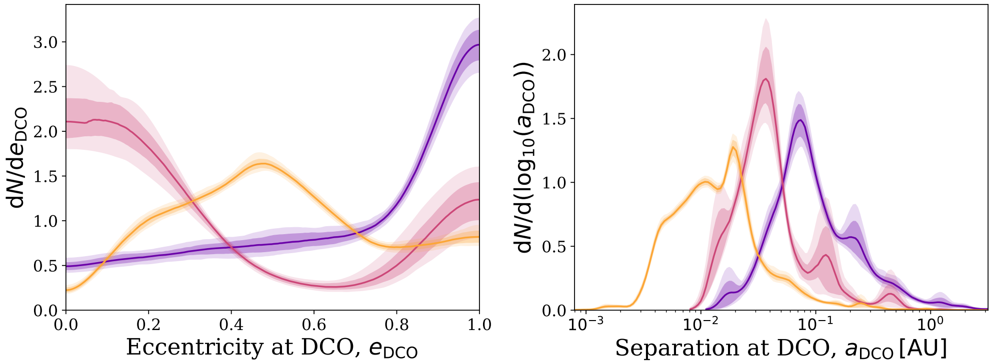

Figure 3.16: As Fig. 3.3, but for detectable system properties at DCO formation.

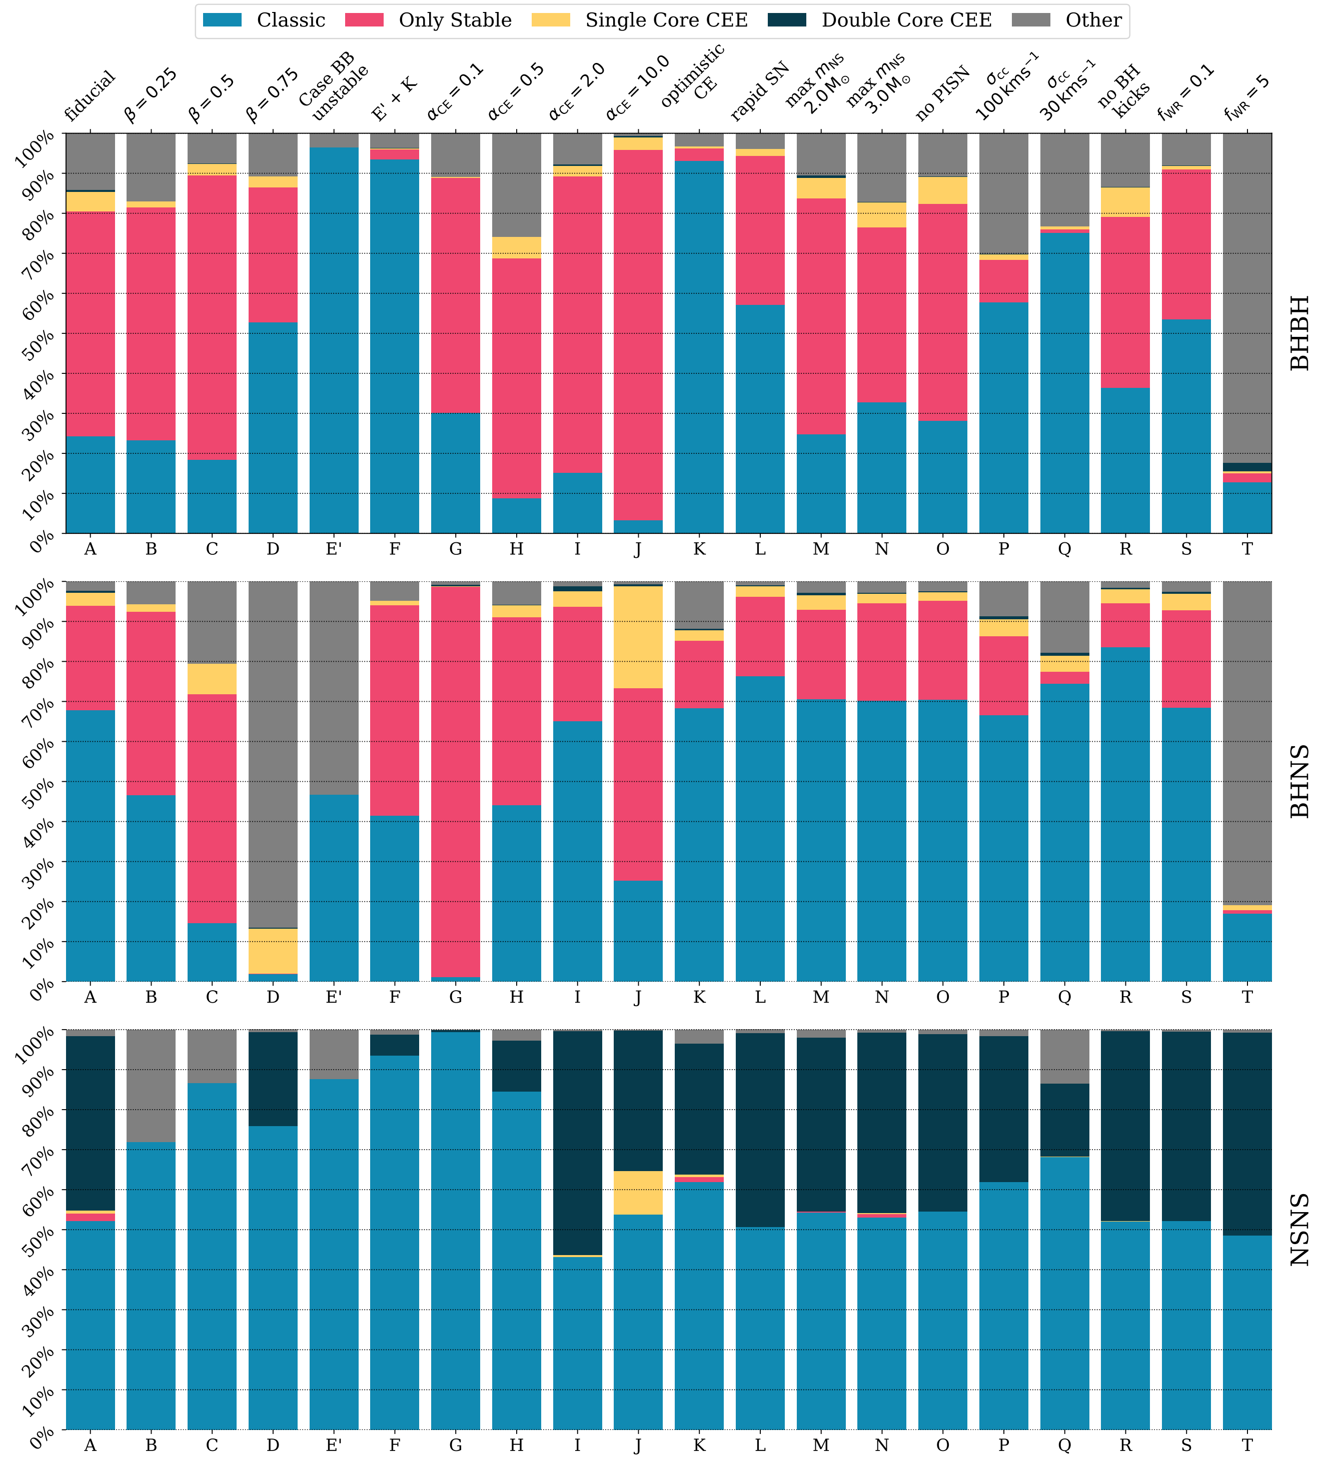

Figure 3.17: Fraction of each DCO type that is formed through different formation channels for all physics variations. Channels are described in detail in Broekgaarden et al. (2021). The classic, single core CEE and double core CEE channels all require at least one common-envelope event whilst only "only stable" consists of only stable mass transfer and "other" contains the remaining binaries which are mainly formed from case A "classic" binaries as well as "lucky" supernova kicks that shrink the binary.

CHAPTER 4

THE ASTEROSEISMIC IMPRINTS OF MASS TRANSFER

You can never know everything, and part of what you know is always wrong. Perhaps even the most important part. A portion of wisdom lies in knowing that. A portion of courage lies in going on anyway.

– al'Lan Mandragoran, True Blade of Malkier

What appliance can pierce through the outer layers of a star and test the conditions within?

– Sir Arthur Eddington

1 Introduction

The majority of stars are born in binaries and multiple star systems (e.g. Duchêne & Kraus, 2013; Moe & Di Stefano, 2017; Offner et al., 2023), a large subset of which will exchange mass at some point in their lifetime (e.g Podsiadlowski et al., 1992; Sana et al., 2012; de Mink et al., 2014). However, mass transfer, both the process itself and the impact it has on the component stars, is still highly uncertain. Specifically, there are large uncertainties in how much mass and angular momentum is transferred and what part is lost from the system (Packet, 1981; de Mink et al., 2007; Renzo & Götberg, 2021) and how the accretor star adjusts to the incoming mass (Hellings, 1983; Braun & Langer, 1995; Cantiello et al., 2007; Staritsin, 2019; Renzo et al., 2023; Lau et al., 2024). These uncertainties in understanding the process of mass transfer result in uncertainties in evolutionary calculations and predictions, such as the rate of formation of close double compact objects (Toonen & Nelemans, 2013; Marchant et al., 2021; van Son et al., 2022b), stripped stars (Crowther, 2007; Heber, 2016; Götberg et al., 2020), X-ray binaries (Fragos et al., 2013), and gravitational wave sources (e.g. Dominik et al., 2015; Belczynski et al., 2002; Broekgaarden et al., 2022; Iorio et al., 2023).

Asteroseismology probes the internal structure of stars through the analysis of stellar pulsations (e.g., Aerts et al., 2010, and references therein) and so may hold the key to directly probing how accretor stars adjust to gaining mass (e.g. Renzo & Götberg, 2021). In particular, high-order gravity (g) mode pulsations carry information about the deep radiative interiors of stars and the boundary between the convective core and radiative envelope. Main-sequence stars that exhibit g-modes include F-type γ Doradus stars, driven by convective flux blocking (Guzik et al., 2000) and Slowly Pulsating B-type (SPB) stars, driven by the κ/γ-mechanism (Waelkens & Rufener, 1985; Waelkens, 1991; Cox et al., 1992; Pamyatnykh, 1999). These g-mode pulsators have been used to provide insights into many aspects of stellar structure and evolution, such as the masses of stellar cores (Johnston, 2021; Pedersen, 2022) internal mixing processes (Pedersen et al., 2018; Michielsen et al., 2021), and angular momentum transport (Aerts et al., 2019; Ouazzani et al., 2020; Salmon et al., 2022; Burssens et al., 2023; Mombarg, 2023; Moyano et al., 2024). Recent work has also suggested that g-modes can be used to probe the binary evolutionary history of stellar mergers in the evolution products of low- and intermediate-mass (Rui & Fuller, 2021) and high-mass stars (Bellinger et al., 2024).

These insights into stellar properties are possible due to the intricate dependence of period spectrum of g-mode pulsations on the size of the convective core, and the chemical composition gradient and structure outside of the core (e.g. Dziembowski et al., 1993; Miglio et al., 2008; Hatta, 2023). Most current works only consider single-star evolution when inferring stellar properties. However, mass transfer can profoundly influence the structure and composition gradients of accreting stars even after thermal re-adjustment, as indicated by numerous studies using 1D stellar evolution codes (Braun & Langer, 1995; Renzo & Götberg, 2021; Miszuda et al., 2021). These changes in structure are usually the result of the rejuvenation and growth of the convective nuclear burning core. As the frequencies of stellar pulsations are finely tuned by the internal structure of stars, asteroseismology holds the potential to identify the signature of previous mass transfer in various classes of pulsating stars. Furthermore, assuming single star evolution for a star that has undergone accretion may result in misleading inferences of its stellar properties from asteroseismology.

Earlier works have explored the asteroseismic modelling of stars in post mass-transfer binaries, with various degrees of accounting for the history of mass transfer. In particular, Guo et al. (2017a,b) and Chen et al. (2021) performed asteroseismic analysis of pulsating stars in post mass-transfer binaries observed by *Kepler*. While Guo et al. (2017a,b) found suitable solutions using only single star evolution models, Chen et al. (2021) compared solutions from binary and single star evolution models. Although Guo et al. (2017b) briefly mention that their solutions may be impacted by not considering the different composition and g-mode cavity resulting from mass transfer, they make no further analysis of these effects. Each of these studies conclude that, for their particular systems, single star models were sufficient.

Miszuda et al. (2021, 2022) investigated the instability of p-modes in two post mass-transfer binaries computed with `MESA`. In their model, the second star accretes nearly conservatively. At the late stages of accretion, the deeper helium-rich layers of the donor star are transferred. The resulting accretor becomes heavily enriched in helium in its outermost layers. They find that this affects overall structure, the mode excitation and pulsation frequencies. In their models, the enriched layer of helium is only present at the surface leading to an inversion of average particle mass (mean molecular weight, see Figure 10 of Miszuda et al. 2021). One would expect that mixing processes, such as the Rayleigh-Taylor instability, thermohaline mixing and also rotational mixing (Kippenhahn et al., 1980; Cantiello et al., 2007), would mix the helium-rich material with the layers below. Both models and observations suggest this overabundance of helium-rich material is erased even before the end of Roche-Lobe overflow (Renzo & Götberg, 2021).

In this work, we study the longer-lived changes to the internal structure of a rejuvenated accretor. We investigate the impact on the asteroseismic signals of g-modes in accretor star. We focus on late B-type main sequence stars, which have a relatively high binary fraction and are commonly observed to pulsate as SPB stars.

We model accretion in a binary system using `MESA` (§2) and demonstrate the difference in evolution and internal structure between an accreting star and an equivalent single star, even when accreting non-enriched material (§3). Using the `GYRE` stellar oscillation code we then show how this influences the period spacing pattern of the accreting star (§4). We highlight how the properties of the star can be inferred inaccurately if single stellar models are used (§5). All code to reproduce the results and figures in this paper is available on GitHub[1] and Zenodo[2]. Interactive versions of several figures are available online[3].

2 Model and numerical setup

In this section we outline the setup of our `MESA` binary model and specify the numeric setup of our `GYRE` calculations.

---

[1] https://github.com/TomWagg/mass-gainer-seismology

[2] https://zenodo.org/records/10011675

[3] https://www.tomwagg.com/html/interact/mass-gainer-asteroseismology.html

2.1 `MESA` model setup

We use Modules for Experiments in Stellar Astrophysics (`MESA`, Paxton et al., 2011, 2013, 2015, 2018, 2019; Jermyn et al., 2023) version r23.05.1 (Paxton, 2023) to simulate non-rotating models for a binary system, as well as a grid of single stars against which to compare. Our full inlists, template folders and our model outputs are available on Zenodo[2].

In particular, the most pertinent settings that we use for this work are as follows: We adopt the Ledoux (1947) criterion to account for the presence of a chemical gradient when determining the stability of convection. We include semiconvection following Langer et al. (1983) with a scaled efficiency of $\alpha_{\rm sc} = 0.1$. We use exponential core overshooting from Herwig (2000), setting $(f, f_0) = (0.01, 0.005)$ (Claret & Torres, 2017), where, in terms of pressure scale height, f determines how far beyond the convective core overshooting extends and f_0 determines how far back into the core the overshooting starts (see Herwig, 2000; Paxton et al., 2011, for more details). We set a minimum diffusive mixing coefficient of $20\,{\rm cm^2\,s^{-1}}$; this smooths out any numerical discontinuities in the composition gradients and partially accounts for the lack of rotational mixing in our models. We motivate our choice of $20\,{\rm cm^2\,s^{-1}}$ and highlight the effect of changing the mixing coefficient on our results in Section 8.1. We do not account for any rotation in our models. For more details on the input physics and settings, see Section 8.2.

Our binary model has a donor with an initial mass of $4\,{\rm M_\odot}$ and an accretor with an initial mass of $3\,{\rm M_\odot}$, such that after accretion we form a star in the typical mass range of SPB stars (Waelkens & Rufener, 1985; Waelkens, 1991; Kurtz, 2022). We chose an orbital period of 5 days such that the donor will fill its Roche-lobe shortly after leaving the main sequence and undergo so-called Case B mass transfer. We account for non-conservative mass transfer with a mass transfer efficiency of $\beta = 0.5$. We allow the secondary to accrete $0.5\,{\rm M_\odot}$, which corresponds to about 30% of the mass that the donor star loses. The remaining mass is lost from the system taking away approximately the specific angular momentum similar to that of the orbit of the accreting star. We only allow accretion of the outer most layers of the donor star, which are not yet enriched in helium. After the mass transfer phase ends and the donor retreats within it Roche-lobe, we further evolve the accretor until central hydrogen depletion.

For comparison, we also evolve a grid of single stars with masses from 2-6 ${\rm M_\odot}$ until the end of helium core burning, using the same physical assumptions. All our models are calculated for a metallicity of $Z = 0.02$. For a demonstration of the numerical convergence with spatial and temporal resolution of our models, see Section 8.3. See also Section 6 for a discussion of the model limitations and caveats.

2.2 `GYRE` setup

We use the `GYRE` stellar oscillation code (version 7.1, Townsend & Teitler, 2013; Townsend et al., 2018; Goldstein & Townsend, 2020; Sun et al., 2023) to calculate the periods of the $\ell = 1, m = 0$ g-modes for both our accreting star and equivalent single star models. We calculate the adiabatic eigenfrequencies for dipole ($\ell = 1$) modes for each model in our `MESA` grid and scan 2000 frequency bins on an inversely sampled grid from 0.25 to $10\,{\rm d^{-1}}$ (equivalent to periods from 0.1 to 4 days). We set the outer boundary condition to use a vacuum boundary condition and solve the full 6th order dimensionless stellar oscillation equations (Dziembowski, 1971; Christensen-Dalsgaard, 2008) using the Colloc scheme `MAGNUS_GL6`. We use the same `GYRE` setup for our single and binary star models. Although excitation physics are an interesting avenue to study in future work, we do not consider non-adiabatic calculations in this work.

3 Binary Stellar Evolution

In this section, we describe the evolution of our model system, both across the Hertzsprung-Russell diagram and in terms of its internal structure.

3.1 Hertzsprung-Russell diagram evolution

In Figure 4.1 we show the evolution across the Hertzsprung-Russell diagram of both the donor and accretor of our binary model, with a subset of our single stellar models in the background. Below we explain these tracks, but we note that binary evolution of this nature has been described in many classic papers (e.g. Morton, 1960; Smak, 1962; Paczyński, 1966; Kippenhahn, 1969; Yungelson, 1973; van der Linden, 1987) and more recent works (e.g. Yoon et al., 2010; Claeys et al., 2011; Eldridge et al., 2013; Tauris et al., 2015; McClelland & Eldridge, 2016; Yoon et al., 2017; Götberg et al., 2017; Renzo & Götberg, 2021).

The evolution of the donor (starting at D1) initially follows the $4\,\mathrm{M}_\odot$ single star track, expanding across the main sequence, exhausting core hydrogen and moving across the Hertzsprung gap. During the expansion on the Hertzsprung gap, at point D2, the donor overflows its Roche-lobe and diverges from the single star track. As it loses mass, it is driven out of thermal equilibrium. It decreases in luminosity as a fraction of the photons produced deep inside are now used to do work to expand the outer layers. The orbit shrinks slightly at first, but quickly starts to widen (Renzo et al., 2019). At point D3 the donor star is still transferring mass but it starts to regain thermal equilibrium. At point D4 the donor has lost its entire hydrogen-rich envelope and contracts within its Roche-lobe. It keeps contracting until it ignites helium at the point marked D5, where it resides a compact subdwarf. As noted in Section 2.1, we limit the accretion onto the companion to $0.5\,\mathrm{M}_\odot$. Although the donor continues to transfer mass beyond this point, none of this mass is accreted by the companion. We mark this point on the donor's evolutionary track with a grey circle and highlight subsequent evolution in grey.

The evolution of the accretor (starting at A1) follows the $3\,\mathrm{M}_\odot$ single star track initially, but early into its main sequence evolution it starts to accrete mass from the donor (at point A2). This drives the accretor out of thermal equilibrium causing it to expand and increase in luminosity. Once we cease the mass transfer (at point A3) the accretor returns to thermal equilibrium at point A4 and proceeds with its main sequence evolution. At this stage it closely resembles the evolution of a $3.5\,\mathrm{M}_\odot$ single star.

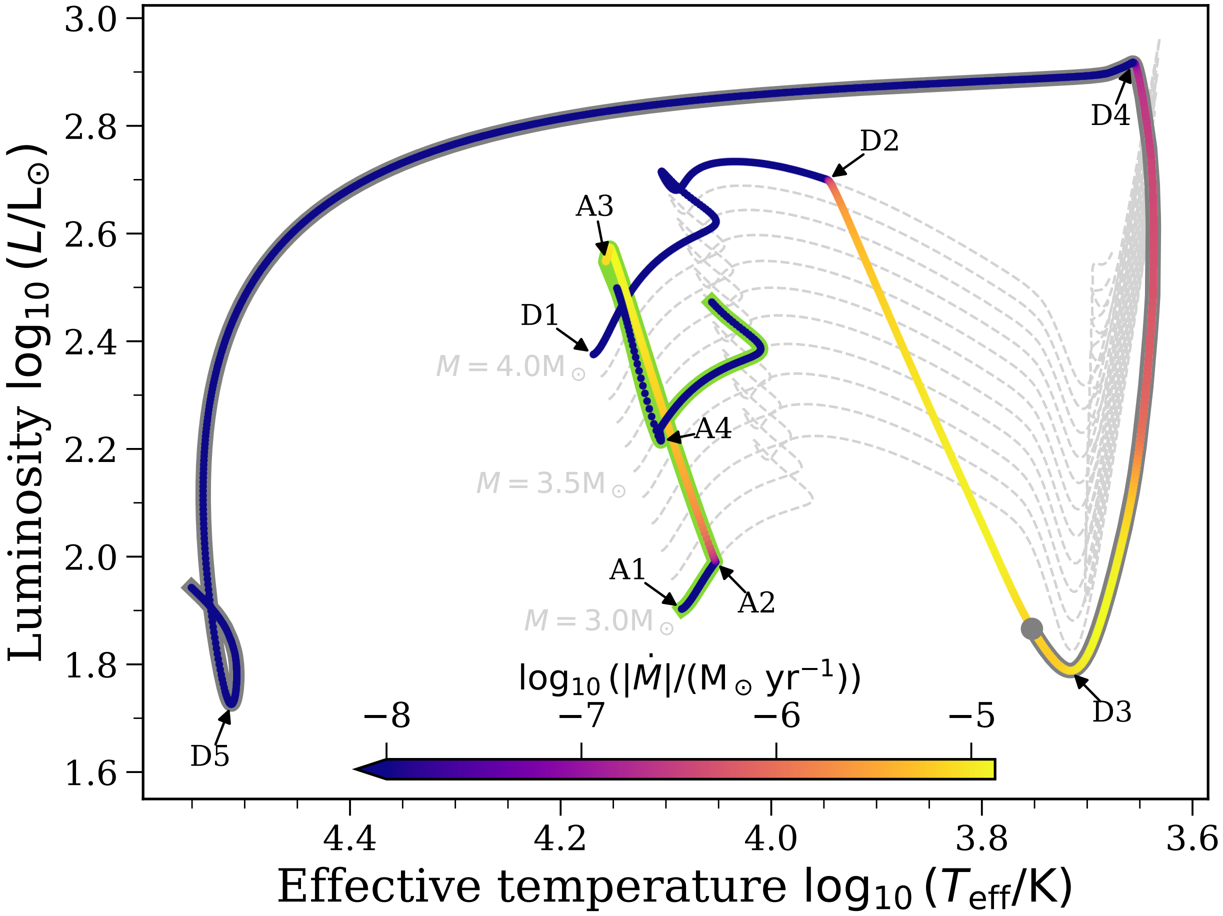

Figure 4.1: Hertzsprung-Russell diagram showing the evolution of the binary. Tracks are coloured by the mass loss and mass accretion rate for the donor (marked D1-5) and accretor (marked A1-4), see Section 3.1 for an explanation of the labels. We limit the accretion on the companion to $0.5\,M_\odot$ and follow its further evolution until core hydrogen exhaustion, see Section 2.1 for details on the technical implementation. The point at which $0.5\,M_\odot$ of material has been accreted is noted with a grey circle and subsequent evolution is highlighted in grey. For reference, we show single star tracks as dashed light grey curves, with masses from 3 to $4\,M_\odot$ in $0.1\,M_\odot$ intervals. Interactive plot available online.

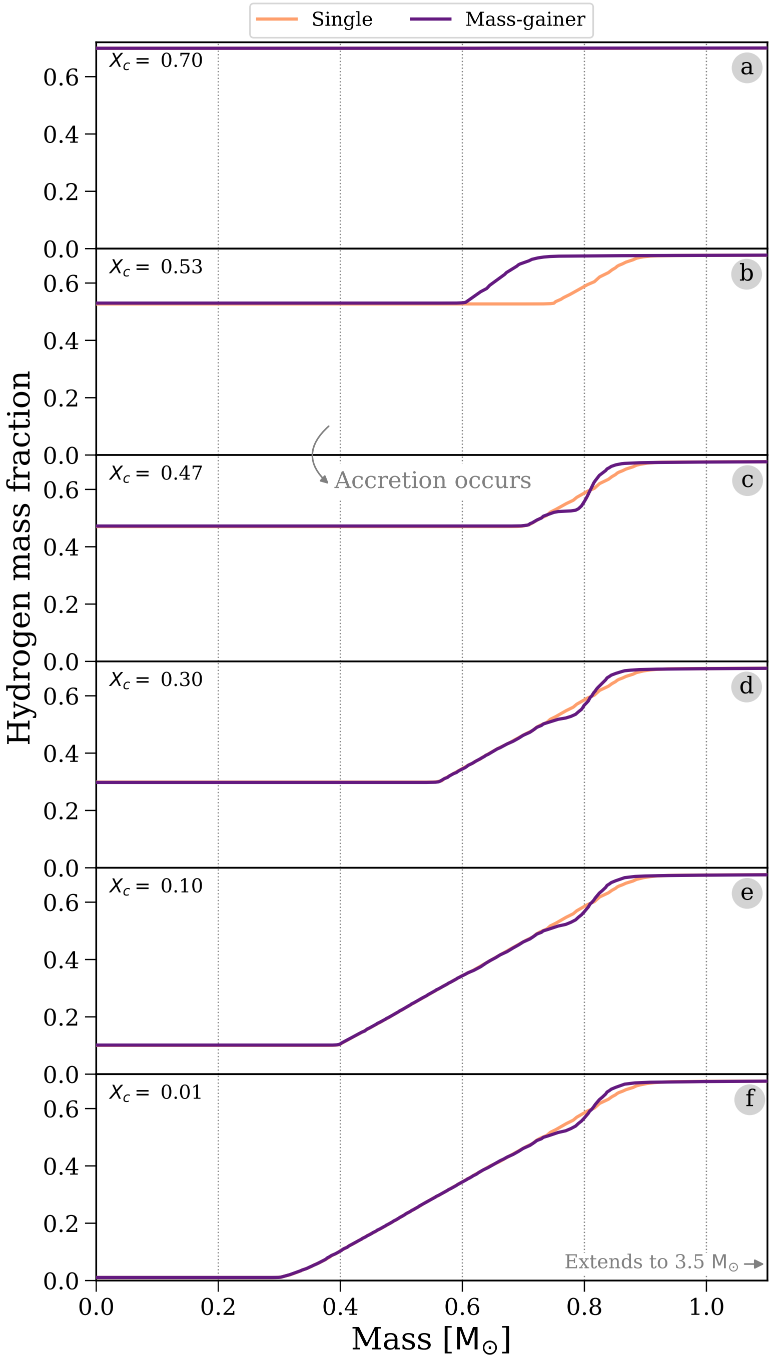

Figure 4.2: Comparison of the hydrogen abundance profiles between (i) an initially $3\,\mathrm{M}_\odot$ star that accretes mass from a companion and (ii) a single star with the same final mass of $3.5\,\mathrm{M}_\odot$. Each panel compares the stars at the same central hydrogen abundance, which is annotated in each panel. Interactive plot available online

3.2 Rejuvenation and chemical gradients

Although the accretor closely follows a more massive single star track in the Hertzsprung-Russell diagram, its *internal* structure has been altered to support the incoming mass after a fraction of its nuclear evolution has already elapsed, leading to enhanced mixing as the convective core expands in mass coordinate (**?**Hellings, 1983; Renzo et al., 2023). This process leaves behind a signature in the hydrogen abundance profile of the star, which we plot in Figure 4.2. In each panel we compare the accretor of our binary model with a 3.5 $\mathrm{M}_{\odot}$ single star, thus the stars have the same final mass. Mass transfer occurs between panels b and c and we discuss the differences below.

First, we consider the evolution of the convective stellar core and the abundance profile for the single star. As the star evolves, it burns hydrogen in its core, decreasing the central hydrogen abundance. The reduced hydrogen abundance decreases the opacity of the core, allowing radiation to travel more freely, leading to a recession of the convective core in mass coordinate (Mitalas, 1972; Crowe & Matalas, 1982; Miglio et al., 2008; Silva Aguirre et al., 2011; Xin et al., 2022). As the core recedes it has a decreasing hydrogen abundance, and therefore it imprints a composition gradient in its wake in the abundance profile. We see these trends in Figure 4.2 as the single star (shown in orange) evolves.

For the mass-gainer the evolution initially proceeds in a similar manner. In panel 4.2b, the shape is similar to that of the single star, though with a smaller convective core due to the star's initially lower mass. Between panels 4.2b and 4.2c, mass transfer occurs. As mass transfer proceeds the accretor increases in luminosity to compensate for the additional mass. This leads to an increase in the convective core size, which one can see as the profiles move outwards in mass coordinates in Figure 4.3. At the same time, this expansion of the core leads to enhanced convective boundary mixing and a rejuvenation of the accretor as more hydrogen is mixed into the core, increasing the central abundance (**?**). The expansion of the core into the region through which it previously receded sharpens the composition gradient, resulting in the 'kink' in the abundance profile relative to the single star for the remaining panels of Figure 4.2. The origin of this feature is shown in Figure 4.3, where we see the hydrogen abundance increase and extend outwards as the core rejuvenates, thus washing away the previous gradient. We emphasise that this occurs even in the absence of rotation and associated mixing. Returning to Figure 4.2, the evolution of the abundance profile after mass transfer proceeds similarly to that of a single star, with subsequent recession of the core and a resulting composition gradient. Critically, however, the feature arising from mass transfer remains throughout the main sequence, albeit marginally smoothed by internal mixing.

4 Asteroseismic Signals

In this Section we demonstrate how the differences in internal structure between the accretor and single star lead to altered asteroseismic signals. We first consider how the Brunt–Väisälä (buoyancy) frequency profile is changed, before showing how this influences the period spacing patterns.

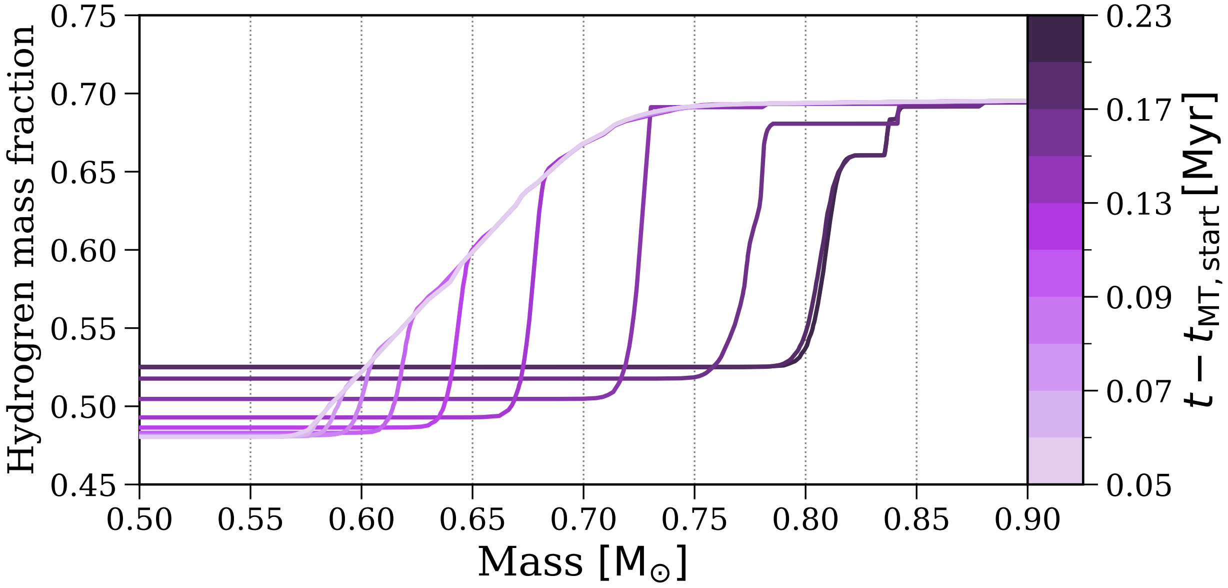

Figure 4.3: Hydrogen abundance profile of our accretor model during mass transfer. Each line is coloured by its time after the start of mass transfer. This plot shows more time-resolved evolution between panels b and c of Figure 4.2.

4.1 Brunt–Väisälä frequency profile

The Brunt–Väisälä frequency (Väisälä, 1925; Brunt, 1927), N, defines the regions in which convective instabilities can occur, such that $N^2 < 0$ indicates a convective region, and $N^2 > 0$ a radiative region in which g-modes can propagate[4]. Physically, it can be understood as the frequency at which a small element of vertically displaced material will oscillate within a radiative region. The Brunt–Väisälä frequency directly determines the period distribution of g-mode oscillations, and thus it is pertinent to consider the impact of mass transfer on it. For an ideal gas, the frequency can be approximated as

$$N^2 \approxeq \frac{g^2\rho}{P}(\nabla_{\mathrm{ad}} - \nabla + \nabla_\mu), \tag{4.1}$$

where ρ is the density, g is the local gravitational acceleration, P is the pressure, $\nabla_{\mathrm{ad}} \approx 2/5$ is the adiabatic temperature gradient and assumed to be a constant, $\nabla \equiv \frac{\mathrm{d}\ln T}{\mathrm{d}\ln P}$ is the temperature gradient and ∇_μ is the chemical composition gradient. Although many of the terms in this expression are similar for our accretor model and equivalent single star model, the density profile and, as noted in Section 3.2, the composition gradient ∇_μ show significant differences and as such we expect similar differences in the Brunt–Väisälä frequency profile.

In Figure 4.4 we compare the Brunt–Väisälä frequency profiles for the accretor star model and single star model with the same final mass. Each panel is for the same central hydrogen content as in Figure 4.2 for a simple comparison. We additionally show the profile as a both

---

[4] The Brunt–Väisälä frequency was originally derived in meteorology and only later applied to stellar evolution. It was first derived in German by Väisälä (1925) and, despite the typical ordering of the names, independently two years later in English by Brunt (1927).

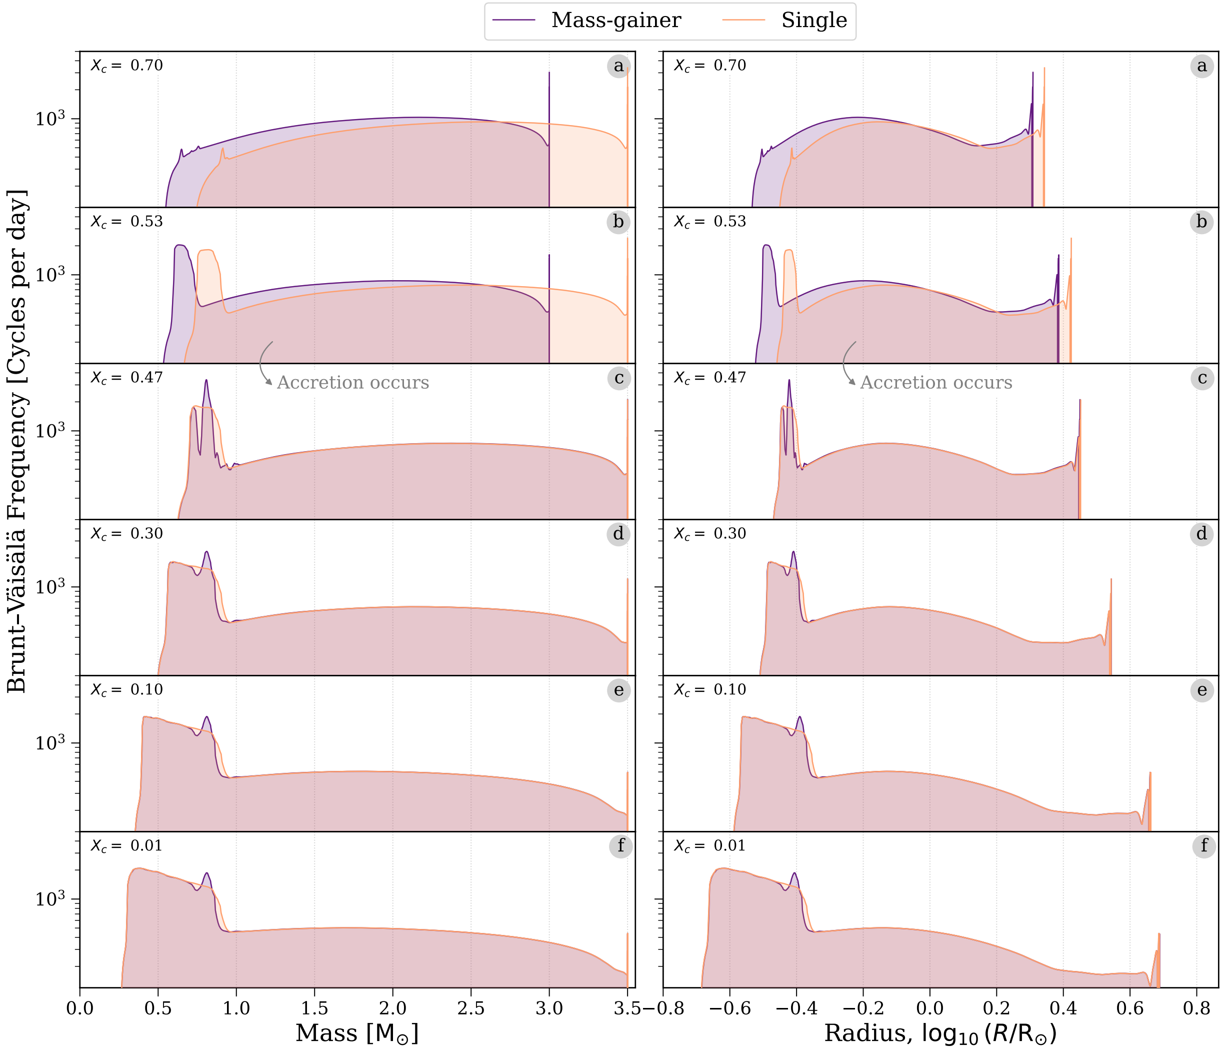

Figure 4.4: As Figure 4.2, but showing the Brunt–Väisälä frequency profile for the same evolutionary timesteps. **Left:** as a function of mass coordinate, **right:** as a function of the radial coordinate. Interactive plot available online

a function of mass coordinate and radial coordinate in the two columns.

Considering first the single star model, we see that initially the convective core ($N^2 < 0$) extends to $\sim 0.75\,\mathrm{M}_\odot$ (or $\sim 0.3\,\mathrm{R}_\odot$) and the frequency profile changes smoothly across the star. As the star evolves, the core recedes, leaving behind a chemical gradient; a peak then emerges in the Brunt–Väisälä frequency profile that extends between the core and the unmixed outer regions of the star. This peak is directly due to the chemical composition gradient (∇_μ) imprinted on the star by the receding core during the main sequence. As the star evolves, the peak extends in concert with the recession of the core, in line with the composition gradient.

For the accretor model, we see similar evolution in panels 4.4a and 4.4b (before mass transfer occurs). Immediately following mass transfer (in panel 4.4c), the convective core radial extent and mass coordinate align with the single star model and several distinct features emerge outside of the core, arising due to the kink in the composition gradient visible in panel 4.2c. As the star evolves, chemical mixing smooths these features to some extent, but importantly the star retains a double-peaked Brunt–Väisälä frequency profile for the rest of its main sequence evolution.

We are unaware of any process in single star evolution that would result in an equivalent Brunt–Väisälä frequency profile for the accretor. In single star evolution we expect a smooth monotonic change in the chemical composition gradient due to the recession of the convective core. Therefore, the profile would always have a smooth, unimodal peak. The occurrence of an increase in convective core size in a rejuvenated accretor results in a change in the chemical composition gradient not possible in single stars (Renzo et al., 2023).

Aside 4.1: The history of the Brunt–Väisälä frequency

The Brunt–Väisälä frequency was originally derived with *Earth's* atmospheric stability in mind and only later applied to astrophysics.

Cloud formation on Earth is strongly linked to stability. An unstable atmosphere (similar a convective region in a star) is much more conducive to cloud formation, as unstable parcels of air can rise quickly and reduce their capacity to hold moisture. So, the measurement of the Brunt–Väisälä frequency in a particular region can provide a useful indicator of the expected rate of cloud formation, which first motivated the derivation of the frequency.

Despite his name appearing second in "Brunt–Väisälä", the derivation was actually first published by Vihlo Väisälä in 1925 (Väisälä, 1925). This Finnish meteorologist explained how wind fluctuations (oscillating with a frequency that he derived) affected pilot observations using weather balloons.

Two years later in 1927, Sir David Brunt, a Welsh meteorologist, published a paper, which performed a similar derivation to arrive at the same frequency (Brunt, 1927). It includes a note acknowledging that he had been made aware of the Väisälä paper but that "[Brunt's derivation] formed a portion of one of a series of lectures on Dynamical Meteorology delivered at the School of Meteorology during the years 1921-1924".

Nowadays, the typical ordering of the names is "Brunt–Väisälä", but given the publication order one might consider whether this is more related to David Brunt publishing in English and being more well known, than from completing the derivation first.

4.2 Period spacing patterns

All differences between the mass-gainer and equivalent single star that we have noted so far are within the internal structure, and so are not directly observable. Therefore, we now consider the impact of these internal structure changes on the observable period spacing pattern.

The period spacing pattern is defined as the difference in period between modes of the same spherical degree, ℓ, and neighbouring radial order, n. Under the assumption of spherical symmetry and high radial order ($n \gg \ell$), this difference is constant and follows the asymptotic g-mode period spacing given by Tassoul (1980):

$$\Delta P_g = \frac{\pi^2}{\sqrt{\ell(\ell+1)}} \left[\int_{r_0}^{r_1} \frac{N}{r} \, \mathrm{d}r \right]^{-1}, \tag{4.2}$$

where ℓ is the spherical degree, N is the Brunt–Väisälä frequency (see Eq. 4.1) and r_0 and r_1 are the boundaries of the g-mode oscillation cavity, which in our model correspond to the convective core boundary and the outer edge of the star respectively.

Deviations from the asymptotic period spacing occur due to abrupt shifts in the Brunt–Väisälä frequency profile, which trap particular modes in certain regions of the star, altering their periods relative to the regular pattern (e.g. Dziembowski et al., 1993; Miglio et al., 2008). The sensitivity of these deviations to the Brunt–Väisälä frequency thus makes the period spacing pattern a useful observable for probing the internal structure of a star (e.g. Aerts et al., 2010).

In Figure 4.5 we compare the period spacing pattern of a mass-gainer to that of an equivalent single star at different stages during their evolution. At the zero-age main sequence (in panel 4.5a), the period spacing pattern closely follows each star's asymptotic period spacing (denoted as dotted lines) due to the lack of any composition gradients. Early during the main sequence, immediately prior to mass transfer (in panel 4.5b), the pattern now displays some oscillation around the asymptotic value due to the chemical composition gradients that have developed outside of the core. Since this is currently pre-accretion, the stars have different masses and thus convective core sizes, resulting in an offset between their asymptotic period spacings.

In subsequent panels (4.5c–f) there are several differences in the period spacing pattern, despite the fact that the stars now have the same mass and convective core size. The two main differences can be expressed in terms of the amplitude and phase of oscillations in the period spacing pattern. Frequently in the stars' later evolution, the amplitude of deviations from the asymptotic spacing are larger for the mass-gainer (for example between ∼0.5–1.0 days in panel 4.5e). This is because the mass-gainer contains regions with steeper chemical composition gradients (Renzo et al., 2023), which more strongly impact the Brunt–Väisälä frequency and thus the period of oscillations.

In addition, we find that the oscillations in the period spacing pattern shift phase in certain regions for the mass gainer. This is most apparent in panel e, in which the patterns are out of phase for periods between ∼1.5–3.2 days and in-phase otherwise. These period-dependent shifts arise due to difference between the Brunt–Väisälä frequency profiles occurring in the region of changing chemical composition. Certain modes are more sensitive to certain regions of the star than others. Modes that are more sensitive to the region of changing chemical composition are shifted and so move out of phase, whilst other modes are less sensitive to the differences from a single star and thus oscillate with the same periods, remaining in-phase.

5 Fitting accretors with single stars

We have demonstrated that, compared to an equivalent single star, a mass-gainer shows significant differences in its period spacing pattern as a result of accretion altering its chemical composition gradient. Therefore, given that around 20% SPB stars are expected to be in interacting binaries (Sana et al., 2012; de Mink et al., 2014), modelling that assumes single star evolution may result in incorrect inferences.

We test how incorrect these inferences may be by fitting the period spacing pattern of our mass-gainer model assuming single star evolution. We use `GYRE` to compute the periods of $\ell = 1, m = 0$ g-modes for our grid of single star models between 2 and 6 $\mathrm{M}_\odot$ across the entire main sequence. We then perform a χ^2 fit for a given mass-gainer period spacing pattern with every single star model, at every timestep.

We match the periods of models *independently* of radial order, since, in reality, the exact radial order of an observed pulsation is not known a priori. This means that the best fitting period for a given radial order in the mass-gainer model may actually be from a different radial order in the single star model. We therefore need to determine the optimal matching

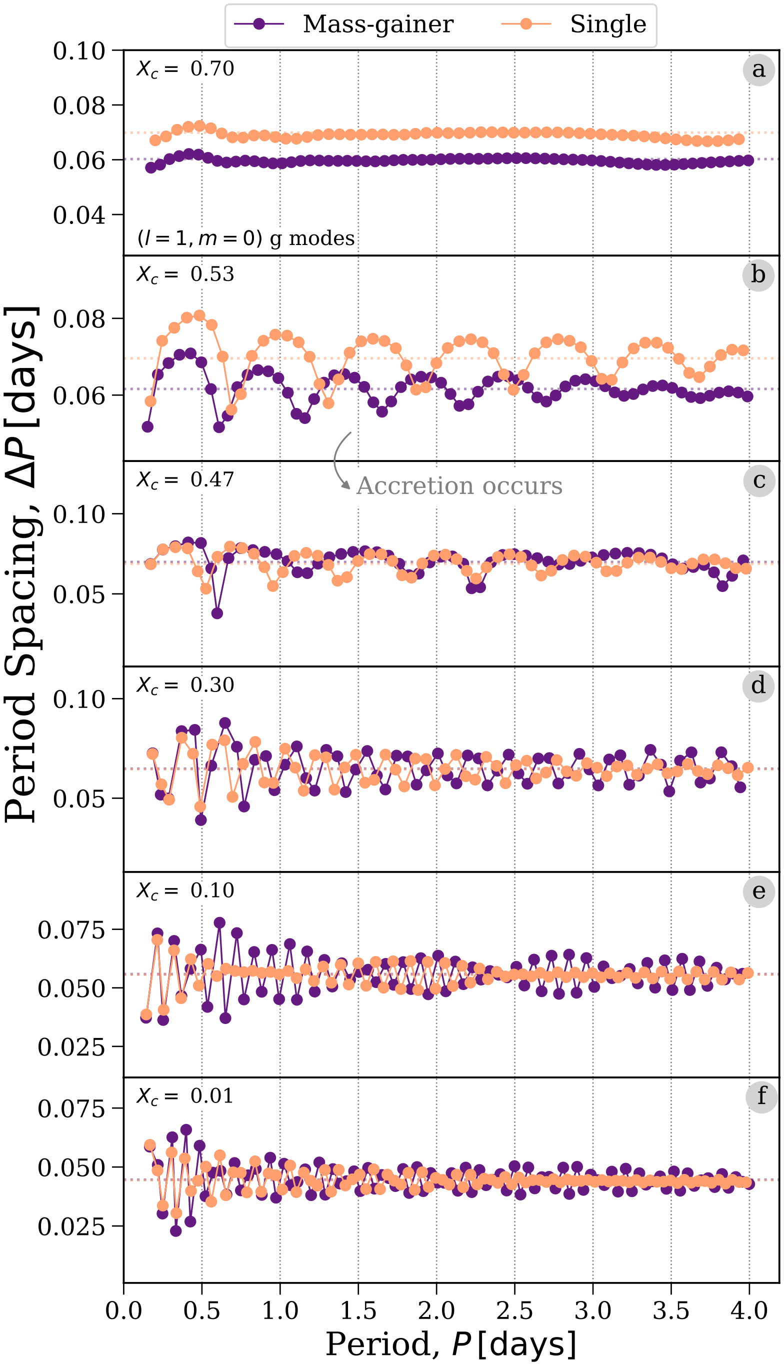

Figure 4.5: As Figure 4.2, but showing the period spacing patterns of the $\ell = 1, m = 0$ g-modes. Asymptotic period spacings for each model are shown as dotted lines. The y-axis limits vary by panel. Interactive plot available online

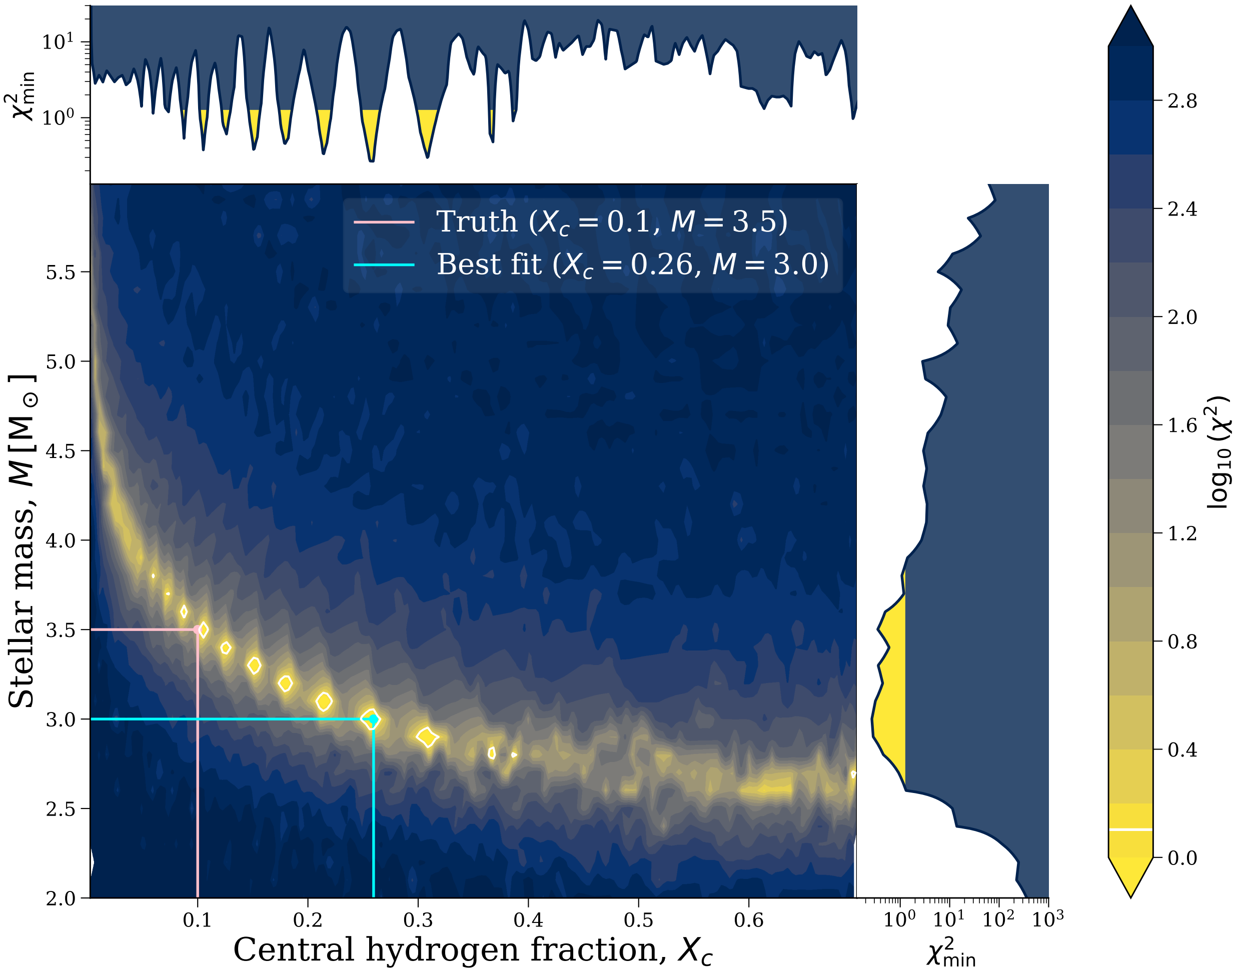

Figure 4.6: Stellar properties could be inferred incorrectly when assuming single star evolution if one does not consider the full multimodal posteriors. Main panel shows the χ^2 values for fitting our $M = 3.5\,\mathrm{M}_\odot$ mass-gainer model at $X_c = 0.1$ with our entire grid of single star models. The true and best fit values are highlighted with lines. We add white contour lines at $\chi^2_{\mathrm{best}} + 1$ to highlight models that are statistically compatible with the true value. Side panels show marginal distributions (the minimum χ^2 marginalised over a given axis), where any values with $\chi^2_{\min} \leq \chi^2_{\mathrm{best}} + 1$ are filled with yellow to mark them as statistically compatible.

between these sets of periods. We make the assumption that the 'observed' period spacing pattern is continuous and monotonic in radial order. This simplifies the matching process, as for each single star model we must now only determine the offset in radial order, ϵ, from the mass-gainer model, such that a period $P_{\mathrm{mg},i}$ in the mass-gainer corresponds to a period $P_{\mathrm{s},(i+\epsilon)}$ in the single star.

In summary, we calculate the χ^2 for each model as

$$\chi^2 = \sum_{i}^{N} \frac{(P_{\mathrm{mg},i} - P_{\mathrm{s},(i+\epsilon)})^2}{\sigma_i^2}, \tag{4.3}$$

where P_{mg} and P_{s} are the g-mode periods of the mass-gainer and single star models respectively, and σ_i is the uncertainty on the measurement. For the purposes of this investigation we adopt a frequency uncertainty of 1/1150 days based on the Kepler time baseline.

In Figure 4.6, we show an example of this χ^2 fitting. In this plot the mass-gainer model has a mass of $3.5\,\mathrm{M}_\odot$ and central hydrogen content of $X_c = 0.1$. The pit-like features in the χ^2 values are the result of degeneracies in the underlying models. Specifically, as the core mass decreases with decreasing X_c, the average ΔP values decrease, thus creating a degeneracy between higher mass models with lower X_c and lower mass models with high X_c (e.g., Buysschaert et al., 2018; Mombarg et al., 2019).

The best fitting model when assuming single star evolution underestimates the mass at $3.0\,\mathrm{M}_\odot$ and overestimates the central hydrogen content at $X_c = 0.26$, more than twice the true value. However, this best-fitting value is found along the trough of degeneracy between mass and hydrogen content. Each value within the white contour lines, which includes the true value, is each statistically compatible with the mass-gainer.

Therefore, when fitting a potential mass-gainer with single star models, it is critical to consider the full multimodal posterior distributions. For example, Basu et al. 2012 suggests selecting all models with likelihoods over 95% of the value of the best-fitting model. This allows one to account for the large degeneracies in models. We highlight that if one were to simply take the mean and standard deviation of the mode of the posterior, then one would infer inaccurate properties for the star (Figure 4.6).

We repeated this fitting procedure throughout the post-accretion evolution of the mass-gainer model, as shown in Figure 4.7. Each column in this Figure shows the same information as Figure 4.6 at a different evolutionary stage. At evolutionary stages recently after accretion has ceased (larger values, close to $X_c = 0.47$), the best-fit values of the mass and X_c from the single star modes are considered statistically compatible with the 'true' values from the accretor model. However, as the main sequence proceeds ($X_c < 0.1$) the lowest χ^2 model diverges from the true value, tending to overestimate X_c and underestimate M. Furthermore, we highlight that the range of the parameter space that is consistent with observations is generally large, and shows multiple optima. This stresses the importance of considering the full posterior distributions.

5.1 Stellar ages

Even when stellar mass and X_c are well fit by single star models, inferences of stellar ages for mass-gainers will be incorrect. A mass-gainer is initially less massive and thus evolves more slowly before accretion occurs. For this reason, our mass-gainer model is ∼72 Myr older than the equivalent ($3.5\,\mathrm{M}_\odot$) single star model for all values of X_c post-accretion (Figure 4.8). The difference in age is dependent on how much mass is accreted and at what stage during the main sequence it is accreted. A star that accreted more mass, or accreted it later in its main sequence, would have a more significant difference in age. Overall, using asteroseismology and single star models to infer the age of mass-gainers will consistently underestimate stellar ages.

5.2 Improved uncertainties

For our assumed frequency uncertainty (based on the Kepler time baseline), we have shown that a single star model with the same stellar mass and X_c is statistically compatible with our mass-gainer model. However, with only moderate improvements (∼15%, see Figure 4.6) to the uncertainties one can rule out this single star model to 2σ.

Therefore, as the quality of data improves with future detectors we find that mass-gainers will be increasingly distinguishable from single star models with the correct mass and X_c. This provides an opportunity to better study the effects of mass transfer and accretion. Yet at the same time, this highlights the increasing importance in considering a star's accretion history when using asteroseismology to infer its stellar properties.

6 Discussion

Binary parameters and the treatment of mass transfer — The mass-gainer model that we use was allowed to accrete only $0.5\,\mathrm{M}_\odot$ of material. This was partially motivated by the idea that an accretor with rotation is expected to quickly reach critical rotation and prevent further accretion (e.g. Petrovic et al., 2005), however others do argue that accretion is possible even beyond critical rotation (Popham & Narayan, 1991; Paczynski, 1991). The amount of accreted material will likely affect the quantitative results. If the star had accreted significantly more, its new convective core may have further expanded, possibly erasing the chemical composition gradient and thus tracers of mass transfer and its asteroseismic signals. If the star had accreted significantly less, the imprints on the structure would be milder and probably harder to detect.

Another parameter is the evolutionary state of the secondary star at the onset of mass transfer (which is governed by the initial mass ratio and, to a lesser extent, the initial period of the binary system). In our model, the accretion occurred when the accretor was still relatively unevolved (at a central hydrogen mass fraction of 0.53). If the secondary had been more evolved, it would had a more pronounced internal chemical gradient, which would have been harder to fully erase. This is not only because the gradient then covers a larger range in mass coordinate, but also since the presence of a steep gradient may have an inhibiting effect to mixing (Braun & Langer, 1995).

Variations on the amount of accreted material and the evolutionary stage of the secondary should be explored systematically in future work.

Rotation — We have limited the scope of our investigations by neglecting rotation in each model. For slow to moderate rotating g-mode pulsators, many recent studies only add the

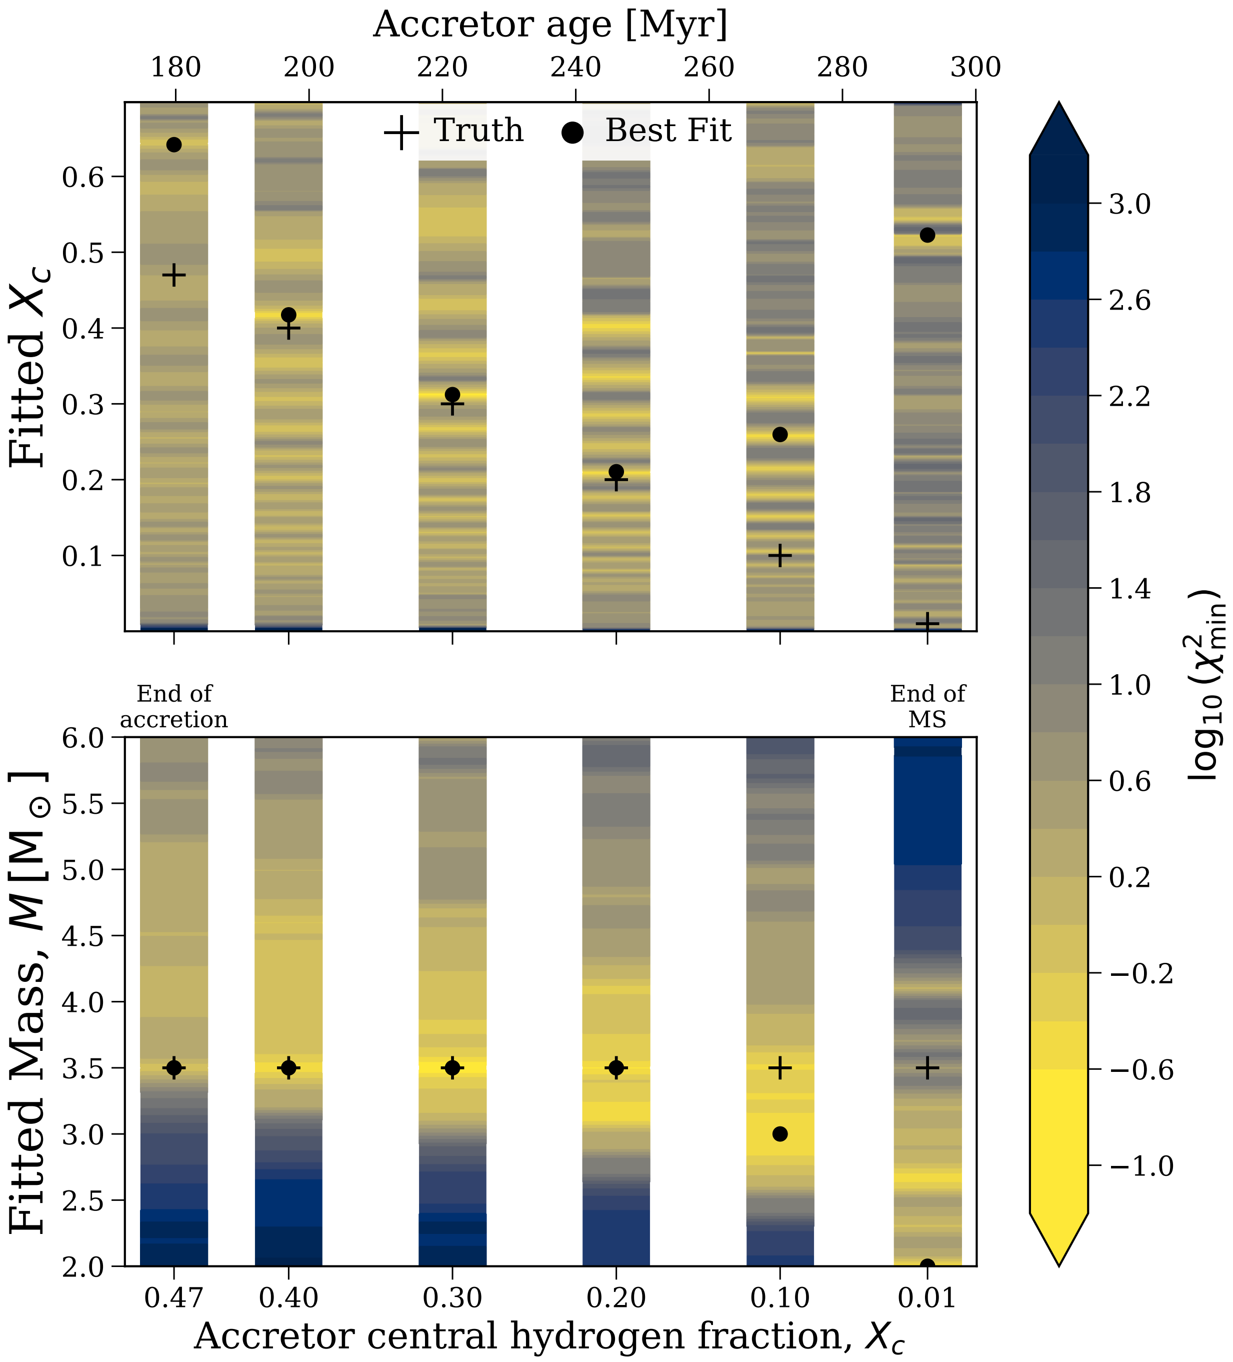

Figure 4.7: Assuming single star evolution for stars that have accreted mass can lead to an overestimation of central hydrogen fraction, X_c, and an underestimation of mass when not accounting for full multimodal posteriors. Comparison between fitted single star models and the 'true' values of the mass-gainer model for X_c (upper panel) and mass (lower panel) as a function of X_c from the end of accretion until the end of the main sequence. Each column shows, for a given X_c, a contour plot with the minimum χ^2 values (similar to Figure 4.6) annotated with markers for the true value and absolute best fit.

Figure 4.8: A single star model with the same mass and X_c as a mass-gainer will underestimate the age. Comparison of the stellar age of our mass-gainer model and $3.5\,\mathrm{M}_\odot$ single star model during post-accretion evolution.

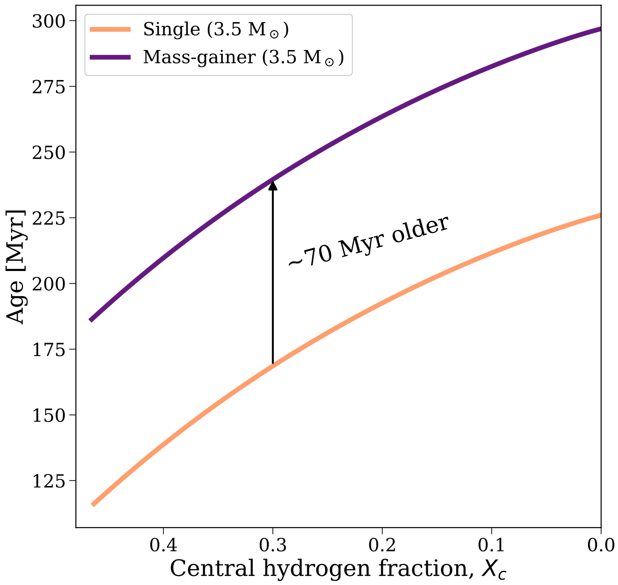

effects of rotation at the stage of calculating pulsation frequencies (e.g. Michielsen et al., 2021). Yet, mass-gainers could reach close-to-critical rotation during mass transfer due to the exchange of angular momentum (Packet, 1981; de Mink et al., 2013; Renzo & Götberg, 2021), and this could impact several aspects of our results, such as the prevalence of the signature of mass transfer in the chemical composition gradient and observed period spacing patter (e.g. Aerts & Tkachenko, 2024, and references therein). In practice, such high rotation rates are not found in observations and accretors typically have spins that are only 10-40% of the critical rate (e.g., Dervişoğlu et al., 2010, see however e.g., Zehe et al., 2018), but this may still significantly impact the asteroseismic signatures.

The two dominant effects that could impact our results concern rotational mixing and the shifting of g-modes due to the introduction of the Coriolis force. Rotation is expected to enhance internal chemical mixing throughout the star through mechanisms such as shear mixing, dynamical instabilities, and Eddington-Sweet circulation (Maeder & Meynet, 2000). These mechanisms are expected to act in the vicinity of the core as well and thus alter the chemical composition gradient, which is the key difference between the mass-gainer and single star models. By introducing the Coriolis force, rotation actively shifts the periods of g-modes (Townsend, 2003), introducing a characteristic tilt in the period spacing pattern (Bouabid et al., 2013). In addition to the effects of modified chemical mixing profile, rotation will further complicate the construction of observed period spacing patterns. However, this does open up the opportunity to investigate the core rotation of stars that have undergone mass transfer, as was done by Guo & Li (2019). Future work should investigate the role of rotation in modifying these predictions.

Treatment of steep chemical gradients — As a result of mass accretion and rejuvenation, accretors develop a steep chemical composition gradient. The treatment of this in a 1D evolutionary code is uncertain and numerically challenging (e.g. Lau et al., 2014). We enforce a minimum level of diffusive mixing in our models that smooths the chemical composition gradient. This has an impact on our results for the Brunt–Väisälä frequency and period spacing pattern. However, as we demonstrate in Section 8.1, differences between the mass-gainer and single star models are still present for a wide-range of choices for this parameter. Crucially, we note that this may artificially suppress physical signatures in the model that result from rejuvenation. Since mass transfer is a rapid and ill-understood phenomenon, we cannot say for certain whether or not all of the signatures that it imparts on the stellar structure are numerical or physical. Specifically, we refer to the points marked as 'glitches' in the Brunt–Väisälä profile in Figure 4.9, which are preferentially smoothed as `min_D_mix` increases. We further note that increased `min_D_mix` values can serve as a proxy for the effects of slow to moderate rotation on chemical mixing in the stellar interior. Thus, while mass transfer may impart some small scale features in the chemical profile of the accretor, we can expect them to be suppressed over time by chemical mixing induced by slow to moderate rotation.

Non-enrichment of accreted material — In previous works it was suggested that differences in the asteroseismic signals of accretors (compared to single stars) could be attributed to an enrichment of the accretor's outer layers as a result of the accretion of helium-rich material (e.g. Miszuda et al., 2021). With our models we show that these difference are found even when non-enriched material is accreted, and are in our case a result of changes in the accretor's internal structure due to rejuvenation. This is particularly encouraging given that one would expect that mixing processes, such as the Rayleigh-Taylor instability, thermohaline mixing and also rotational mixing (Kippenhahn et al., 1980; Cantiello et al., 2007), would quickly mix the helium-rich material with the layers below. Both models and observations suggest this overabundance of helium-rich material is erased even before the end of Roche-Lobe overflow (Renzo & Götberg, 2021).

7 Summary & Conclusions

We have presented new simulations investigating the impact of mass transfer on the asteroseismic signals of slowly pulsating B stars. We used `MESA` to evolve a binary star system and computed the asteroseismic properties of the accretor star using `GYRE`. We compared the internal structure, and rejuvenation, of the accretor, as well as its period spacing pattern, to an equivalent single star. Our main conclusions are as follows:

1. **Mass transfer produces a distinct asteroseismic imprint**
   Our accretor star model shows a significantly different internal chemical composition gradient (e.g. Renzo & Götberg, 2021), and hence Brunt–Väisälä frequency (see Figures 4.2–4.4), even though we only consider non-enriched hydrogen being accreted. While Miszuda et al. (2021) demonstrated that mass transfer impacts the frequencies of p modes in lower mass δ Scuti type stars, they considered helium enriched materials being accreted. Even when considering only un-enriched material, the modified chemical gradient selectively traps certain g-modes and therefore produces a measurably different asteroseismic signal (see Figure 4.5) compared to an equal mass star that has undergone single star evolution.

2. **The asteroseismic signature of rejuvenation persists throughout the main-**

sequence evolution of the mass gainer
At times soon after rejuvenation, there is only a small region of the star with a non-zero chemical composition gradient (e.g. Figure 4.2c). Therefore, fewer modes are shifted by the steep chemical gradient. As the star evolves and the region with a non-zero chemical gradient increases, more modes are trapped by the gradient, and even though mixing modifies the remaining gradient, it is present until the terminal-age main sequence. As such, we can see that the differences in the period-spacing patterns of the mass-gainer and single star persist until the end of the main sequence as well. In particular, we notice that the mass gainer has increased oscillatory amplitude in the period-spacing pattern compared to the single star, likely due to the steeper gradient feature.

3. **Single-star models cannot robustly reproduce period spacing patterns from mass accreting models**
We demonstrate that, given realistic observational uncertainties on pulsation period-spacing patterns, we cannot uniquely identify the matching accretor model when using a single-star evolution model (Figure 4.6). While the individual maximum a-posteriori point estimates correctly identifies the stellar mass and X_c in many cases, the range of models that are statistically valid covers a much wider parameter range in all cases (Figure 4.7). Moreover, even models that correctly estimate mass and X_c significantly underestimate the stellar age (Figure 4.8).

4. **The asteroseismic signature of rejuvenation can be identified with detailed modelling**
While we have demonstrated that single star models are able to recover the bulk properties of mass gaining models, there is still clear structure in the period-spacing patterns that can be modelled. As accretor stars are necessarily in binary systems, this provides a unique opportunity to leverage the high-precision fundamental stellar parameter estimates ($< 1\%$) in the asteroseismic modelling procedure (e.g. Torres et al., 2010; Johnston et al., 2019a; Sekaran et al., 2021). With an independent and precise estimate of the mass gainer stars mass, radius, and age, the asteroseismic analysis will be driven by small difference in the trapping pattern as opposed to matching the asymptotic pattern value (Johnston et al., 2019b). Fortuitously, there are numerous examples of stars that have undergone mass transfer in eclipsing binaries at all masses.

This work demonstrates that asteroseismology can be used to probe the structural impact of mass accretion in binary interaction. These results have immediate implications as this opens the door to providing observational constraints on poorly-understood binary evolution processes that are otherwise unobservable. In particular, there are several classes of known or strongly favoured post mass-transfer systems, such as Algol variables (e.g. Shi et al., 2022) and Be-type stars (e.g. Baade, 1982; Bodensteiner et al., 2020; Labadie-Bartz et al., 2022), that are: i) known g-mode pulsators, and ii) have existing long time-base photometry with telescopes such as Kepler (Borucki et al., 2010) and TESS (Ricker et al., 2015). Future work will investigate the sensitivity of asteroseismic signals to the rate of mass transfer, as well as the mixing processes that accompany mass transfer. Finally, detecting the signatures of mass transfer in currently available asteroseismic data will bridge the exploits of asteroseismology

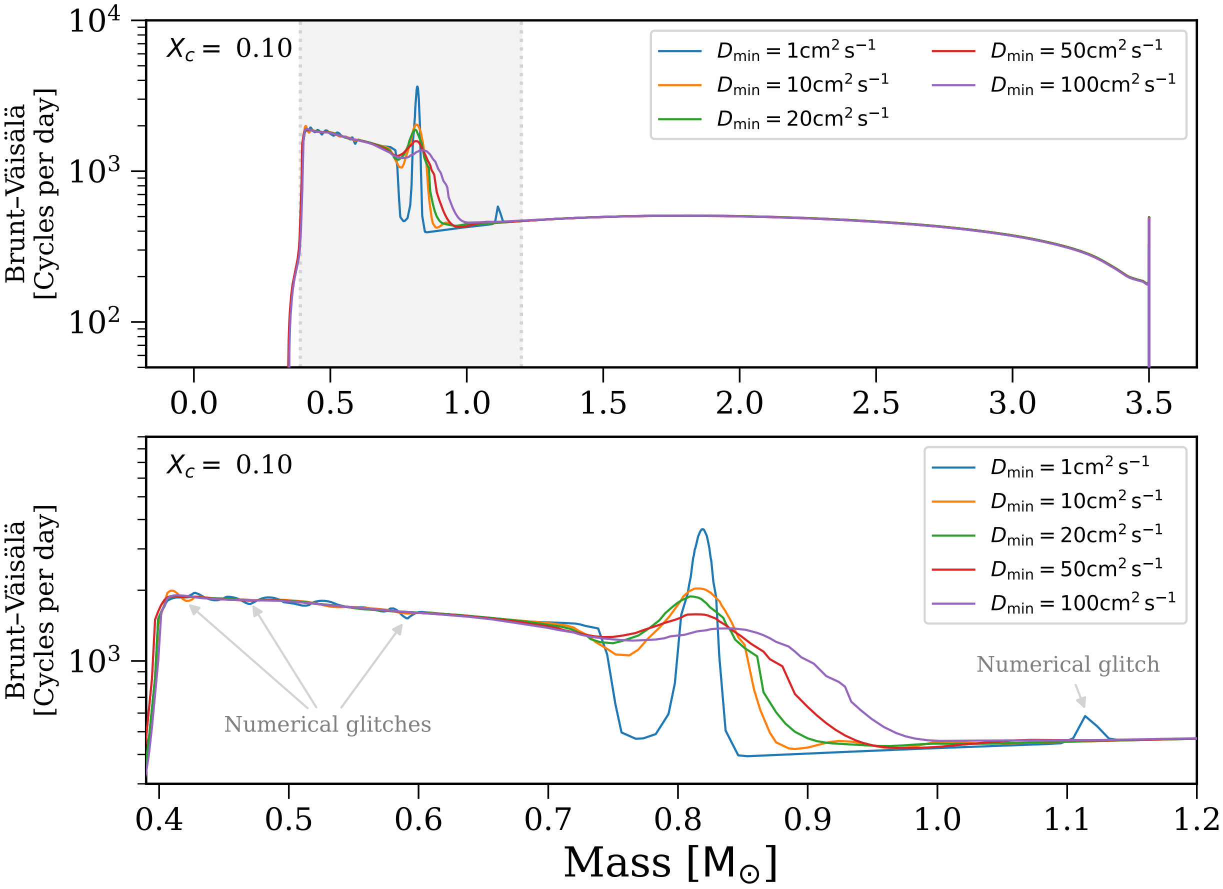

Figure 4.9: Comparison of the impact of changing the `MESA` minimum diffusive mixing parameter, $D_{\rm min}$, on the Brunt–Väisälä frequency profile. Bottom panel zooms in on the highlighted range in the top panel. Annotations highlight numerical glitches in low $D_{\rm min}$ models.

to help calibrate predictions of gravitational wave progenitor populations which require at least one episode of mass transfer in their evolution to interpret the observed distributions of gravitational wave progenitors and events (Abbott et al., 2023; Renzo et al., 2023).

8 Appendix: Numerical considerations

8.1 Importance of choice of minimum diffusive mixing

In our `MESA` models we set a minimum diffusive mixing coefficient, $D_{\rm min}$ (or `min_D_mix` in `MESA`), in order to account for mixing processes not included in our model and to mix over unphysically sharp composition gradients resulting from the 1D approximation of stellar structure. For our models analysed in this paper we set $D_{\rm min} = 20\,{\rm cm^2\,s^{-1}}$, in this Section we explore the impact that this choice has on our results.

We repeated our binary `MESA` simulations for four additional choices of $D_{\rm min}$, ranging from $1 - 100\,\mathrm{cm}^2\,\mathrm{s}^{-1}$. In Figure 4.9 we compare the Brunt–Väisälä frequency profiles of these models at a central hydrogen abundance of $X_c = 0.1$, where the lower panel zooms in on the highlighted region in the upper panel. There are significant differences in the profiles between the different models. As one may expect, lower mixing coefficients lead to steeper composition gradients and therefore sharper features in the Brunt–Väisälä frequency and stronger signals in the period spacing pattern.

However, an overly low choice of D_{min} leads to numerical glitches in the composition gradient and the Brunt–Väisälä frequency profile. These glitches are a result of `MESA` discretising a 3D gradient in a 1D spherical model, which can lead mesh-point wide spikes in the composition gradient and Brunt–Väisälä frequency, even in higher resolution models. We highlight this in Figure 4.9, where glitches are clearly visible in both the $D_{\mathrm{min}} = 1\,\mathrm{cm}^2\,\mathrm{s}^{-1}$ and $D_{\mathrm{min}} = 10\,\mathrm{cm}^2\,\mathrm{s}^{-1}$ models.

We explored a more dense grid of D_{min} models and found that $D_{\mathrm{min}} = 20\,\mathrm{cm}^2\,\mathrm{s}^{-1}$ was the smallest level of mixing that still removed numerical glitches, which informed our selection of this model as the default choice in this paper.

Many processes are expected to induce extra mixing, such as wave mixing and in particular induced rotation in the accretors (Packet, 1981). While we do not attempt to directly and explicitly model these phenomena, the choice of applying a D_{min} to prevent numerical artifacts partially compensates for this.

We stress that the qualitative differences between the mass-gainer and single star in the period spacing patterns remain the same for all choices of D_{min} that we explored. To highlight this point we show the period spacing pattern for the model with $D_{\mathrm{min}} = 100\,\mathrm{cm}^2\,\mathrm{s}^{-1}$ in Figure 4.10. Despite slight differences to the exact shape of the pattern, we still find the same features of (i) stronger ΔP for mass-gainers and (ii) regions in which the period spacing is in-phase and regions in which it is out of phase between the mass-gainer and single star. This confirms that our arbitrary choice of D_{min} does not affect the main conclusions of this study.

8.2 `MESA` input physics

The MESA equation of state (EOS) is a blend of the OPAL (Rogers & Nayfonov, 2002), SCVH (Saumon et al., 1995), FreeEOS (Irwin, 2004), HELM (Timmes & Swesty, 2000), PC (Potekhin & Chabrier, 2010), and Skye (Jermyn et al., 2021) EOSes. Radiative opacities are primarily from OPAL (Iglesias & Rogers, 1993, 1996), with low-temperature data from Ferguson et al. (2005) and the high-temperature, Compton-scattering dominated regime by Poutanen (2017). Electron conduction opacities are from Cassisi et al. (2007) and Blouin et al. (2020). Nuclear reaction rates are from JINA REACLIB (Cyburt et al., 2010), NACRE (Angulo et al., 1999) and additional tabulated weak reaction rates (Fuller et al., 1985; Oda et al., 1994; Langanke & Martínez-Pinedo, 2000). Screening is included via the prescription of Chugunov et al. (2007). Thermal neutrino loss rates are from Itoh et al. (1996). Roche lobe radii in binary systems are computed using the fit of Eggleton (1983). For accretors we include thermohaline mixing once they finish accretion following Kippenhahn et al. (1980) with an efficiency of $\alpha_{\mathrm{thm}} = 1$. We follow the Kolb & Ritter (1990) prescription for the mass transfer rate in Roche lobe overflowing binary systems with an implicit scheme.

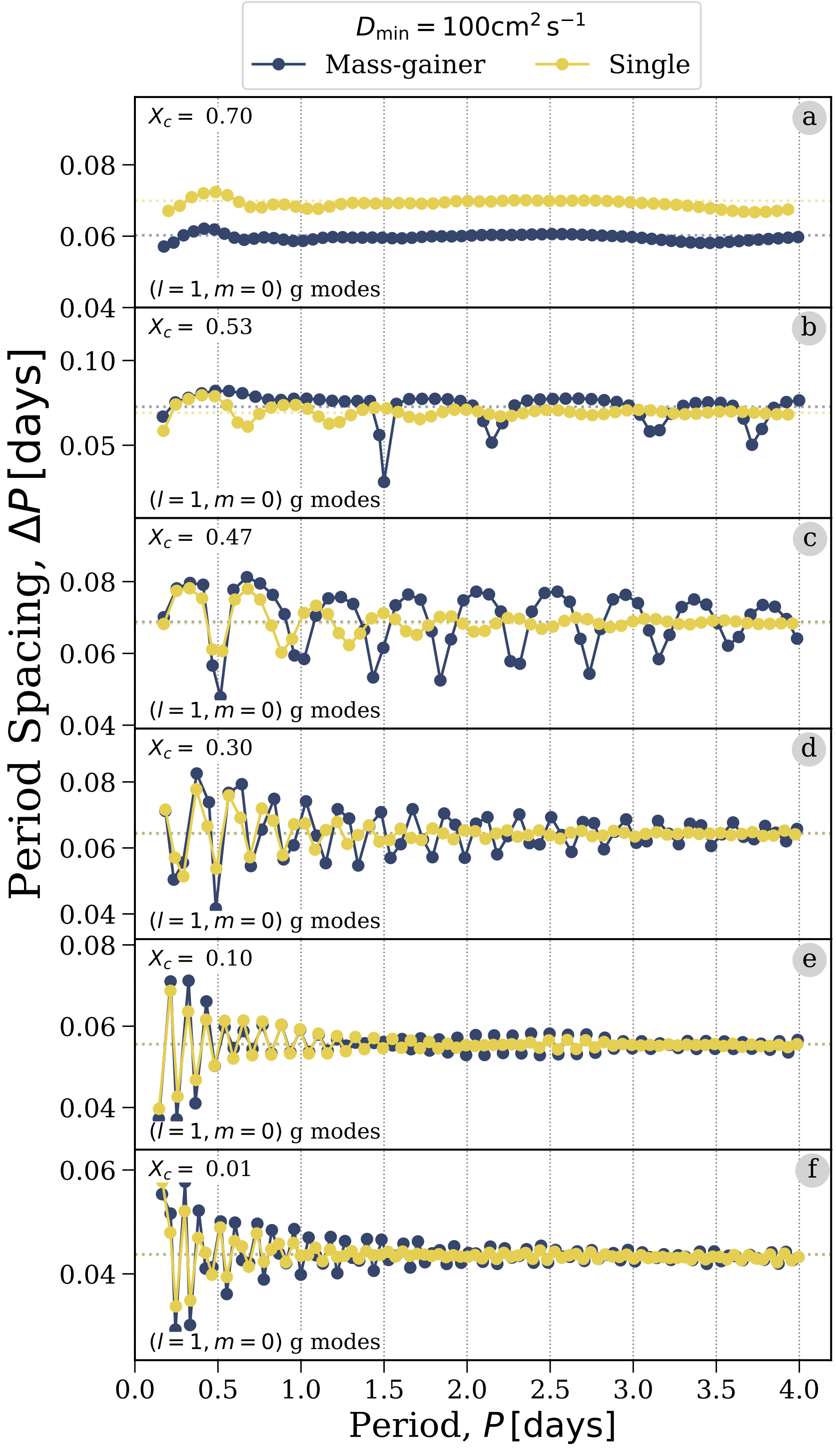

Figure 4.10: As Figure 4.5, but with the minimum diffusive mixing coefficient set to $D_{\min} = 20\,\mathrm{cm}^2\,\mathrm{s}^{-1}$.

8.3 `MESA` & `GYRE` Convergence Tests

We assessed the numerical convergence of our `MESA` models by increasing both the number of timesteps and the number of mesh points. Our default model uses `delta_mesh_coeff` = 0.4 and `delta_time_coeff` = 1.0. We decreased both of these by a factor of two (0.2 and 0.5 respectively) and re-ran the analysis for our mass-gainer model.

In the left panel of Figure 4.11 we show the hydrogen abundance profiles with an additional line for the higher resolution model. There are slight changes from our default model, mainly a small offset in mass coordinate, but critically the kink in the distribution is still present throughout the star's evolution. We additionally compare the period spacing pattern for the default and higher resolution models in the right panel of Figure 4.11. The differences here are mainly in terms of amplitude (except in panel c in which a different set of modes are trapped). However, the main shape and features present are still very similar, and compared to the single star model we still see the same key features, primarily in terms of regions with phase offsets after mass transfer. Therefore, given our results are a proof of principle and not quantitative we confirm that our simulations are sufficiently numerically converged and our findings are robust.

For `GYRE`, as described in Section 4.2, our frequency scanning range is between 0.25 and 10 cycles per day (equivalent to 0.1 and 4 days) with 2000 steps. Our choice of 2000 steps was determined by iteratively increasing the number of steps until the predicted frequencies of the pulsation modes output by `GYRE` converged.

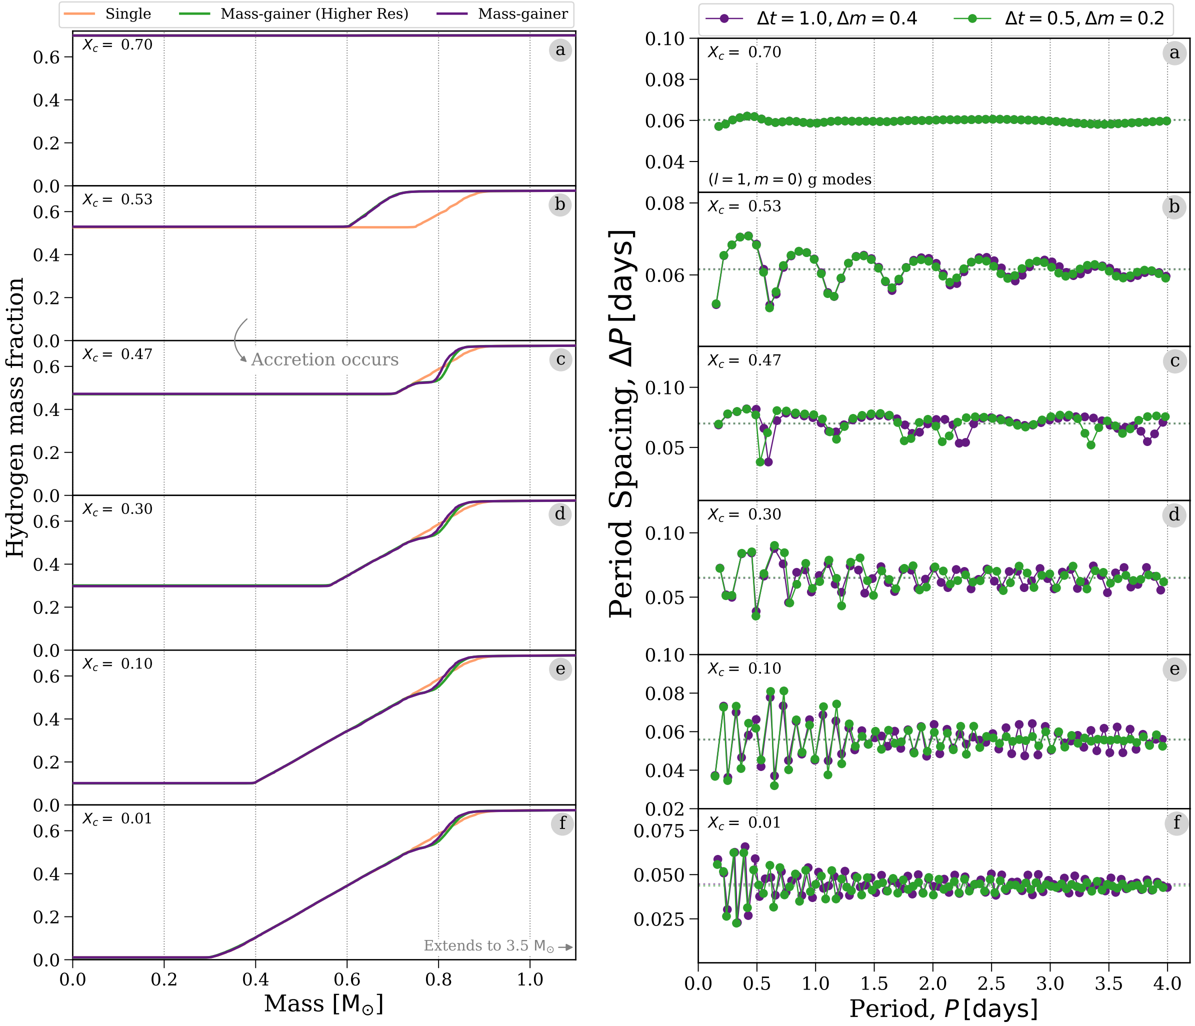

Figure 4.11: Convergence tests. **Left:** As Figure 4.2, but with an additional profile at higher resolution (`delta_mesh_coeff` = 0.2 and `delta_time_coeff` = 0.5) shown in green. **Right:** As Figure 4.5, but now comparing two mass-gainer models at different resolution in timestep and mesh size.

CHAPTER 5

SELF-CONSISTENT POPULATION SYNTHESIS AND GALACTIC DYNAMICS

> *And so, does the destination matter? Or is it the path we take? I declare that no accomplishment has substance nearly as great as the road used to achieve it. We are not creatures of destinations. It is the journey that shapes us. Our callused feet, our backs strong from carrying the weight of our travels, our eyes open with the fresh delight of experiences lived.*
>
> – Nohadon in *The Way of Kings*

1 Introduction

The majority of stars are born in binaries and multiple star systems (e.g., Kroupa, 1995; Duchêne & Kraus, 2013; Moe & Di Stefano, 2017; Oliva & Kuiper, 2020; Offner et al., 2023), a large subset of which will exchange mass at some point in their lifetime (e.g Podsiadlowski et al., 1992; Sana et al., 2012; de Mink et al., 2014). These massive stars play a critical role in the formation and evolution of galaxies as a result of their feedback (e.g., Dekel & Silk, 1986; Hopkins et al., 2012; Nomoto et al., 2013; Somerville & Davé, 2015; Naab & Ostriker, 2017).

However, binary evolution remains uncertain, with many parameters such as common-envelope efficiency, mass transfer efficiency, angular momentum loss due to mass transfer and the mean magnitude of supernova natal kicks unconstrained across several orders of magnitude (e.g., Janka, 2012; Ivanova et al., 2013; Katsuda et al., 2018; Ivanova et al., 2020; Röpke & De Marco, 2023; **?**).

Single massive stars are not expected to migrate far from their birth location before reaching core-collapse due to their short lifetimes ($\lesssim 50\,$Myr, e.g., Zapartas et al., 2017). However, binary stars may disrupt after an initial supernova event, ejecting the secondary star from the system with its orbital velocity (e.g., Blaauw, 1961; Eldridge et al., 2011; Renzo et al., 2019). Thus, close massive binaries that disrupt can lead to the displacement of secondary stars significantly farther from star-forming regions. The present-day positions and kinematics of massive stars and binary products are therefore strongly impacted by changes in binary physics that alter the separation prior to supernova. This means that comparing simulations of positions and kinematics of stars and compact objects to observations will enable constraints on binary stellar evolution parameters.

The use of positions and kinematics as tracers of binary evolution has been considered in the past. Recent work has shown the importance of accounting for the galactic potential, which can change the velocity of kicked objects (e.g., Disberg et al., 2024). It is also important to consider the inclination or timing of a supernova kick relative to the galactic orbit, since, for example, a kick out of the galactic plane at an object's highest galactic vertical position will have a strong effect on its final position. Failing to consider impacts from both a galactic potential and kicks (i.e. velocity impulses) will lead to misleading conclusions regarding the final spatial distributions of the population. Some studies have considered using the Galactic potential at the present-day position of objects to place a lower limit on the peculiar velocity at birth and constrain supernova kicks (Repetto et al., 2012; Repetto & Nelemans, 2015; Repetto et al., 2017; Atri et al., 2019), but the accuracy of this method is debated (Mandel, 2016). Other work has considered the impact of the Galactic potential for individual special cases, rather than at a population level. For example, Evans et al. (2020) considered the orbits of hyper-runaway candidates evolving through the Milky Way potential, whilst Neuhäuser et al. (2020) developed software for tracing the motion of stars to investigate the recent nearby supernovae that ejected ζ Ophiuchi. Additionally, Andrews & Kalogera (2022) considered galactic orbits of synthetic populations to place constraints on black hole natal kicks based on observations of a microlensed black hole. Previous work has also worked towards connecting population synthesis and galactic dynamics. Through a combination of the `COMPAS` population synthesis code (Riley et al., 2022) and `NIGO` galactic orbit integration tool (Rossi, 2015) to make predictions for binary neutron stars, black hole neutron star binaries and pulsars (Chattopadhyay et al., 2020, 2021; Song et al., 2024). Moreover, the `BiPoS1` code accounts for the effect of dynamical encounters on internal binary orbits and can be used to synthetise a galactic stellar population (Dabringhausen et al., 2022).

In this paper we present `cogsworth`, a new open-source tool for self-consistent population synthesis and galactic dynamics simulations. `cogsworth` provides the theoretical infrastructure for making predictions for the positions and kinematics of massive stars and compact objects, placing these systems in the context of their host galaxy and its gravitational potential. The code is applicable to a wide range of binary products, both common and rare, from walkaway and runaway stars to X-ray binaries, as well as gravitational-wave and gamma-ray burst progenitors.

The paper is structured as follows: in Section 2 we explain the functionality of `cogsworth` and describe its primary features and capabilities. We demonstrate these capabilities in a series of example use cases in Section 3. We use `cogsworth` to predict the spatial distribution of compact objects and runaways in both dwarf and Milky-Way-like galaxies. Using a cluster from a hydrodynamical simulation, we show how supernovae can change the orbits of stars in several ways. We predict the separation of disrupted binary stellar companions on the sky, as well as create a synthetic *Gaia* colour-magnitude diagram. In Section 4, we discuss the current limitations of the package and we outline planned additional future developments in Section 5.

2 `cogsworth`

`cogsworth` is a code that combines binary population synthesis simulations (via `COSMIC`, Breivik et al., 2020b) with galactic dynamics (via `Gala`, Price-Whelan, 2017) to self-consistently use stellar and orbital evolution to rapidly derive present day positions, kinematics and demographics for complete populations of binary stars and their descendants.

Our code is fully open-source and openly-developed (available on GitHub[1]), pip installable (`pip install cogsworth`) and indexed on Zenodo. In this paper we describe v2.0.0 of the code. We wrote `cogsworth` in Python to make it convenient and accessible, but its core dependencies are written in Fortran and C (via `COSMIC` and `Gala` respectively) for efficiency. We use automated testing via a detailed suite of unit tests, with full code coverage. Additionally, we have written thorough documentation of `cogsworth`, including ∼20 tutorials covering full usage of the code, several short examples and a series of longer in-depth case studies, all of which is available online[2].

We describe the specific capabilities in the following subsections, and illustrate an overview of the code in Figure 5.1. The first subsections focus on core functionality of `cogsworth`, which is accessed via a `Population`, with which one can:

§2.1 Sample initial galactic positions, velocities, birth times and metallicities from a star formation history model

---

[1] `https://github.com/TomWagg/cogsworth`

[2] `https://cogsworth.readthedocs.io`

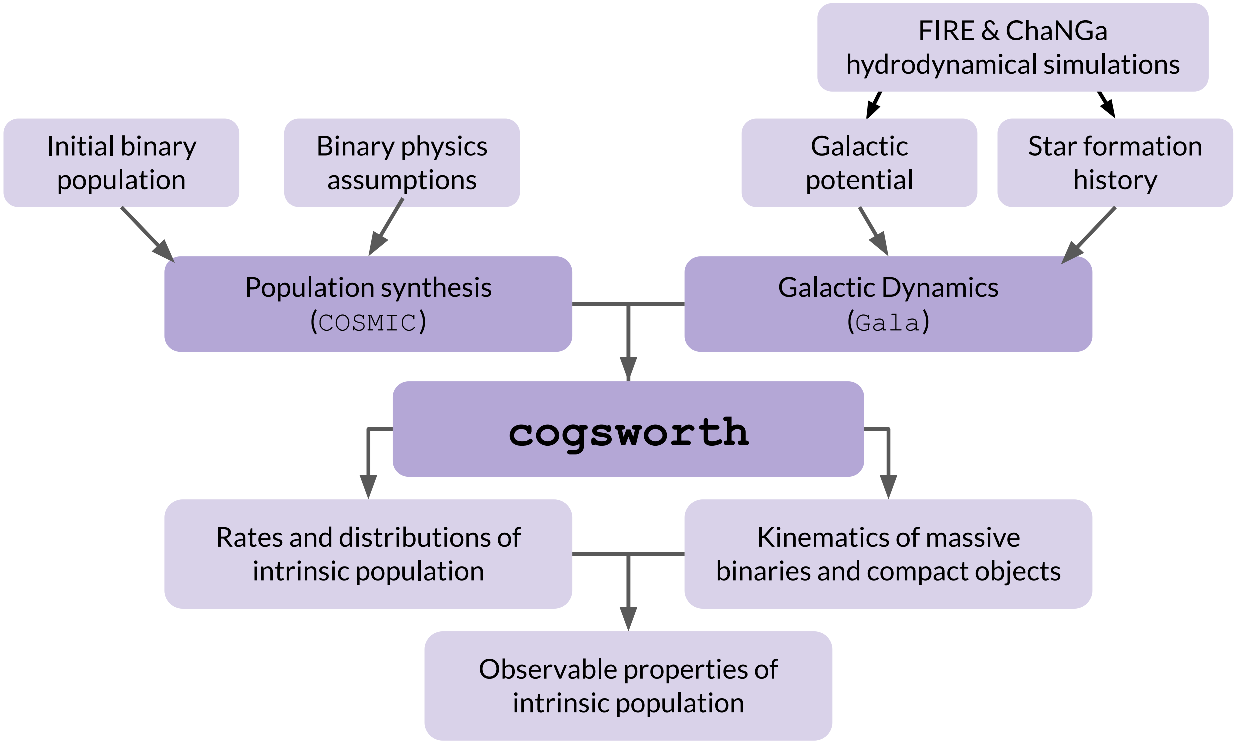

Figure 5.1: `cogsworth` combines population synthesis and galactic dynamics self-consistently to produce rates, distributions, observables and kinematics of massive binaries and compact objects. A schematic of the input options and outputs of `cogsworth` simulations.

§2.2 Sample and evolve a (binary) stellar population until present day

§2.3 Integrate the orbits of each binary through the galaxy, accounting for supernova kicks and disruptions

§2.4 Identify observable constituents of the present day intrinsic population

Each of these features are flexible and can be tuned to a particular use case. In addition, in §2.5 we describe how one can alternatively use `cogsworth` to initialise and evolve populations based on hydrodynamical zoom-in simulations. In the later subsections we detail `cogsworth`'s visualisation functionalities (§2.6), details of its typical runtime (§2.7) and data storage (§2.8), and its ability to create custom citation statements for a given simulation (§2.9).

2.1 Galactic star formation histories

A galactic star formation history (SFH) defines the distributions of times, locations and metallicities at which stars are formed in a galaxy. In `cogsworth`, one can flexibly adjust the SFH of a simulation with a `StarFormationHistory` Python class. By default, `cogsworth` uses `Wagg2022`, an empirically motivated analytic model for the Milky Way (Wagg et al., 2022c). This model contains a low-[α/Fe] (a.k.a. “thin") disc, high-[α/Fe] (a.k.a. “thick") disc, and bulge, each with their own spatial and temporal distributions. Metallicity is calculated as a function of birth time and galactocentric radius. This SFH accounts for effects such as the inside-out growth of the galaxy and radial migration (Frankel et al., 2018, 2019). We explore this SFH further in Section 3.4 and the full details of the model are given in Section 2.2.1 of Wagg et al. (2022c).

Beyond the default SFH, we include simpler parameterised SFHs, such as `BurstUniformDisc`, in which stars are formed in a single burst of star formation with a fixed metallicity in a uniform disc. Additionally, we include action-based SFHs using `Agama` (Vasiliev, 2019), such as `QuasiIsothermalDisc`, which represents a quasi-isothermal distribution function for the Milky Way disc as described in Sanders & Binney (2015).

Each SFH class is designed to be modular and flexible and as such, they can be entirely customised. Users can overwrite individual distributions in a given SFH (such as changing the birth time distribution of the bulge component), or define their own entirely custom SFH. We explain how to accomplish this in a tutorial.

For each binary or single star, i, in a population, we use the SFH to draw the initial galactic parameters

$$g_i = \{\tau, R, z, \phi, Z\}, \tag{5.1}$$

where τ is the lookback time (the time before the present day when the system formed), R is the initial galactocentric radius, z is the initial height above the plane, ϕ is the azimuthal angle, and Z is the metallicity. For an SFH that doesn't explicitly define a distribution for galactocentric velocities, we assign the initial galactocentric velocity of a system, i, as follows

$$\vec{\mathbf{v}}_i = \vec{\mathbf{v}}_{\rm circ}(R_i) + \vec{\mathbf{v}}_{\rm disp}, \tag{5.2}$$

where $\vec{\mathbf{v}}_{\rm circ}$ is the circular velocity for the population's galactic potential and $\vec{\mathbf{v}}_{\rm disp}$ is an isotropic velocity dispersion, which is an input option that by default has a magnitude of $5\,{\rm km\,s^{-1}}$.

2.2 Stellar population sampling and evolution

`cogsworth` uses the open-source and community-driven rapid binary population synthesis code `COSMIC` to perform the sampling of the initial (binary) stellar population and the (binary) stellar evolution (Breivik et al., 2020b). `COSMIC` uses fitting formulae based on single stellar tracks originally developed for the Binary Stellar Evolution (`BSE`) code (Tout et al., 1997; Pols et al., 1998; Hurley et al., 2000, 2002) and allows the user to rapidly sample and evolve populations of binaries with a variety of physics assumptions. With `COSMIC` a user can specify dynamic time resolution conditions for its outputs based on binary parameters which, for example, could be used to increase the number of outputted timesteps once a star is stripped, or during mass transfer onto a compact object to investigate X-ray binaries. One can also easily access the initial conditions of a population evolved with `COSMIC`, allowing for convenient reproducibility of simulations. This is additionally useful for re-running identical initial populations with alternative binary physics settings to ascertain how the evolution changes with different settings. Each of these features are directly inherited by `cogsworth`.

When sampling an initial binary stellar population, a user can specify their choice of initial mass function (IMF) for drawing primary masses, a binary fraction and distributions of initial orbital period, eccentricity and mass ratio for a given `Population` of binaries (see Section 2.1.1 of Breivik et al., 2020b for all available options). Metallicities are set based on the chosen SFH model (see Section 2.1). Using these distributions, one can either draw a fixed number of systems, or specify a total mass to sample. This tutorial explains in detail how to change these settings.

`cogsworth` evolves this initial binary population from their individual birth times until present day with a user-specified choice of binary physics. Any binary stellar evolution parameter that can be supplied to `COSMIC` can also be specified in a `Population` in `cogsworth`. These parameters cover a range of binary physics including stellar wind mass and accretion, mass transfer through Roche-lobe overflow, common-envelopes, supernova kicks, remnant mass prescriptions and tides. For a full list of the available parameters see the `COSMIC` documentation and learn about changing them in `cogsworth` with this tutorial.

2.3 Galactic orbit integration

`cogsworth` applies the galactic dynamics package `Gala` for the galactic orbit integration of binaries (Price-Whelan, 2017). This package allows users to integrate the orbits of sources rapidly with user-friendly functions wrapped on low-level code (primarily C) for fast computations. One can choose from numerous, flexible potentials (or even define custom potentials) through which to integrate orbits. `cogsworth` uses `Gala` to integrate the full orbit of each binary in a population through a given galactic potential. By default, `cogsworth` uses the `MilkyWayPotential2022` potential, which is fit to observations of the Milky way rotation curve, the shape of the phase-space spiral in the solar neighbourhood and a compilation of recent mass measurements of the Milky Way (Eilers et al., 2019; Darragh-Ford et al., 2023). `cogsworth` users can apply any model for the galactic potential that is available in `Gala`. Each model has flexible parameters (such as mass and scale height) and can be combined into composite potentials, such as combining an NFW profile (Navarro et al., 1997) and a Plummer potential (Plummer, 1911). Additionally, one can fit an arbitrary galactic mass distribution using the self-consistent field method implemented in `Gala` based on Hernquist & Ostriker (1992) and Lowing et al. (2011). This tutorial describes how one can change the galactic potential used in `cogsworth` simulations.

For systems that experience supernovae, `cogsworth` accounts for the resulting changes in velocity. `COSMIC` logs the velocities imparted by Blaauw (Blaauw, 1961) and natal kicks (Katz, 1975; Janka, 2012, 2017) and whether a binary is disrupted by a supernova (e.g., Renzo et al., 2019). In each case, `cogsworth` transforms the resulting velocities to galactocentric coordinates (uniformly sampling a random orbital phase, θ, of the binary and inclination, ι, of the binary relative to the galaxy) and updating the orbit of the system. In the case of disruptions, a second orbit is produced for the secondary, tracking the binary position until the disruption and then the subsequent motion of the secondary.

Overall, this allows users to track the location of either star in a binary system at *any* point in its evolutionary history. This can be used to, for instance, predict the location of supernovae or track the sites of r-process enrichment from binary mergers.

2.4 Observables estimation

Key to applying `cogsworth` to realistic problems and constraining our models is being able to compare simulations to observations. In this Section, we explain how users can transform intrinsic `cogsworth` populations into observables.

2.4.1 Electromagnetic observations

We have implemented functionality to translate intrinsic stellar parameters in `cogsworth` populations (such as mass, luminosity and galactic position) into observables (such as fluxes and colours). Currently, `cogsworth` focuses on producing predictions for *Gaia* observables, but we intend to build on this with other instruments in future (see Section 5.3).

`cogsworth` can compute the magnitude of sources in arbitrary filters by applying bolometric corrections and dust extinctions, achieved through a combination of the `dustmaps` and `isochrones` packages (Green, 2018; Morton, 2015). We match isochrones based on the effective temperature, surface gravity and [Fe/H] metallicity of the evolved stars. We note that the MIST isochrones do not vary [α/Fe] and as such we effectively assume a solar abundance pattern in these predictions (Dotter, 2016; Choi et al., 2016). You can follow this tutorial to learn how to compute photometry for `cogsworth` sources. Within the Milky Way, the interplay between distance, the 3-dimensional dust distribution, and the *Gaia* scanning pattern leads to a complex selection function, but one that can be captured through the empirical selection function made available through `gaiaunlimited` (Cantat-Gaudin et al., 2023). `cogsworth` is therefore capable of predicting whether a given source (either a bound binary or star from a disrupted binary) would be detectable by *Gaia*. This tutorial explains how to make predictions about the observable *Gaia* population.

2.4.2 Gravitational waves

In addition to electromagnetic observations, we consider gravitational wave detections from the inspiral of double compact objects. Stellar-mass binaries in the Milky Way will be detectable by LISA via millihertz gravitational wave emission (Amaro-Seoane et al., 2017). The `LEGWORK` package allows users to compute gravitational-wave strains and SNRs for binaries, and calculate the evolution of binary separations and eccentricities due to gravitational wave emission (Wagg et al., 2022a,b). We connect `cogsworth` to `LEGWORK`, allowing users to quickly calculate the LISA SNR of each binary in a population, as well as the time until its merger with a single, simple function call. This tutorial shows an example of calculating LISA SNRs for a `cogsworth` population of double white dwarfs.

2.5 Building off hydrodynamical zoom-in simulations

Using an analytic model for a galactic SFH can work well for longer-lived populations, which can be expected to have erased all memory of their initial positions. However, it is unrealistic for younger stars, and in particular short-lived massive ones, which should retain significant initial spatial clustering and correlations with the surrounding ISM (e.g., Sarbadhicary et al., 2023). Such correlations are particularly important when attempting to constrain aspects of binary physics by comparing predicted present-day locations from `cogsworth` models to observations of recently formed binaries.

Motivated by a need for more detailed initial spatial clustering, we include the option to initialise a `cogsworth` population using hydrodynamical zoom-in simulations. These simulations are not only used to set the locations and times of star formation, but also the galactic gravitational potential.

2.5.1 Compatible simulations

We currently support the post-processing for two different suites of hydrodynamical zoom-in simulations. One can use any of the public `FIRE` simulations (Wetzel et al., 2016; Hopkins et al., 2018b; Sanderson et al., 2020), which have been connected to population synthesis and cluster models successfully in the past (e.g., Lamberts et al., 2018; Chawla et al., 2022; Grudić et al., 2023; Rodriguez et al., 2023; Thiele et al., 2023). Additionally, one can use simulations from `ChaNGa`, such as the MARVEL-ous dwarfs and DC Justice League simulations (Applebaum et al., 2021; Christensen et al., 2023). These simulation suites directly resolve the formation of giant molecular clouds and the interstellar medium (ISM), and thus capture the characteristic spatial clustering of star formation. The simulations additionally explicitly account for the feedback from stars, following predictions laid out by stellar population synthesis models - though neglecting the impact of binary evolution on the timing and positioning of supernovae (see Chapter 7).

2.5.2 Snapshot preparation

Hydrodynamical simulations record snapshots of the state of the simulation a specific times. `cogsworth` provides a wrapper over `pynbody` (Pontzen et al., 2013) in order to prepare simulation snapshots for use as initial conditions to simulations. This functionality centres snapshots on the primary halo, using either an automatically detected halo catalogue or applying a shrinking sphere method to iteratively refine an estimate of the centre of the mass of the simulation. It additionally rotates the halo to be edge-on and then face-on, and converts data to physical units.

Galactic potential — As with a regular `Population`, before initialisation one needs to know the galactic potential and SFH of the galaxy. We provide functionality for computing a galactic potential from a simulation snapshot, accounting for stars, gas and dark matter, using the self-consistent field method implemented in `Gala` based on Hernquist & Ostriker (1992) and Lowing et al. (2011). This method fits the galactic mass distribution using a basis function expansion in spherical harmonics.

Initial stellar positions — The formation locations of star particles are necessary for sampling the initial positions of binary stellar populations. `cogsworth` can identify these formation locations by backwards integrating the orbits of star particles through the galactic potential derived from the simulation. Note that this step is only necessary for `FIRE` simulations, since `ChaNGa` simulations store formation locations.

For more information on processing simulation snapshots in `cogsworth` we refer interested readers to this tutorial.

2.5.3 Population initialisation and evolution

A `cogsworth` `Population` based on the star particles and galactic potential of a hydrodynamical zoom-in is called a `HydroPopulation`. Each star particle in a hydrodynamical simulation can represent many 100–1000s of solar masses. Given this, we use `COSMIC` to sample binary stellar populations from each star particle, assigning each system the same formation time and metallicity as the star particle. Based on a user-defined star particle radius, each sampled system is assigned a random position, $\vec{\mathbf{p}}_i$, from a Gaussian centred on the parent star particle such that

$$\vec{\mathbf{p}}_i = \mathcal{N}(\vec{\mathbf{p}}_{\mathrm{sp},i}, r), \tag{5.3}$$

where $\vec{\mathbf{p}}_{\mathrm{sp},i}$ is the position of star particle from which system i was sampled and r is the user choice of radius. Similarly, based on a user's choice of virial parameter, α_{vir} (the ratio of kinetic and gravitational energy of a cluster, as defined in Bertoldi & McKee, 1992) a velocity dispersion of each star particle is determined and used for sampling initial velocities, $\vec{\mathbf{v}}_i$, such that

$$\vec{\mathbf{v}}_i = \vec{\mathbf{v}}_{\mathrm{sp},i} + \vec{\mathbf{v}}_{\mathrm{disp},i}, \tag{5.4}$$

where $\vec{\mathbf{v}}_{\mathrm{sp},i}$ is the velocity of the star particle from which system i was sampled and

$$\vec{\mathbf{v}}_{\mathrm{disp},i} = \sqrt{\frac{\alpha_{\mathrm{vir}} G M_{\mathrm{cl},i}}{5r}}, \tag{5.5}$$

where $M_{\mathrm{cl},i}$ is user-defined mass of the star cluster from which system i was formed. Beyond this initial sampling, a `HydroPopulation` has the exact same functionality and methods available as a regular `Population`. A demonstration of evolving a population sampled from a snapshot is given in this tutorial.

2.6 Visualisation

`cogsworth` offers several methods for visualising the evolution and end-states of binaries evolved in populations. `COSMIC` and `Gala` already provide useful tools for investigating the binary evolution history and galactic orbits, but `cogsworth` expands these to aid in the interpretation of simulation data.

Binary evolution—For each binary evolved in `cogsworth`, one can dynamically generate a cartoon timeline of its evolution, sometimes called a Van den Heuvel diagram (van den Heuvel, 1976), with the function `plot_cartoon_binary()`. This timeline will show the evolutionary history and is capable of illustrating: the masses of each star, orbital period, mass transfer, common-envelope events, contact phases, mergers and supernovae. The distance between each star in the plot is directly scaled by the orbital separation of the binary (on a log-scale). This functionality provides a simple way to interpret the evolution of a binary without needing to know the meanings of each number representing stellar types and evolutionary stages in `COSMIC` output tables. We show an example of this in Figure 5.2 for a randomly sampled binary that we evolved for 100 Myr. Initially, as the primary loses mass to stellar winds the orbit widens very slightly. However, once the primary ends its

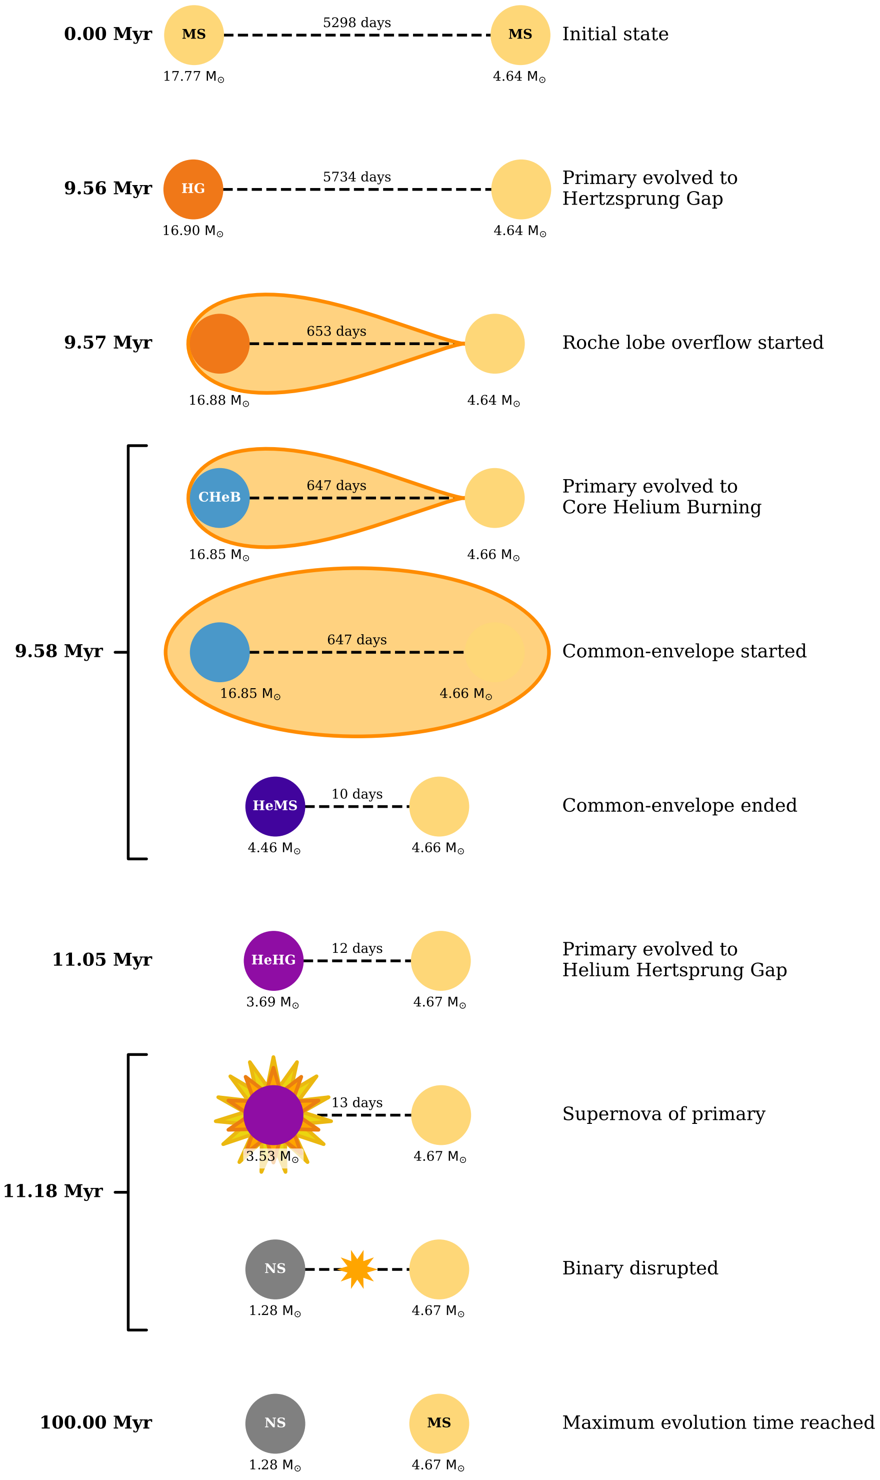

Figure 5.2: A dynamically generated cartoon binary evolution timeline. Each row shows an evolutionary step, labelled with its time and the event that occurred. Circles are shown for each star, annotated with their masses and the binary's orbital period. A dashed line indicates a bound binary, whilst a lack thereof indicates two unbound stars that were previously in a binary.

Figure 5.3: The scaling of `cogsworth`'s runtime with the number of processes used for a fixed population size of 10,000 binaries. The black points indicate the total runtime, whilst the shaded areas show the relative contribution from the stellar evolution and galactic orbit integration. The right y-axis and grey dashed line shows the relative speedup compared to the simulation using a single process.

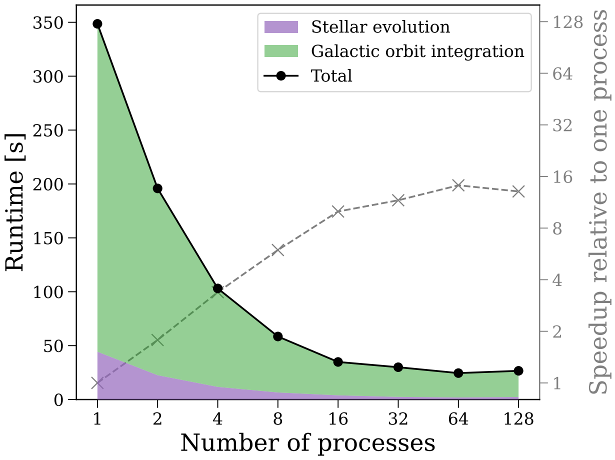

main sequence and expands across the Hertzspung gap, it initiates mass transfer. The star continues to evolve during mass transfer, eventually expanding at such a rate to make the mass transfer unstable, leading to a common-envelope and the stripping of the primary star. The primary star reaches supernova after around 11 Myr and the resulting kick disrupts the binary, ejecting the newly formed neutron star and its (prior) companion across the galaxy.

Galactic orbits — `cogsworth` provides a wrapper, `plot_orbit()`, on the `Gala` orbit plotting functionality to allow users to plot projections of a given binary's orbit in galactocentric coordinates. This allows users to plot the orbit of any evolved system, with markers indicating the location of each supernova. For a disrupted system, an additional line will be plotted for the secondary star after the binary disrupts.

Sky maps — For Milky Way simulations, users can plot their simulated populations on the sky with `cogsworth`. This is possible either with a simple scatter plot of right ascension and declination, or with a HEALPix Molleweide heatmap via `healpy` (Zonca et al., 2019; Górski et al., 2005). For HEALPix maps, one can customise the plots in several ways, including choosing the coordinates to plot (celestial, galactic, or equatorial) and the resolution of the map. See Section 3.6 for a demonstration of plotting a simulated population on the sky.

2.7 Multiprocessing scalability

Each binary in `cogsworth` is assumed to evolve independently of all other binaries in the simulation (we do not account for dynamical N-body interactions, see Section 4). An advantage of this is that each system can be efficiently parallelised. Leveraging this fact, `cogsworth` uses a multiprocessing pool for the evolution of each binary system. This means that the runtime of simulations scales well with the number of processes used.

We demonstrate the scaling of `cogsworth`'s runtime with the number of processes used in Figure 5.3. We first sample a fixed population of 10,000 binaries to ensure consistency across runs. We do not include this sampling in the runtime analysis since it is typically negligible relative to the evolution. We computed the runtime for a `cogsworth` simulation when using our default settings and varying the number of processes used.

Figure 5.3 shows that increasing the number of processes can significantly decrease the runtime of a simulation. However, this only continues up to a point, beyond which increasing the number of processes yields diminishing returns. The reason that adding processes does not always increase the runtime is that small subsets of the population take longer to run, such as binaries with multiple interactions between the two stars, or orbits that pass close to the galactic centre requiring finer time-sampling for the orbital evolution. Since we do not spread these equally among processes, additional resources are left idle while the complicated subset runs on a limited number of cores. In the case of this test, using more than 16 processes does not improve the performance, though we note that this threshold is dependent on the population (both its size and demographics) and the settings chosen by a user. We recommend that users perform similar tests to ascertain the optimal number of processes for their use case.

In Figure 5.3 we additionally show the relative contributions to the runtime from performing the stellar evolution and galactic orbit integration. For this example, the galactic orbit integration typically dominates the runtime by around a factor of $\sim$5. The relatively higher cost of the orbit integration is expected given that `COSMIC` relies on pre-computed stellar models for stellar evolution, whilst `Gala` fully integrates galactic orbits. The relative contributions from the two phases depend on the simulation that is run. One with more binaries that have many interactions between companions (e.g. a population focused on massive stars) may increase the runtime of the stellar evolution. Conversely, a simulation using a more complex galactic potential than the smooth `MilkyWayPotential2022` (such as one computed from a hydrodynamical zoom-in simulation) may increase the runtime of the galactic orbit integration. We highlight that the overhead added by `cogsworth` in connecting these two aspects is negligible.

2.8 Data storage

`cogsworth` includes functions for saving and loading an evolved `Population`. Simulations are efficiently stored in a single HDF5 file using `h5py` (Collette, 2013), which contains: simulation input settings, a SFH and sampled initial galactic variables, the chosen galactic potential, the initial state, final state and full evolutionary history of the binary stellar population, and the orbits of each system through the galaxy. This provides all the information necessary for reproducing a simulation and analysing its outputs.

We implement lazy-loading for `cogsworth` simulations. This means that not all data is immediately loaded and is instead only loaded as it is needed. For example, one can load a population without the full galactic orbits of the binaries, select a subpopulation of interest (runaway stars, for instance) and plot their orbits, at which point `cogsworth` will load only the data necessary for these systems on the fly.

The file size of a simulation depends on various factors. As one would expect, increasing the number of binaries or single stars simulated, increases the size of the output file. Additionally, since the output log includes a row for each significant evolutionary stage, binaries with more interactions between companions result in larger files. The length of time simulated is also an important factor, since it not only allows more time for binaries to experience complex evolution, but also requires more integration timesteps for galactic orbits. Users can reduce the file size of a simulation by specifying that `cogsworth` use larger integration timesteps, or even that `cogsworth` should only retain the final position of each star rather than its full galactic orbit. As some examples, the population used in Section 3.3 consists of $\sim$500 binaries, which is evolved for $\sim$50 million years, results in a file size of $\sim$5Mb. The population used in Section 3.1 is larger, containing 30,000 massive binaries and 30,000 massive single stars, and is evolved for 100 Myr, producing a file of $\sim$300Mb.

2.9 Citations

Given the breadth of work upon which `cogsworth` depends, we make it simple for users to ensure they fully credit the work that went into a given simulation. `cogsworth` is capable of creating a custom citation statement for any given `Population` using the `get_citations()` function.

For example, the simplest `cogsworth` simulation may cite only `cogsworth` itself, `COSMIC` and `Gala`. But if you compute the observable features of your population, `cogsworth` will add citations for the dust maps, isochrones, and selection function that you use. This is similarly true for citations regarding star formation histories and hydrodynamical zoom-in simulations.

3 Use cases

The following subsections each demonstrate a particular use case of `cogsworth`, showcasing its capabilities in binary stellar evolution, galactic dynamics, observables estimation, and integration with hydrodynamical simulations. Each 💻 icon links directly to a page in our online documentation that guides users through using `cogsworth` to reproduce a given figure.

Unless otherwise specified in the individual use cases, each `cogsworth` simulation uses the `Wagg2022` SFH, the `MilkyWayPotential2022` Galactic potential and the default binary physics settings from `COSMIC` v3.4.16. Additionally, primary masses are sampled using the Kroupa (2001) IMF, a uniform mass ratio distribution is assumed and the initial orbital period and eccentricity distributions follow Sana et al. (2012). These assumptions are appropriate for massive stars, which can experience core-collapse events and be kicked and disrupted.

3.1 The importance of binary evolution and the galactic potential

In this use case, we demonstrate the need for accounting for binary interactions and a galactic potential simultaneously. We use `cogsworth` to simulate a population of massive stars formed in the most recent 1 Gyr in the Milky Way assuming a 50% binary fraction. We then repeat the orbital integration for the binaries in this population, but without a galactic potential. In this way, binaries with no core-collapse events remain in their birth locations, whilst those that receive kicks continue at their ejection velocity indefinitely.

We compare the present-day positions of single stars and binary stars in Figure 5.4, additionally showing the distribution of binary stars when neglecting the Galactic potential. First, comparing the single and binary stars in the presence of a potential, we note that the tails of the binary distribution are significantly extended. For single stars, the fraction of the population at $|z| > 1\,\mathrm{kpc}$ is only 0.9%, whereas for binaries it increases by a factor of 2x to

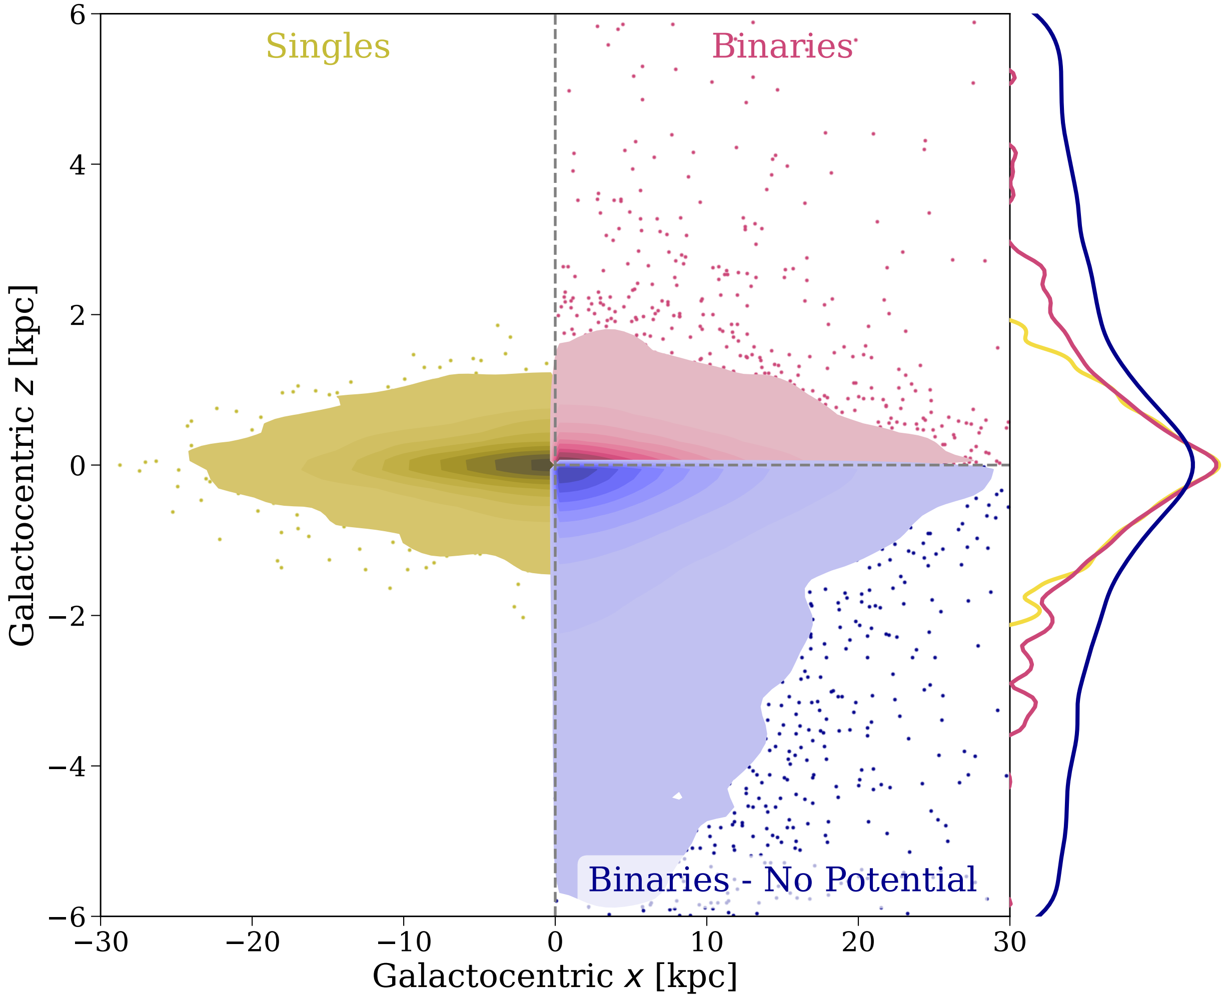

Figure 5.4: Binary evolution and the presence of a galactic potential both have a significant impact on the spatial distribution of stars. Each panel shows the spatial distributions of massive stars at present day from a `cogsworth` population of single stars (left), binaries (right, top) and binaries evolved without a galactic potential (right, bottom). Density distributions are computed using kernel density estimators (KDE) and shown (in linearly spaced isolevels) up to the 98th percentiles, with the remaining stars plotted as scatter points. Marginal KDEs are shown on the right with a log scale.

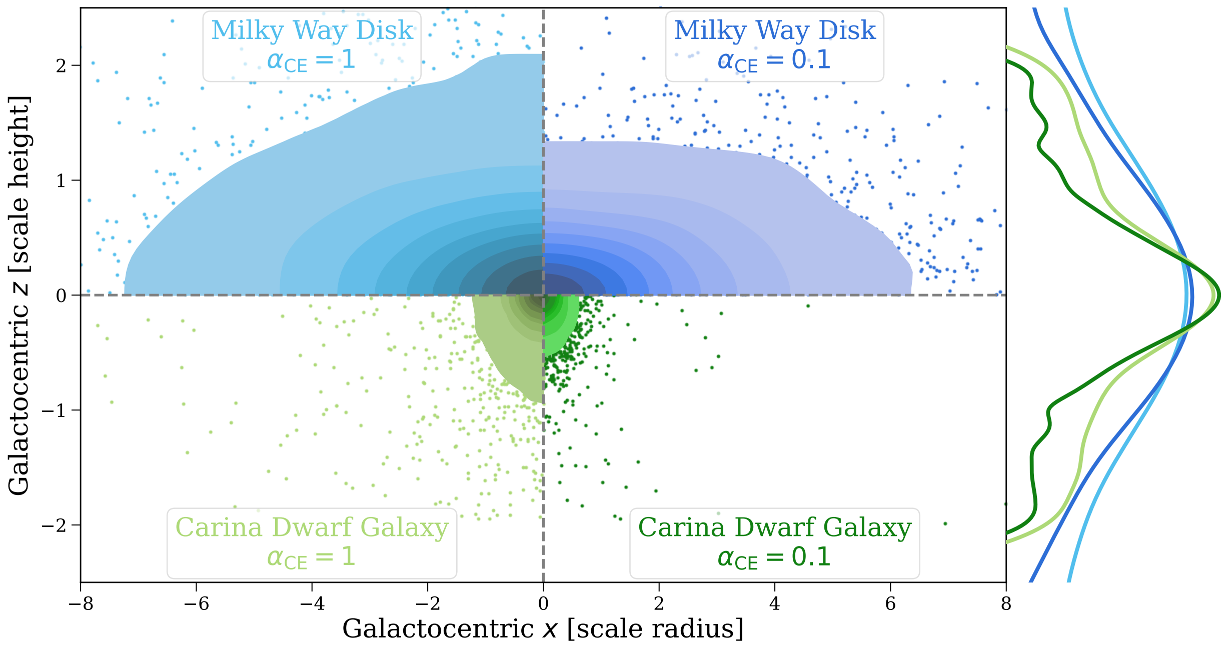

Figure 5.5: Similar to Figure 5.4, but comparing the impact of varying binary physics and the galactic potential. The axes are now in units of scale radius and scale height for better comparison between the two galaxies. Each quadrant corresponds to a different combination of common-envelope efficiency and galactic potential (annotated in each).

1.9%. Previous work has investigated the spatial distribution of compact objects without accounting for binary interactions , assuming that these interactions do not change the overall distributions (Sweeney et al., 2022). However, as Figure 5.4 shows, our `cogsworth` simulations demonstrate that binary interactions can significantly alter the spatial distributions of massive stars and compact objects.

Moreover, comparing the binary population with and without the Galactic potential, it is clear that neglecting the Galactic potential results in misleading conclusions regarding spatial distributions. In particular, the population without a potential is broader in both width and height, and the fraction at $|z| > 1\,\mathrm{kpc}$ is 8.6%, more than 4x the fraction when accounting for the potential.

3.2 Comparing the impact of common-envelope efficiency and galactic potentials on spatial distributions

A key feature of `cogsworth` is its ability to self-consistently account for the effect of binary interactions and galactic potentials. We explore this capability by determining the relative impact of varying supernova natal kicks and galactic potentials on the spatial distribution of a population of binary stars.

We vary the efficiency of common-envelope phases from our default choice of $\alpha_{\rm CE} = 1.0$ to an extreme choice of $\alpha_{\rm CE} = 0.1$, where $\alpha_{\rm CE}$ is the fraction of orbital energy that is available to unbind the envelope (Webbink, 1984; de Kool, 1990). This means that when $\alpha_{\rm CE}$ is lower equivalent systems that survive the phase are tighter after a common-envelope, such that they have higher orbital velocities and therefore ejection velocities. However, binaries are also more likely to merge as a result of a common-envelope phase. For galactic potentials, we compare the Milky Way's disk (using parameters from Sanders & Binney 2015) to a spheroidal dwarf galaxy with parameters matching the Carina dwarf galaxy (Pascale et al., 2019). The latter is assumed to be in isolation, such that there is no tidal stripping from the Milky Way as is observed in the Carina galaxy.

We sample an initial binary population (assuming a binary fraction of 100%) from the most recent 1 Gyr of star formation, retaining only systems with primary stars more massive than $7\,\mathrm{M}_\odot$, since here we are only interested in binaries that experience a core-collapse event. `COSMIC` reports the precise initial conditions of sampled populations, making them easily reproducible with different evolution settings. Using this feature, we evolve identical initial populations with the two different common-envelope efficiencies, then integrate the orbits of each variation through the two different potentials. This results in 4 different populations of present-day positions.

In Figure 5.5, we compare these populations, showing the positions of all massive stars that experienced a core-collapse event, or were a companion to a star that did. To make the comparison across galaxies easier, we plot the positions in terms of the scale radius, r_s, and scale height, z_s, of each galaxy, where for the Milky Way we assume $r_s = 2.6\,\mathrm{kpc}$ and $z_s = 0.3\,\mathrm{kpc}$ (Bland-Hawthorn & Gerhard, 2016) and for the Carina galaxy we assume $r_s = z_s = 1\,\mathrm{kpc}$ (Pascale et al., 2019). Given the variety of results from the four panels, we highlight that both binary physics and the choice of galactic potential can have a strong effect on the resulting spatial distribution of massive stars. In particular, less efficient common-envelope phases result in reduced galactic scale heights and radii for massive stars. This is because a significant fraction of binaries that experience a common-envelope merge during the event, thus reducing the number of stars that are ejected from their binaries. Although the width of the overall distributions are difficult to statistically distinguish, the tails of the distribution are more strongly affected. The fraction of objects at $|z| > z_s$ decreases from 5.1% to 3.0% for the disc population, and from 1.4% to 0.4% for the dwarf galaxy population. Therefore, observing the outliers in a galactic height distribution of massive stars can be more informative than the overall width of the distribution for inferring the efficiency of common-envelopes. One would need to also investigate other parameters that could affect these distributions, in particular the critical mass ratio for the stability of mass transfer and the efficiency of stable mass transfer (e.g., Evans et al., 2020).

3.3 Evolution of binary orbits in a star cluster

`cogsworth` is capable of producing populations of binaries based on hydrodynamical simulations (Section 2.5). We use `cogsworth` to post-process the FIRE m11h simulation (El-Badry et al., 2018; Wetzel et al., 2023), an intermediate-mass halo with a strong disc component, fitting a galactic potential and rewinding all star particles formed in the past 150 Myr to their formation locations. We then sample a binary population from each star particle, matching its formation time, metallicity and mass. For this example, we examine one random star particle and the orbits of its constituents. This star particle was formed $\sim$43 Myr before present-day, with a mass of $\sim 6000\,\mathrm{M}_\odot$ and a metallicity of $Z \approx 0.0137$. The sampled pop-

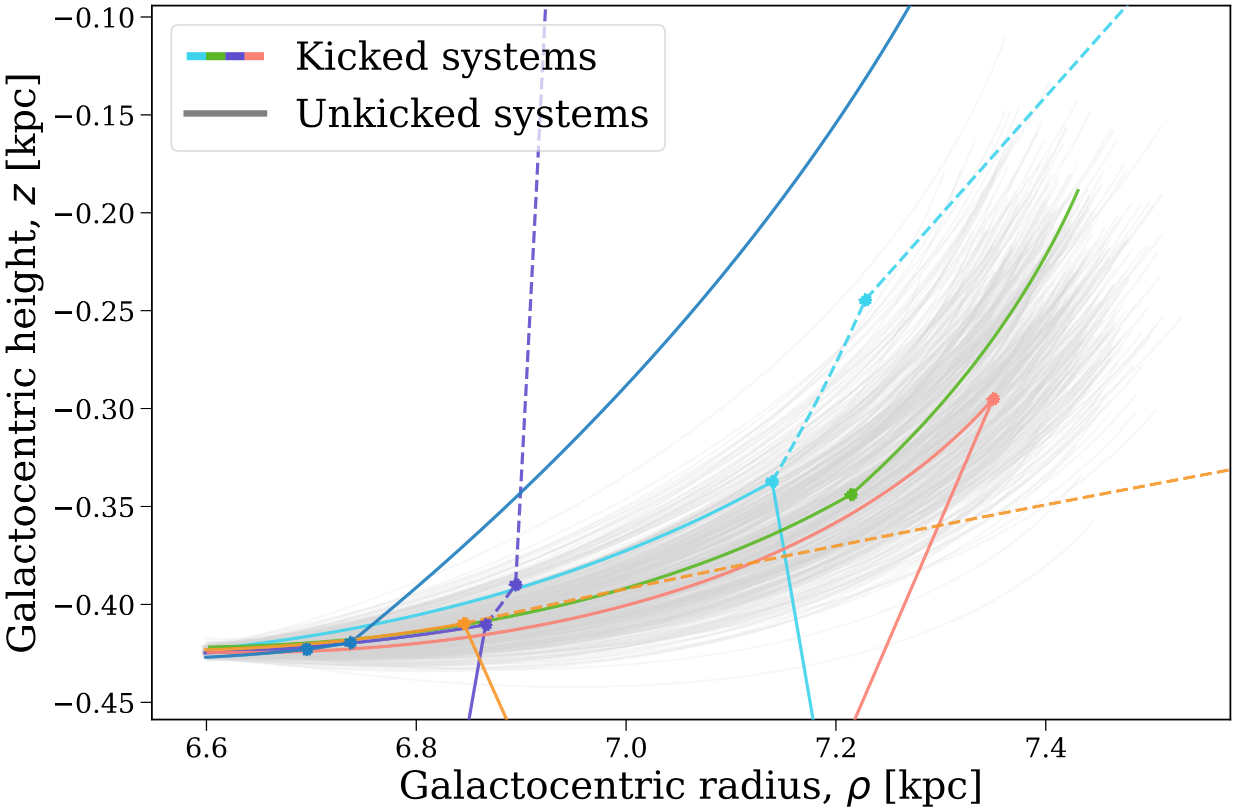

Figure 5.6: An example star cluster, sampled from a single star particle in the FIRE m11h hydrodynamical zoom-in simulation, evolved within a `cogsworth` simulation. Each line shows the orbit of a star in the cluster. Coloured lines are for stars that had supernova events (highlighted by scatter points), while the rest are shown in grey. For binaries that disrupt, an additional dashed line is shown for the subsequent orbit of the ejected companion.

ulation consists of ~ 7600 systems, split evenly between single stars and binary stars (as one would expect given our assumption of a 50% binary fraction). Note that we neglect the self-gravity of the cluster.

In Figure 5.6, we plot the orbits of a representative subset of 500 of the binaries sampled from the star particle. In grey, we show the orbits of binaries that experienced no supernovae events, which are by far the majority since the IMF favours low mass stars. The cluster is formed in the lower left (at $\rho = 6.6\,\mathrm{kpc}, z = -0.42\,\mathrm{kpc}$) and evolves to larger ρ. One can note the dissolution of the cluster over time in the grey orbits, which occurs as a result of the initial velocity dispersion (Eq. 5.5).

The coloured lines show the more eventful orbits of binaries that experienced supernovae. In many cases this leads to the disruption of the binary orbit and so we show the orbit of the subsequent evolution of the ejected companion with a dashed line. With `cogsworth`, one can examine the detailed evolution of each binary to understand its orbit. The examples shown include several scenarios involving bound, disrupted and merged binaries - we discuss each in detail below.

The earliest core-collapse event occurs for the dark blue binary after 4 Myr, which is indicated by closest scatter point to the cluster origin in the lower left. This forms a $12\,\mathrm{M}_\odot$ black hole and, due to a fallback fraction of 92% for the black hole, much of the explosion asymmetry is negated, resulting in a relatively weak natal kick of $11\,\mathrm{km\,s^{-1}}$, which allows the binary the remain bound. The companion to this binary reaches core-collapse 1.5 Myr later, forming a slightly less massive black hole of $6\,\mathrm{M}_\odot$, with a stronger natal kick of $70\,\mathrm{km\,s^{-1}}$. Yet the binary is much tighter at this point (with a separation of $45\,\mathrm{R}_\odot$) and thus has a higher binding energy. This means that it remains bound and is ejected from the cluster as a binary black hole.

The first supernova in the orange binary occurs 9 Myr after the cluster birth and forms a neutron star with a natal kick of $447\,\mathrm{km\,s^{-1}}$. This kick disrupts the binary orbit, such that both stars are ejected from the cluster. The secondary is a lower mass star of $3.6\,\mathrm{M}_\odot$ and so experiences no supernova, but it is ejected from the cluster at $160\,\mathrm{km\,s^{-1}}$ and as such is now a runaway star.

The purple binary experiences its first supernova at 9.3 Myr and, similar to the orange binary, this forms a $1.6\,\mathrm{M}_\odot$ NS with a natal kick of $415\,\mathrm{km\,s^{-1}}$ that unbinds the binary. Interestingly, the mass ratio of this system is inverted, such that the companion forms a more massive $5.3\,\mathrm{M}_\odot$ black hole after its core-collapse 2 Myr later. This inversion occurred as a result of significant, near conservative mass transfer from the primary star during its Hertzsprung gap phase 1.2 Myr before it reached core-collapse. Both supernovae for the light blue binary form neutron stars (of 1.3 and $2.2\,\mathrm{M}_\odot$ respectively) with strong kicks (of 406 and $728\,\mathrm{km\,s^{-1}}$ respectively), which disrupt the binary and eject both neutron stars rapidly from the cluster.

The primary star in the green binary reaches core-collapse 26 Myr after the cluster birth, forming a neutron star of $1.27\,\mathrm{M}_\odot$. The star explodes as an electron-capture supernova and thus its kick is assumed to be weaker, in this case the drawn natal kick is only $35\,\mathrm{km\,s^{-1}}$ and the binary remains bound. In addition, as a result of the angle of the kick relative to the binary's orbit, this only results in a $3.4\,\mathrm{km\,s^{-1}}$ change to the systemic velocity of the binary. As a result, the binary remains bound to the cluster for its subsequent evolution.

Finally, after ∼33 Myr, the primary star in the red binary finishes its main sequence, As it expands during its Hertzsprung gap phase, it overflows its Roche lobe, causing unstable mass transfer which leads to a merger. The merged star then reaches supernova 4 Myr later, forming a neutron star which is ejected by its strong natal kick of $819\,\mathrm{km\,s^{-1}}$, in almost the opposite direction to the cluster's centre of mass motion.

3.4 Examining metallicity-radius-time relations in star formation histories

`cogsworth` can be used to sample detailed star formation histories independently of evolving binary stars or performing galactic orbit integration (see Section 2.1). In this use case, we explore the `Wagg2022` SFH model in more detail. We sample 500,000 points (which could be designated as a single or binary star) from the SFH, which each have an associated position, lookback time and metallicity.

In the left panel of Figure 5.7, we plot the distribution of Galactocentric radii and metallicities for each sampled point. As a general trend, one can note that stars closer to the

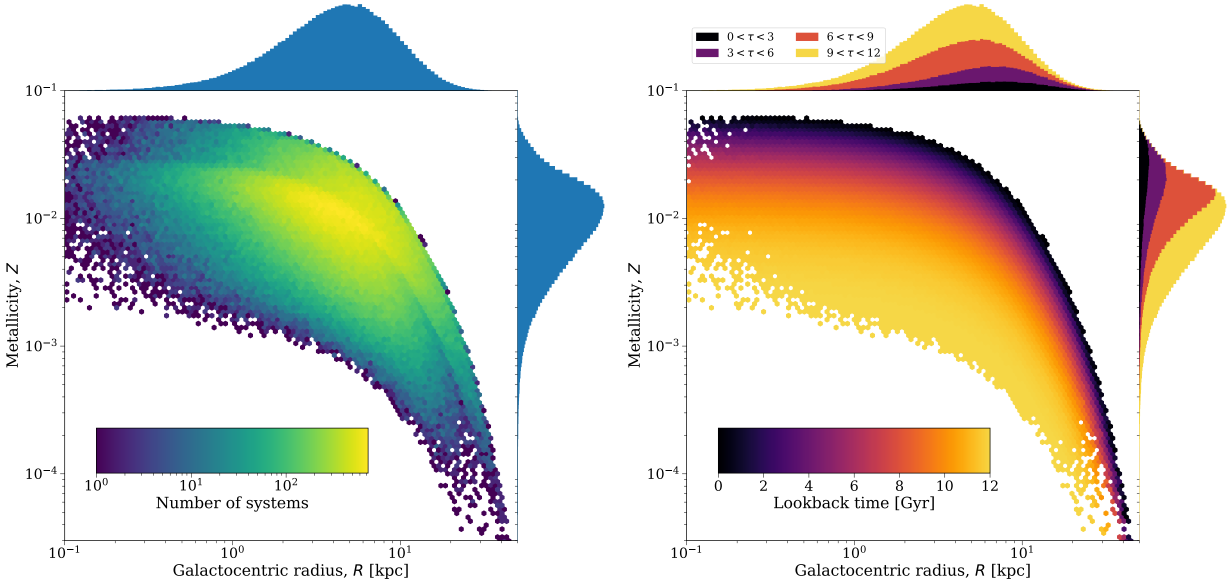

Figure 5.7: The relationship between galactocentric radius, metallicity and lookback time in the `Wagg2022` SFH. **Left:** The main panel shows a 2D histogram of the number of stars formed at a given radius and metallicity, with marginal distributions for each parameter shown as histograms on the side panels. **Right:** The main panel shows a 2D histogram in which each bin is coloured by the average lookback time of associated stars. Marginal distributions are now shown as stacked histograms, grouped by lookback time.

centre of the Galaxy are more metal-rich than those on the outer edges. This is due to the inside-out growth of the Galaxy (e.g., Fall & Efstathiou, 1980; Frankel et al., 2019), which is accounted for in this SFH following the model of Frankel et al. (2018). Additionally, the discontinuity in the distribution (occurring at inner radii at $Z \approx 0.03$) is a result of the multi-component nature of the model. The upper right portion above the discontinuity comes is primarily from the low-[α/Fe] disc component, which forms stars from 8 Gyr ago until present-day, while the lower portion is primarily from the high-[α/Fe] disc, which form stars from 12 Gyr ago until 8 Gyr ago. The bulge component contributes to both portions, though only at small radii.

For a given radius, there is a wide variation in the metallicity of sampled stars. This is because of the birth time of each star, which we demonstrate in the right panel of Figure 5.7. The 2D histogram now shows the average lookback time, τ, of the stars in each bin (where $\tau = 0$ corresponds to present-day). The clear gradient shows that over time the Galaxy as a whole becomes more metal-rich as it is enriched by stellar evolution. This is additionally visible in the marginal distribution of metallicities, where no high metallicity stars are formed at early birth times. The marginal distribution of radii again demonstrates the inside-out growth, as older stars were formed closer to the Galactic centre.

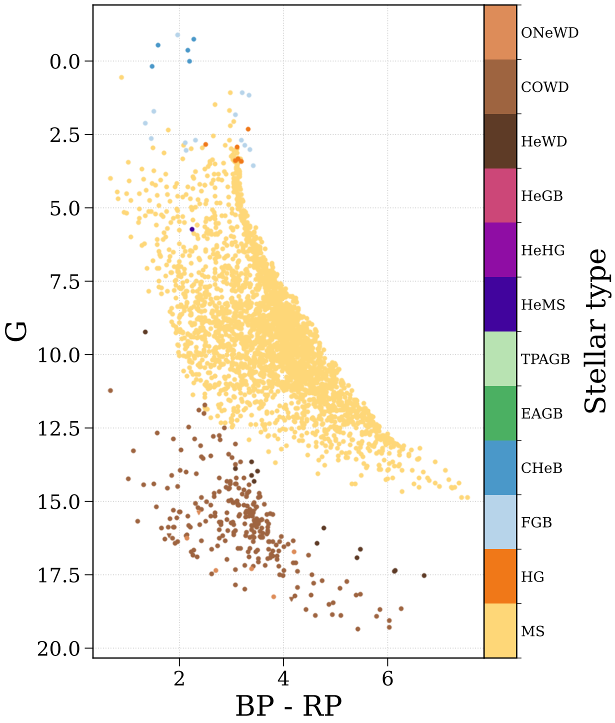

Figure 5.8: An example simulated *Gaia* colour-magnitude diagram. Each scatter point corresponds to a binary or disrupted star, coloured by its stellar type (using the stellar type of the brighter component for binaries). The stellar types abbreviated in the colourbar follow those defined in `BSE` (see Section 4 of Hurley et al., 2000).

3.5 Simulating a *Gaia* colour-magnitude diagram

In this use case, we highlight `cogsworth`'s ability to transform an intrinsic population into a simulated observable population (see Section 2.4). We sample 2500 binary systems over the full SFH of the Milky Way and evolve them until present day.

We compute observables for this population using the `Population.get_observables()` function in `cogsworth`. We first calculate the absolute magnitude of each star and determine which star is brighter in a binary. `cogsworth` then converts these magnitudes to the *Gaia* filters G, BP and RP using the MIST isochrones to apply bolometric correction with the `isochrones` package (Morton, 2015; Dotter, 2016; Choi et al., 2016; Paxton et al., 2011, 2013, 2015). Finally, `cogsworth` uses the `dustmaps` package to account for dust extinction through the application of the Bayestar 2019 dust maps (Green, 2018; Green et al., 2019).

We plot the resulting colour-magnitude diagram (CMD) in Figure 5.8, colouring each bound system by the stellar type of the brighter component in the G band. These systems cover a range of metallicities, distances and ages and hence have a wide spread in the CMD. In addition to these systems, 59 isolated (either from mergers or binary disruptions) neutron stars and black holes are present in the evolved population. This same simulation could be easily repeated with a starburst localised in one specific place to model a specific cluster CMD instead.

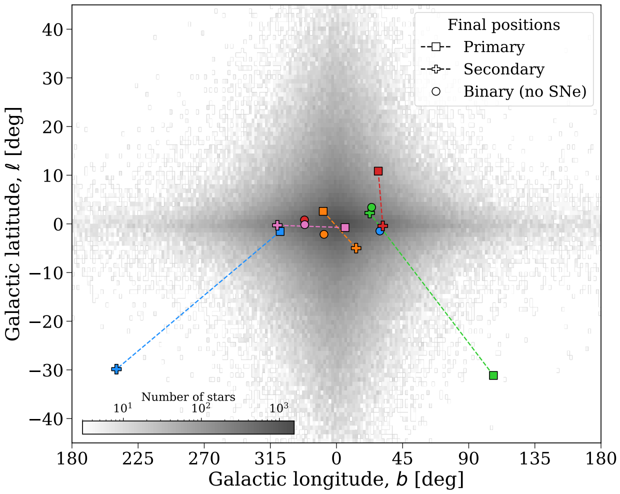

Figure 5.9: Present-day sky positions of disrupted binaries, with and without supernova kicks. A square (plus) marker shows the location of the primary (secondary) star from a disrupted binary, connected to its companion by a dashed line. Circular markers indicate the location of the binary had no supernova kicks occurred. The background density shows a 2D histogram of 500,000 stars sampled from the same SFH (`Wagg2022`) for comparison. Note that disk appears extended in latitude because we limit the axes to the relevant region.

3.6 Comparing present-day sky locations with and without supernova kicks

`cogsworth` can report the present-day sky position of each source, in addition to its evolutionary history. In this use case we demonstrate how one can track the relative positions of binary companions after they disrupt, as well as consider where they would be found if no supernova kick had occurred.

We sample and evolve 100 random binaries in the Milky Way with our default assumptions, except we set the minimum mass of the IMF to $3\,\mathrm{M}_{\odot}$. This preferentially samples more massive stars, which are more likely to reach core-collapse and cause a binary disruption. From this population, we subselect five binaries that are disrupted. These binaries experienced at least one core-collapse event which disrupted the orbit and led to the separation of the two unbound companions. We use `cogsworth` to compute the present-day sky location of both companions for each binary, before re-integrating the binary's orbit without accounting for the impact of supernova kicks. `cogsworth` returns the final coordinates of stars as an `Astropy SkyCoord` (Astropy Collaboration et al., 2013, 2018, 2022), which allows for simple transformation between coordinate frames.

In Figure 5.9, we compare the present-day sky locations (in galactic coordinates) of each compact object from the disrupted binaries to the location of the binary had it experienced no supernova kicks. As expected, when neglecting supernova kicks, binaries are generally concentrated close to the galactic midplane and centre. However, when accounting for the effect of supernova kicks on the internal and galactic orbits, the present-day sky locations are often significantly different. In many cases, companions are located not only far from one another, but also far from the position of the binary had no kicks occurred. Several of the compact objects from the disrupted binaries are found well beyond the typical sky locations of galactic sources (shown as a histogram in the background), though all remain bound to the galaxy. In particular, the secondary from the blue binary is first ejected from the binary as a runaway star (travelling at $\sim 35\,\mathrm{km\,s^{-1}}$) after the primary reaches core-collapse. After spending 10 Myr as an O-type runaway star, the secondary reaches core-collapse and is kicked onto an orbit that takes it even further from the typical galactic population. The green binary starts on a relatively wide initial orbital period of $\sim$30, 000 days. The binary widens by nearly 10% due to the stellar winds of the primary before this star's core-collapse 32 Myr after the birth of the binary. The primary forms a neutron star that receives a natal kick of $325\,\mathrm{km\,s^{-1}}$, such that it takes large excursions from the galactic plane during its orbit. However, due to the large orbital period at the supernova, the secondary star is only ejected at $16\,\mathrm{km\,s^{-1}}$. This star eventually forms a white dwarf and receives no further kick, hence its present-day location is relatively similar to that of the binary had no kick occurred.

4 Limitations

Although `cogsworth` has a wide range of features and capabilities, there are still some limitations to the code that users should be aware of.

Dynamical interactions — We do not implement dynamical interactions between systems or account for any N-body dynamics. As mentioned in Section 2.7, each binary in `cogsworth` is evolved independent of all others. This means that dynamical formation channels for different populations (such as interactions in dense clusters forming gravitational-wave progenitors) cannot be simulated in `cogsworth`. However, we do intend to implement a prescription for emulating dynamical cluster ejections (see Section 5.2), such that investigations into runaway stars could consider both channels.

Impact of galactic potential on internal orbits — `cogsworth` accounts for the impact of stellar evolution on the galactic orbits of binaries (i.e. as a result of supernovae). Yet it does not currently account for the inverse case, in which a galactic potential alters the orbit of the binary. For wide binaries the potential can significantly affect the orbit, causing large-amplitude oscillations and potentially drive systems to disrupt or merge (e.g., Weinberg et al., 1987; Heisler & Tremaine, 1986; Jiang & Tremaine, 2010; Modak & Hamilton, 2023; Stegmann et al., 2024). However, for closer binaries the effect is negligible and as such as we do not currently account for it in `cogsworth`.

Population synthesis model uncertainties — `cogsworth` uses `COSMIC` for binary population synthesis, which is a code based on `BSE`. The `BSE` code relies on approximate parametric prescriptions for a limited set of evolutionary tracks of single stars (Pols et al., 1998; Hurley et al., 2000, 2002). Although many of the original prescriptions used in the `BSE` code have been improved in `COSMIC` (Breivik et al., 2020b), the core of the code still relies on the same methodology. In particular, the treatment of mass loss and the stability of mass transfer, as well as the reliability of the most massive progenitor models, is uncertain. However, some of these uncertainties can be alleviated by incorporating information on the internal structure of stars (e.g., Kruckow et al., 2018; Fragos et al., 2023). `COSMIC` is in the process of being integrated with `METISSE`, MEthod of Interpolation for Single-Star Evolution. `METISSE` is an alternative to fitting formulae that allows for the interpolation between pre-computed detailed one-dimensional stellar evolution tracks, while maintaining the same code interfaces as the previously implemented prescriptions of SSE (Agrawal et al., 2020, 2023). By working with updated libraries of pre-computed single star tracks from `MESA` (Paxton et al., 2011, 2013, 2015, 2018, 2019; Jermyn et al., 2023), `METISSE` enables a wide range of investigations of the impact of uncertainties in single-star evolution like convection, rotation, and nuclear reaction rates and how these uncertainties interface with uncertainties in binary interaction physics. Once `METISSE` is fully integrated into `COSMIC`, `cogsworth` will be able to immediately leverage these new improvements.

5 Future developments

We intend to complete further development on `cogsworth` beyond this initial release. In the following subsections we highlight some areas in which we aim to focus.

5.1 Time-evolving galactic potentials

Traditional models using static galactic potentials are not capable of describing the dynamically complex evolutionary history of galaxies, and can lead to misleading results (e.g., Arora et al., 2022). Although this is less relevant for shorter lived populations (such as massive runaway stars), it could have important implications for longer lasting tracers of binary endpoints (e.g., gravitational wave mergers). Thus we also intend to leverage our integration with hydrodynamical simulations to implement a `cogsworth` option for a time-evolving gravitational potential that accounts for the mass growth of a galaxy over time.

5.2 Dynamical cluster ejections

`cogsworth` does not currently account for the impact of dynamical interactions between binaries in dense environments, as noted in Section 4. The interactions can change the initial architecture of binaries (e.g., Fujii & Portegies Zwart, 2011) and create alternate formation channels for binary products. For instance, runaway stars are thought to be formed in two main channels: the disruption of binaries as a result of supernovae (Blaauw, 1961; Eldridge et al., 2011; Renzo et al., 2019) and dynamical ejections from stellar clusters (Poveda et al., 1967). Fully modelling the latter channel would require more complex N-body dynamics that are currently beyond the scope of the code. Instead, we intend to create an approximation in which we will give a fraction of massive stars kicks shortly after their formation. The mass-dependent rate, kick velocity and timing will follow distributions modelled in N-body simulations (e.g., Oh & Kroupa, 2016; Schoettler et al., 2022).

5.3 Other observables

For high-energy, degenerate, and/or accreting sources formed through binary channels (e.g., X-ray binaries, cataclysmic variables, short gamma-ray bursts, type Ia supernovae), the mapping between binary physical parameters and flux is naturally more complex (and sometimes uncertain) than it is for most stars. This means that predictions for other observables (beyond those current implemented) are more complicated, though not out of reach in many cases. For example, prescriptions for the X-ray luminosity of a given binary exist (Misra et al., 2023), and we intend to add this feature to `COSMIC` (and thus also to `cogsworth`) to make predictions for the X-ray binary populations that have been widely observed with

Chandra in nearby galaxies. In the future, we will implement mappings for other missions and observables based on their own selection functions.

6 Conclusions & Summary

In this paper we have presented `cogsworth`, a new open-source code for performing self-consistent population synthesis and galactic dynamics simulations. `cogsworth` provides the theoretical infrastructure necessary to make predictions about the positions and velocities of stars and compact objects. We have demonstrated several use cases of the code, showcasing its capabilities to investigate the impact of binary interactions and galactic potentials on the evolution of stars and compact objects - both for intrinsic and observable populations. `cogsworth` could be applied to a plethora of investigations on a wide-range of populations, including runaway stars, supernova remnants, X-ray binaries, short gamma-ray bursts and double compact objects.

Given its accessibility and flexibility, we hope that `cogsworth` will be a useful tool for the community, enabling and accelerating future studies into binary stars and compact objects.

CHAPTER 6

THE WEAK DEPENDENCE OF STELLAR EJECTION VELOCITIES ON SUPERNOVA KICKS

1 Introduction

Observations of pulsars in the Milky Way indicate that they have a much higher velocity dispersion than their massive star progenitors, and hence that neutron stars acquire high velocities upon formation (e.g., Gunn & Ostriker, 1970; Hobbs et al., 2005; Igoshev, 2020). For compact objects formed via isolated binary evolution, these velocities are expected to be imparted at the moment of core-collapse. Even for a symmetric explosion, the instantaneous mass loss from a supernova (SN) changes the post-SN orbital parameters (Blaauw, 1961; Boersma, 1961). The symmetric mass loss alone is typically insufficient to disrupt the orbit, since the mass that can be ejected is often limited by loss during previous mass transfer phases. However, in some cases (∼16%) symmetric mass loss from supernova ejecta alone can unbind the binary (Renzo et al., 2019). This is typically only possible for wide binaries that have massive SN progenitors, which eject a significant fraction of their mass. Moreover, asymmetry in the explosion can impart a significant natal kick on the newly formed compact object (e.g., Shklovskii, 1970; Lyne & Lorimer, 1994; Janka, 2013, 2017), which in the majority of cases is strong enough to disrupt a binary (e.g., De Donder et al., 1997; Eldridge et al., 2011; Renzo et al., 2019). Similar asymmetry in neutrino emission can result in additional natal kicks, which can dominate in the case of full fallback (e.g., Janka & Mueller, 1994; Burrows & Hayes, 1996; Scheck et al., 2006; Wongwathanarat et al., 2013; Coleman & Burrows, 2022; Vigna-Gómez et al., 2024).

The effect of these SN kicks on the post-SN orbital parameters of a binary system has been considered in detail by many previous works (e.g., Flannery & van den Heuvel, 1975; Sutantyo, 1978; Hills, 1983; Dewey & Cordes, 1987; Wijers et al., 1992; Brandt & Podsiadlowski, 1995; Kalogera, 1996; **?**; Hurley et al., 2002; Pfahl et al., 2002; Belczynski et al., 2008; Kiel & Hurley, 2009). These investigations account for a variety of effects, with the most extensive models addressing the impact of the natal kick on the compact object, including the effect of instantaneous symmetric mass loss, handling eccentric pre-SN systems, the probability of disruption of the system, the SN blast wave impulse on the secondary star (e.g., Liu et al., 2015; Hirai et al., 2018; Ogata et al., 2021; Wong et al., 2024), and collisions between the compact object and companion star (e.g., Davies et al., 1992; Hirai & Podsiadlowski, 2022; Kremer et al., 2022). For a detailed derivation, which explains how each of these effects impacts a binary in clear and compact vector notation, we refer the reader to Appendix B of Pfahl et al. (2002).

SN kicks have a wide-ranging impact on massive stellar populations, for both disrupted and bound binary products. The post-SN eccentricity, semi-major axis and centre-of-mass velocity of bound systems, as well as the rate of unbinding and ejection velocities of stars and compact objects, are sensitive to modelling choices in kick prescriptions. The demographics and rates of bound binaries containing at least one compact object have been shown to be strongly affected by the magnitude and distribution of SN kicks. In particular, kicks affect the prevalence of galactic X-ray binaries (e.g., Tauris et al., 1999; Pfahl et al., 2002; Wong et al., 2012; Vigna-Gómez et al., 2024), dormant black hole binaries (e.g., Breivik et al., 2017), double neutron stars (e.g., Brandt & Podsiadlowski, 1995; Wex et al., 2000; Podsiadlowski et al., 2004; Bray & Eldridge, 2016; Vigna-Gómez et al., 2018), and gravitational-wave sources (e.g., Fryer et al., 1998; Belczyński & Bulik, 1999; Dominik et al., 2013; Broekgaarden et al.,

Code	Supernova kick prescription
`COSMIC`	Kiel & Hurley 2009 (<v3.5.0) Pfahl et al. 2002 (≥v3.5.0)
`COMPAS`	Pfahl et al. 2002
`binary_c`	?

Table 6.1: The prescription for supernova natal kicks used in the population synthesis codes we consider.

2022; Wagg et al., 2022c). Moreover, the observed population of runaway stars has grown significantly in the era of *Gaia* (e.g., Carretero-Castrillo et al., 2023). The relative rate of runaway star production from the binary supernova scenario (Blaauw, 1961; Boersma, 1961) and the dynamical ejection scenario (Poveda et al., 1967) remains debated (e.g., Sana et al., 2022). The ejection velocity of a runaway star is not strongly dependent on its companion's SN kick, as we will show here. However, the fraction of binaries that are disrupted by SNe is highly sensitive to the distribution of SN kicks. Therefore, SN kicks strongly affect the rate and demographics of runaway stars produced by the binary supernova scenario (e.g., De Donder et al., 1997; Eldridge et al., 2011; Renzo et al., 2019).

Binary stellar population synthesis codes are used to make predictions for these binary products that, in tandem with observations, are used to constrain models for SN kicks. These codes apply different choices of SN kick methodologies in their evolution, which have the potential to significantly change the subsequent evolution of binary systems after a SN. In this paper, we compare SN kick prescriptions of three different codes to ensure consistency. In particular, we focus on how the ejection velocities of stellar companions compare for a binary with identical initial conditions evolved in different codes.

2 Methods

2.1 Population synthesis codes and default settings

We use three different open-source population synthesis codes in our analysis. The `COSMIC` population synthesis code (Breivik et al., 2020b) applies the method from Pfahl et al. (2002) (though previously followed Kiel & Hurley (2009), see Section 4). Similarly, `COMPAS` (Riley et al., 2022) uses the method of Pfahl et al. (2002). In contrast, `binary_c` (Izzard et al., 2004, 2006, 2009, 2018; Izzard & Jermyn, 2023; Hendriks & Izzard, 2023) applies **?**. The combination of these codes therefore allows us to cross-validate three implementations of two different algorithms for determining secondary ejection velocities.

These codes share a joint origin of BSE (Hurley et al., 2002), but have since diverged and implemented additional physics, with alternate default settings. As a result, the same binary evolved in each code can result in moderately different evolution. Of particular relevance to this study are the choices regarding remnant mass prescriptions and SN natal kick magnitudes. By default, `COSMIC` and `binary_c` both use the delayed remnant mass prescription from Fryer et al. (2012), which is a deterministic function of the pre-SN CO core mass. In contrast, `COMPAS` instead applies the Mandel & Müller (2020) prescription, which is a probabilistic model that accounts for some of the intrinsic stochasticity in stellar evolution and SNe. This model also draws a natal kick based on the remnant mass. `COSMIC` instead samples a natal kick magnitude from a Maxwellian distribution, with a central parameter of $\sigma = 265\,\mathrm{km\,s^{-1}}$ for core-collapse SNe (e.g., Hobbs et al., 2005), and $\sigma = 20\,\mathrm{km\,s^{-1}}$ for electron-capture SNe (e.g., Igoshev, 2020). The default choice for `binary_c` is to sample natal kicks from a single Maxwellian distribution, with a central parameter of $\sigma = 190\,\mathrm{km\,s^{-1}}$ (Hansen & Phinney, 1997) and to follow the remnant mass prescription from BSE (Hurley et al., 2000, 2002).

Given that our goal is to compare the implementation of different kick prescription algorithms, we enforce consistent choices for natal kick magnitudes and remnant masses to better isolate any differences between the codes. Therefore, for each code, we draw natal kick magnitudes from a Maxwellian with $\sigma = 265\,\mathrm{km\,s^{-1}}$ and apply the Fryer et al. (2012) delayed remnant mass prescription. All other settings we leave unchanged from the code's default choices. As a result, the pre-SN evolution in each code is different for the same initial conditions, which we discuss in Section 3.1.

2.2 Initial conditions

We evolve a representative mass-transferring massive binary with a range of random supernova natal kicks in each code. This binary has the following initial conditions: primary mass $m_1 = 20\,\mathrm{M_\odot}$, secondary mass $m_2 = 15\,\mathrm{M_\odot}$, orbital period $P = 100$ days, eccentricity $e = 0$, and metallicity $Z = 0.02$. We repeat the evolution 50,000 times, each with a different random natal kick magnitude and direction. The kick is assumed to be isotropic, and magnitudes are drawn from a Maxwellian with $\sigma = 265\,\mathrm{km\,s^{-1}}$.

This system is representative of O-type stars that undergo case B mass transfer, which often produce ejected secondary stars. Moreover, case B mass transfer is the most common form of mass transfer in this regime(e.g., Kippenhahn & Weigert, 1967; van den Heuvel, 1969), and this system produces an ejected stellar companion with a typical walkaway star velocity (e.g., Renzo et al., 2019).

3 Results

In this Section we report the pre-SN evolution of the binary (Section 3.1), compare the ejection velocities of secondary stars as calculated by different population synthesis codes (Section 3.2), and explain the relationship between natal kicks and companion ejection velocities in detail (Section 3.3).

3.1 Pre-supernova evolution

The typical evolution of the system is to initiate stable case B mass transfer and unbind after the primary supernova. This evolution is qualitatively consistent across each code, but the quantitative details vary.

We illustrate the typical evolution of this binary in Figure 6.1, for the system evolved in `COSMIC`. The primary star initiates almost fully-conservative mass transfer after 8.3 Myr, as

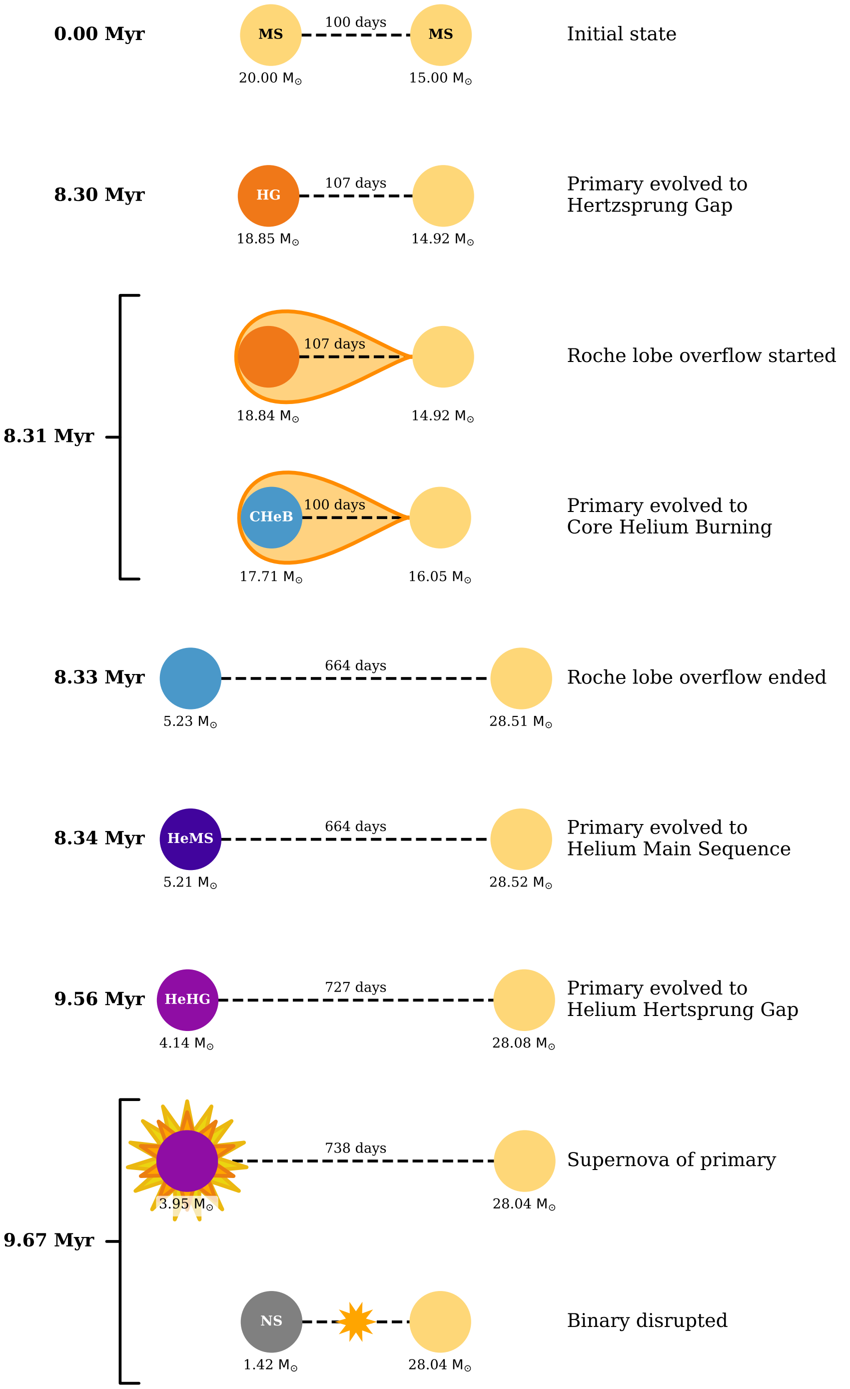

Figure 6.1: A cartoon evolution diagram of the representative binary that we consider. Each star is represented by a circle, where the colour represents the stellar type and their separation is scaled by their orbital period. Each row indicates an important timestep from a `COSMIC` simulation, which is labelled with its time and a description to the sides. Acronyms indicate the stellar type, including main sequence (MS), Hertzsprung gap (HG), core helium burning (CHeB), helium main sequence (HeMS), helium Hertzsprung gap (HeHG) and neutron star (NS).

it expands on the Hertzsprung gap. This transfers 13.6 $\mathrm{M}_\odot$ to the secondary star (increasing its mass to 28.5 $\mathrm{M}_\odot$) and widens the orbital period from ~100 d to ~660 d. The primary proceeds to self-strip via stellar winds, further widening the binary until it reaches core collapse 1.5 Myr later.

The difference in binary physics settings between the codes results in moderately different pre-SN evolutionary states. In `COMPAS` the evolution proceeds similarly, except the mass transfer is much less conservative, such that a large fraction of the transferred mass is lost from the system. As a result, after the mass transfer ends, the companion is much less massive (~18 $\mathrm{M}_\odot$). Therefore, even though and the binary is slightly wider (~800 d) due to the angular-momentum loss from the system, overall increasing the secondary star's orbital velocity increases. For `binary_c`, the mass transfer is almost identical to `COSMIC`; however, the stellar winds on both stars are weaker. Consequently, the binary remains tighter, with $P_{\rm orb} = 555$ d at the end of mass transfer, and each star has a higher mass (5.3 $\mathrm{M}_\odot$ and 29.3 $\mathrm{M}_\odot$ respectively), placing the average ejection velocity between that of `COSMIC` and `COMPAS`. It may be possible to enforce agreement between codes through a detailed selection of binary physics settings but that is beyond the scope of this work.

3.2 Ejection velocity comparison

The ejection velocities of secondary stars closely follow their pre-SN orbital velocities in each code, and have only a weak dependence on the magnitude of natal kicks. In Figure 6.2, we show how the applied natal kick affects the difference between the pre- and post-SN velocity of the secondary star in these simulations. For a circular binary, the pre-SN orbital velocity of the secondary star, $v_{2,\rm preSN}$, is given by:

$$v_{2,\rm preSN} = \frac{m_1}{m_1 + m_2} \cdot v_{\rm orb}, \tag{6.1}$$

where m_1 is the primary mass, m_2 is the secondary mass, and

$$v_{\rm orb} = \sqrt{\frac{G(m_1 + m_2)}{a}}, \tag{6.2}$$

where G is the gravitational constant and a is the semi-major axis. We note that any post-mass-transfer system is typically assumed to be circularised in population synthesis, as is the case for all three codes in this study. The main panels show a 2D kernel density estimator for the disrupted binaries in the 50,000 realisations, while the side panels show marginal histograms for each axis.

The distributions across codes are very similar despite their different pre-SN evolution. The mean ejection velocity in each code (indicated by the white crosses) aligns directly with the pre-SN orbital velocity. Quantitatively, we find that, for each code, the ejection velocity of the secondary star is within ~5 km s^{-1} of its pre-SN orbital velocity, despite the natal kick magnitudes spanning 3 orders of magnitude.

The weak dependence of the secondary's ejection velocity on the natal kick applied to the primary is expected. Due to the strength of the natal kicks, the secondary star experiences a near-instantaneous disappearance of the compact object from its orbit, removing the cen-

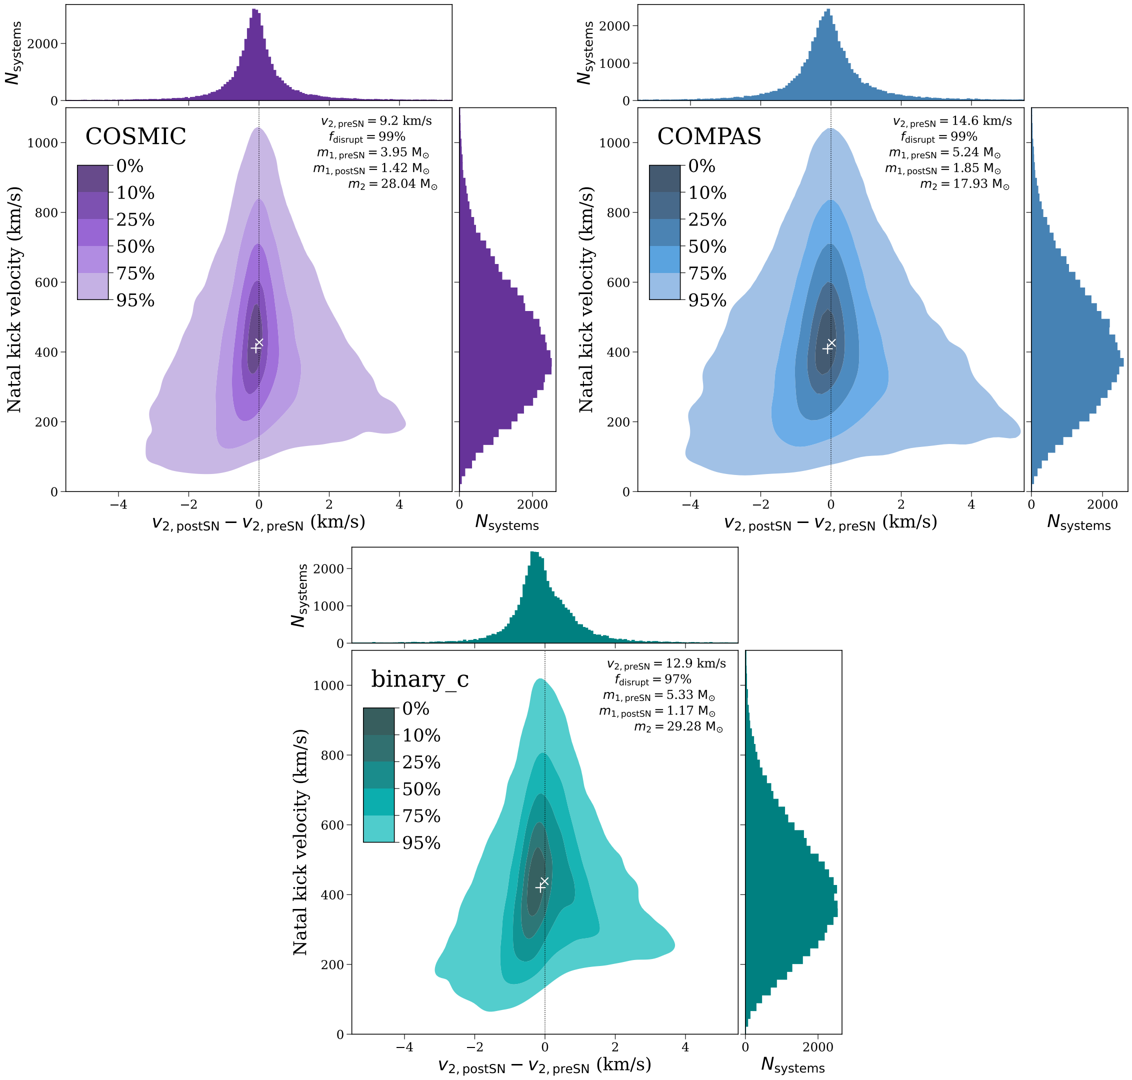

Figure 6.2: Secondary ejection velocities closely follow pre-SN orbital velocities, with extremely weak dependence on natal kick magnitudes, across different population synthesis codes. These plots compare a binary with identical initial conditions ($m_1 = 20\,\mathrm{M}_\odot$, $m_2 = 15\,\mathrm{M}_\odot$, $P_{\mathrm{orb}} = 100\,\mathrm{days}$, $e = 0$, $Z = 0.02$), evolved 50,000 times with different random natal kicks in three codes. For each code, the main panel shows shows the difference between the secondary ejection velocity and its pre-SN orbital velocity on the x-axis, while the y-axis shows the magnitude of the natal kick applied to the primary star. Each main panel shows a 2D kernel density estimator with a colourbar indicating the fraction of the population contain with each contour. Secondary panels show marginal histograms. Small white crosses and pluses indicate the mean and median of each distribution respectively. Panels are annotated with the population synthesis code used for evolution and information about the pre- and post-SN system.

tripetal acceleration and leading to its ejection at close to its pre-SN orbital velocity (see **?** for a more in-depth discussion).

In addition to the ejection velocity, the fraction of instances of the binary that are disrupted, $f_{\rm disrupt}$, is very similar. We find that $f_{\rm disrupt} = 99\%$, 99% and 97% for `COSMIC`, `COMPAS`, and `binary_c` respectively. For a fixed kick angle, the likelihood of disruption is driven by the ratio of the kick velocity to the pre-SN orbital velocity (e.g., Kalogera, 1996). Therefore, the lower disruption fraction for `binary_c` is likely due to the binary being tighter at the onset of core collapse (see Section 3.1).

The overall agreement between the codes validates their implementation of the SN kick routines. Although each code produces a markedly different pre-SN evolution of the system (compare the annotations in each plot in Figure 6.2), the relationship between the pre- and post-SN velocities holds.

3.3 The relationship between natal kick magnitude, angle and companion ejection velocities

The shape of the distributions in Figure 6.2 can be understood through a consideration of the magnitude and angle of the kick, in concert with the geometry of the system, as shown by several earlier works. In Figure 6.3, we show the individual disrupted systems from the `binary_c` simulation as scatter points. In the left panel, each point is coloured by the azimuthal angle of the kick ω, while the right panel is coloured by the polar angle of the kick φ. We include a schematic for the geometry of the system in the upper left corners, which shows the direction of ejection of the companion with an arrow and a coloured wheel indicating the value of ω. As we noted above, the ejection velocity distribution is centred on the pre-SN orbital velocity of the companion star.

There are several features in Figure 6.3 that require explanation: (1) the scatter around the mean ejection velocity narrows for strong natal kicks, (2) certain azimuthal angles produce higher or lower ejection velocities, and (3) for weaker kicks, there is a dearth of systems with ejection velocities above the pre-SN orbital velocity and a general skew to lower ejection velocities.

Firstly, stronger natal kicks lead to a tighter relationship between the pre- and post-supernova velocities of the secondary star. This tightening is because a stronger natal kick more immediately removes the compact object from the system. As a result, the compact object is less able to influence the trajectory of the ejected companion star.

Secondly, the azimuthal angle of the kick influences the difference between ejection velocity and the pre-supernova velocity of the secondary star. Figure 6.3 indicates that the strongest accelerations of the companion occur for angles close to $\omega = 5\pi/6$ (light purple), while the strongest decelerations occur for angles close to $\omega = \pi/3$ (blue). From the geometry of the system, we can see that these angles respectively correspond to a trajectory that leads to the compact object being mostly ahead, or mostly behind, the companion ejection direction. For a compact object that is ahead of the direction of the companion, its will pull the companion towards it, hence the acceleration at these angles, while the inverse reasoning applies for the deceleration. For the same reason, natal kicks pointed directly away from the companion ($\omega = 3\pi/2$) have the weakest effect on the ejection velocity of the companion.

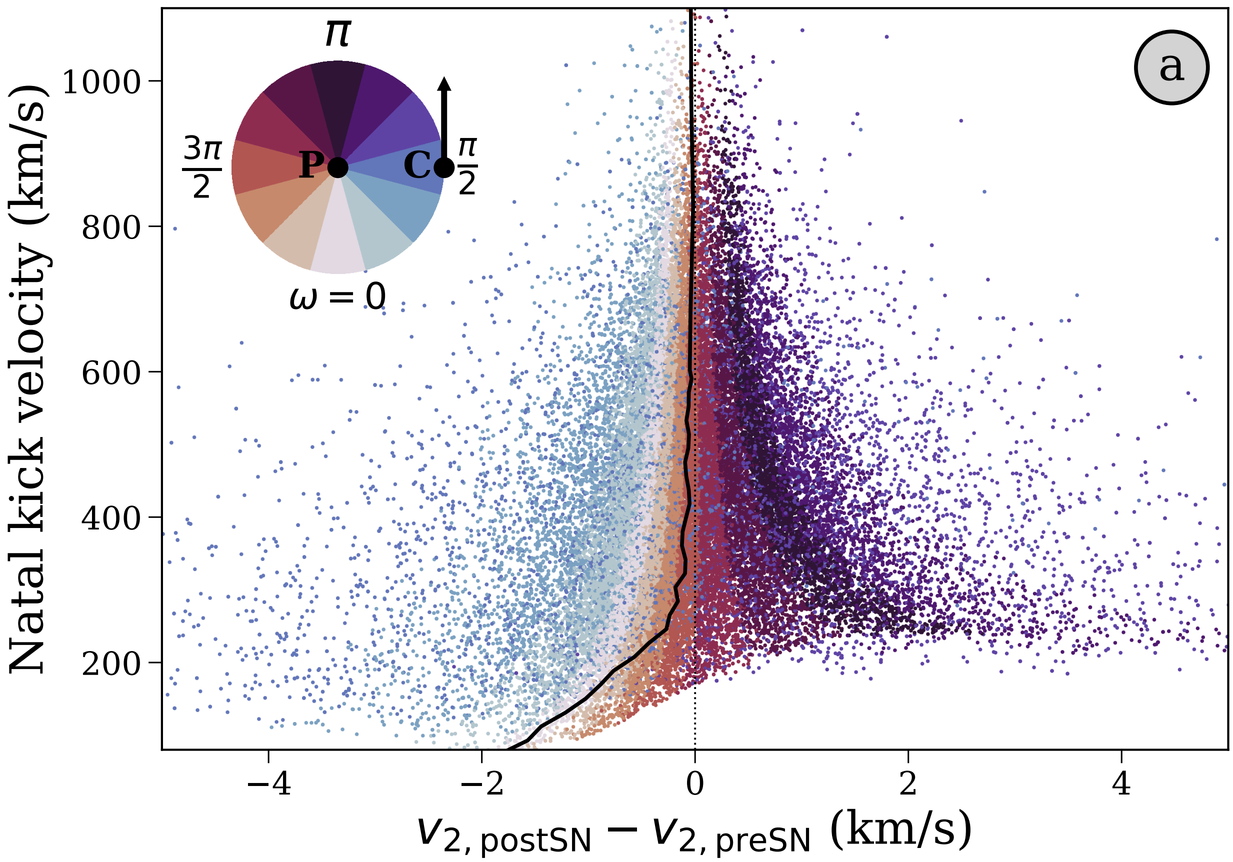

Figure 6.3: The slight variances in the ejection velocity of the secondary star are due to the angle of the natal kick, which are exaggerated for weaker, in-plane kicks. The scatter points show the 50,000 different realisations of the binary run with `binary_c`, each of which are coloured by the azimuthal angle of the kick ω. The solid black line shows the median difference between the pre- and post-SN velocity of the secondary star as a function of the natal kick magnitude. In the upper left corner we include a guide for the geometry of the system, where P is the primary (exploding) star, C is the (ejected) companion star, the arrow indicates the pre-SN orbital velocity of the companion and the wheel of colour indicates the azimuthal angle of the natal kick applied to the primary. The dearth of points in the lower right is where systems remain bound for weak natal kicks with an angle counter to the compact object's orbital motion.

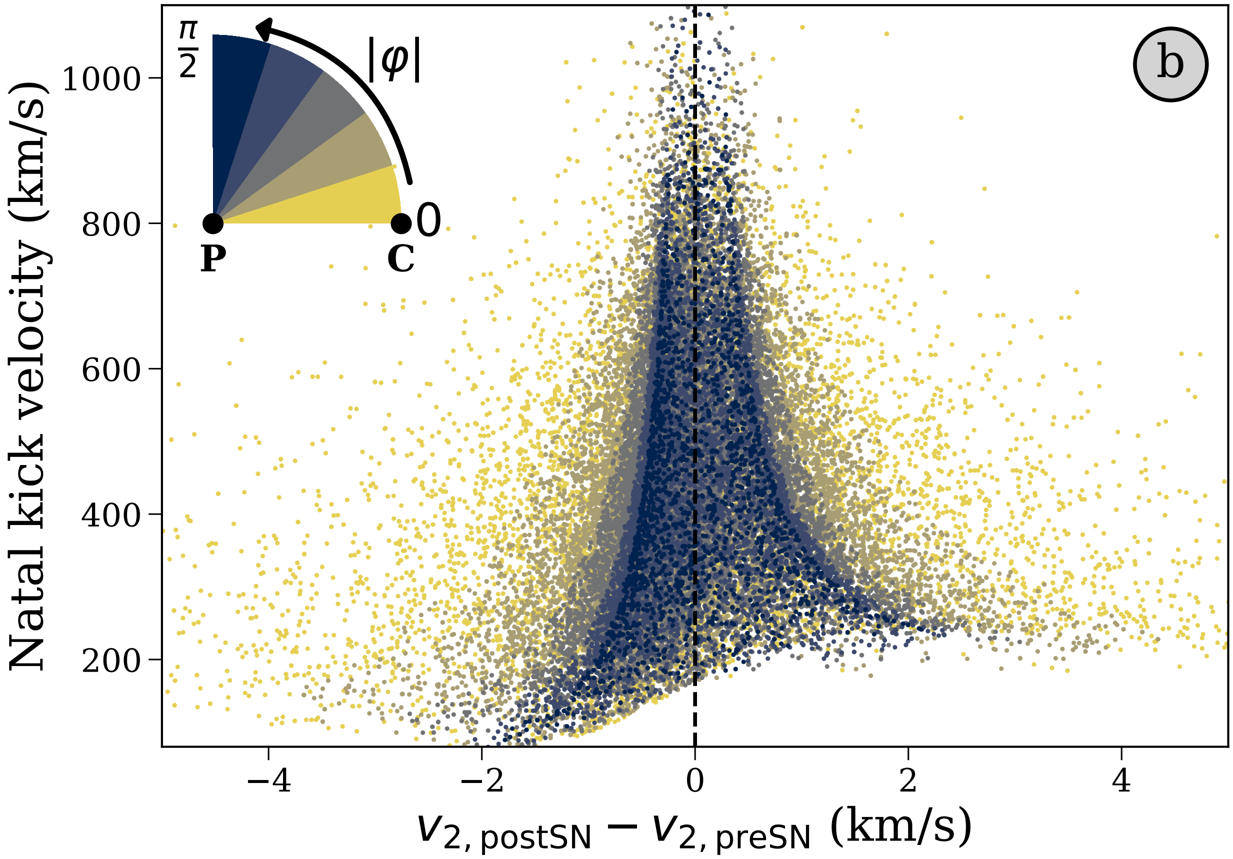

Figure 6.4: The slight variances in the ejection velocity of the secondary star are due to the angle of the natal kick, which are exaggerated for weaker, in-plane kicks. Similar to Figure 6.3, except now each point is coloured by the absolute value of the polar angle of the kick φ, where $\varphi = 0$ indicates an in-plane kick. In the schematic in this plot the orbital motion of the stars is directed in and out of the page.

Lastly, for weaker kicks ($\lesssim 250\,\mathrm{km\,s^{-1}}$) the shape of the distribution shifts because the approximation that the compact object near-instantaneously leaves the system no longer holds. In this regime, a kick directed away from the companion, or behind its ejection trajectory (i.e. $\omega < \pi/2$ or $\omega > 3\pi/2$) can allow the compact object to more significantly decelerate the companion before leaving the system. This is notable in the shift of the distribution to more negative values, meaning an ejection slower than the pre-SN orbital velocity of the companion, for weak kicks. Furthermore, a weaker kick with an angle counter to the compact object's orbital motion (i.e. $\pi/2 \leq \omega \leq 3\pi/2$) can simply decelerate the compact object, tightening the binary, but leaving it bound. The binary remains bound in this way only for approximately 1%, 1% and 3% of kicks applied to the system, for `COSMIC`, `COMPAS` and `binary_c` respectively. The bound systems leave a gap in the distribution of disrupted systems in the lower right of Figure 6.3. As a result of these effects, the median difference between the pre- and post-SN velocity of the secondary star is negative for weak kicks, as shown by the solid black line in Figure 6.3.

Each of these effects are most prominent for kicks that are directed within the orbital plane of the binary. In Figure 6.4, we instead colour each point by the angle of the kick relative to the orbital plane. In-plane kicks ($\varphi = 0$, yellow) produce the full range of deviations from the pre-SN orbital velocity. These kicks can effect a greater change on the ejection velocity because they allow the compact object to remain closer to the companion, and thereby exert a greater gravitational influence, for a longer period of time. For kicks directed perpendicular to the orbital plane ($|\varphi| = \pi/2$, dark blue), the post-SN influence of the compact object is always reduced. Therefore, the maximal deviation from the pre-SN orbital velocity is decreased, with almost no perpendicular kick instances deviating by more than $2\,\mathrm{km\,s^{-1}}$.

4 Correction to Kiel & Hurley (2009)

Earlier versions of `COSMIC` (<v3.5.0) used the methodology described by Kiel & Hurley (2009), hereafter KH09, in calculating the post-SN orbital properties of binaries, as well as the ejection velocities of compact objects and their prior companions. However, when using this method, we found that the ejection velocities of companions were not weakly correlated with their pre-SN orbital velocities as one would expect. Instead, their ejection velocity showed a strong (unphysical) dependence on the natal kick received by the primary. Moreover, the ejection velocity of the compact object was often significantly slower than the natal kick that it received.

We demonstrate the (unphysical) strong dependence of the ejection velocity on natal kick magnitude in Figure 6.5. The grey contour plot shows the distribution of ejection velocities and natal kick magnitudes for our representative binary, evolved 50,000 times using the original KH09 prescription. The triangular shape of the contour indicates that stronger natal kicks also lead to faster ejection velocities. Moreover, the marginal distribution in the upper panel of Figure 6.5 shows there is a wide spread in ejection velocities, extending up to 4 times the pre-supernova orbital velocity. This dependence is unphysical, as the ejection velocity of the companion should be closely related to its pre-SN orbital velocity.

The reason for these effects is that, in KH09, the z-component of the final velocity of each binary constituent neglects the contributions from the velocity at infinity of the hyperbolic orbit, V_∞. The derivation assumes that this velocity is directed entirely within the orbital plane of the binary, but does not fully account for the fact that the orbital plane is no longer aligned with the z axis after the SN. Therefore, the z-component of the final velocity of the compact object (equation 17 from KH09) should be as follows

$$V_{1\mathrm{fz}} = \frac{M_{\mathrm{NS}}}{M_b'} V_{\mathrm{kick}} \sin\phi \ \underbrace{+ V_\infty \sin\nu \sin\gamma \frac{M_2}{M_b'}}_{\text{Term missing from KH09}}, \tag{6.3}$$

where the variables are defined in the same way as KH09, such that M_{NS} is the mass of the resulting compact object, M_b' is the final total mass of the system, V_{kick} is the magnitude of natal kick, ϕ is the polar angle of the natal kick, V_∞ is the velocity at infinity of the hyperbolic orbit and ν and γ are angles accounting for the misalignment of the orbital plane

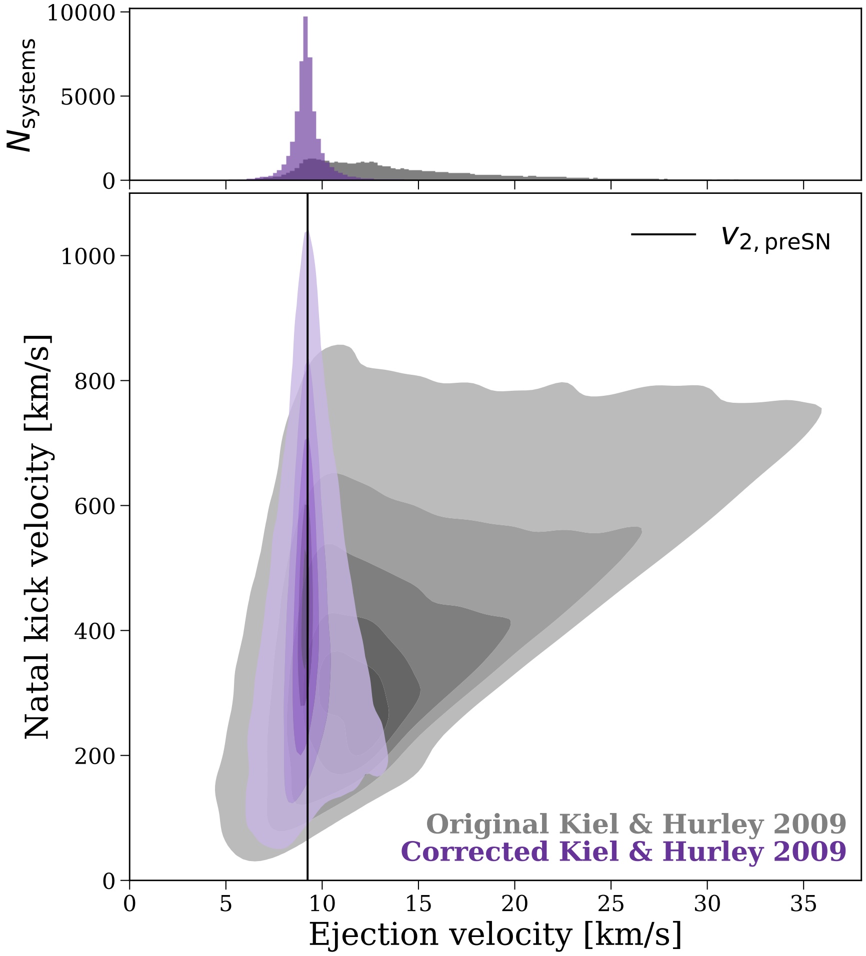

Figure 6.5: Our correction to Kiel & Hurley (2009) removes an unphysical dependence of the ejection velocity of secondary stars on the magnitude of the primary star's natal kick. The main panel shows kernel density estimators for the 2D density distributions run in `COSMIC`, with and without our correction. The upper panel shows the marginal distribution for the ejection velocity.

after the SN (see Section 2.2.1 and Figure 2 of KH09). Similarly, the z-component of the final velocity of the companion (equation 20 from KH09) should be as follows

$$V_{\rm 2fz} = \frac{M_{\rm NS}}{M_b'} V_{\rm kick} \sin\phi \underbrace{-V_\infty \sin\nu \sin\gamma \frac{M_{\rm NS}}{M_b'}}_{\text{Term missing from KH09}} . \qquad (6.4)$$

We have confirmed that including the additional V_∞ terms in these equations removes the unphysical dependence on natal kick magnitude. In Figure 6.5, the purple distributions correspond to the corrected version of KH09, which now strongly correlates with the pre-supernova orbital velocity. This change brings the KH09 method into agreement with that of Pfahl et al. (2002) and **?**.

The latest versions of `COSMIC` (≥v3.5.0) now use the Pfahl et al. (2002) method by default.

We make the original Kiel & Hurley (2009) method accessible (via negative `kickflag` inputs) so that earlier `COSMIC` results can be reproduced. We note that this issue only affected the ejection velocities of stars and compact objects from disrupted systems, but the *conditions* for disruption and post-SN orbital parameters for *bound* systems were correct. Therefore prior work using `COSMIC` to explore bound systems remains unchanged.

5 Summary and importance of open-science comparisons

We have compared the ejection velocities of stellar companions for a representative case B mass-transferring binary across 3 independent open-source rapid population synthesis codes. We have confirmed that each code is consistent in its finding that there is only a very weak dependence of the stellar ejection velocity on the SN natal kick, as one would physically expect. We explain the influence of the direction and magnitude of natal kicks in slightly shifting the companion's ejection velocity from its pre-SN orbital velocity and provide a correction to an earlier prescription for SN kicks.

It is notable that, during the course of this project, we identified problems with either the kick prescription or implementation in each of the codes that we considered. As discussed above, `COSMIC` previously used an erroneous kick prescription, which has been resolved as of v3.5.0[1]. The implementation in `COMPAS` had anomalous component velocities such that the ejection velocity of secondary stars was consistently overestimated by a factor of 2. This problem arose as a result of spurious vector indexing and incorrect rotation algebra in the code. This issue has been resolved as of `COMPAS` v03.09.02[2]. In `binary_c`, the prior implementation did not sample natal kick magnitudes and azimuthal angles in an independent manner, which led to slightly lower ejection velocities on average. This correlated sampling occurred as a result of re-using a randomly generated number from the kick magnitude in the azimuthal angle sampling. This issue has been resolved in `binary_c` for versions after v2.24[3].

Despite the rigorous testing suites of each code that aim to address issues such as these, each bug was not uncovered until comparison with other open-source codes. It is therefore evident that open science, and comparison projects (e.g., Toonen et al., 2014; O'Connor et al., 2018), are essential to identifying and resolving subtle issues in population synthesis.

---

[1] https://github.com/COSMIC-PopSynth/COSMIC/pull/678

[2] https://github.com/TeamCOMPAS/COMPAS/pull/1290

[3] https://gitlab.com/binary_c/binary_c/-/commit/fba20687984e9cfb46bc38838d07655bfc99c6e5

CHAPTER 7

THE IMPACT OF BINARY INTERACTIONS ON THE TIMING AND LOCATION OF SUPERNOVA FEEDBACK

If you're always on time, it implies that you never have anything better you should be doing.

– Kelsier, in *The Final Empire*

So do all who live to see such times, but that is not for them to decide. All we have to decide is what to do with the time that is given us.

– Gandalf, in *The Fellowship of the Ring*

1 Introduction

Stellar feedback from massive stars plays a critical role in galaxy formation and evolution. Without this feedback, simulated galaxies tend to form too many stars and become too dense compared to observed galaxies (e.g., Katz et al., 1996; Hopkins et al., 2018a). While several aspects of stellar evolution are likely to contribute to the regulation of star formation, supernovae (SNe) from massive stars are undoubtedly one of the most important (e.g., Dekel & Silk, 1986; Hopkins et al., 2012; Somerville & Davé, 2015; Naab & Ostriker, 2017).

In the broadest picture of supernova (SN) feedback, the energy of the SN couples to the surrounding gas, driving it to lower densities and higher temperatures that are unfavourable for forming stars. However, the actual process is far more nuanced, with a large body of work clearly demonstrating that the timing and location of each SN has a significant impact on the efficacy of its feedback (e.g., Walch et al., 2015; Girichidis et al., 2016; Hu et al., 2016, 2017, 2019; Smith et al., 2021; Orr et al., 2022). In addition to a strong dependence on SN clustering, which affects how "collectively" SNe can impact the surrounding gas, much depends on the density of the surrounding gas where the SN explodes. SNe that occur near their birthplace are likely to directly disrupt their natal molecular clouds, rapidly shutting down star formation. In contrast, SNe that happen outside molecular clouds may have less direct effect on the efficiency within clouds, but have a greater chance of driving outflows (e.g., Ceverino & Klypin, 2009; Ceverino et al., 2014; Zolotov et al., 2015; Hu et al., 2017; Andersson et al., 2020; Steinwandel et al., 2023). At the extremes, massive runaway stars can be ejected far out of the plane of the galaxy (e.g., Renzo et al., 2019), which are much more effective at boosting the momentum outflow rate of galaxies than SNe which are dispersed within the disc, since vertically displaced SN can more effectively couple to the hot phase of pre-existing wind outflows (Steinwandel et al., 2023).

Simulations have additionally already demonstrated the importance of the timing of SNe for their relative impact on a galaxy (Struck-Marcell & Scalo, 1987; Parravano, 1996; Quillen & Bland-Hawthorn, 2008). In particular, more realistic stellar lifetime distributions, as opposed to fixed values for all stars, can prevent the formation of high density gas clouds and allow more feedback to occur in low density environments (Kimm et al., 2015), as well as reduce the overall clustering of SNe (Smith et al., 2021; Hu et al., 2023).

One important mechanism for setting the timing and spatial distribution of SNe are interactions during binary stellar evolution. The vast majority of massive stars that reach core collapse and produce SNe are born in binaries and multiple systems (e.g., Mason et al., 2009; Almeida et al., 2017; Moe & Di Stefano, 2017; Offner et al., 2023). A large subset of these stars will interact with their companion in their lifetime (e.g, Sana et al., 2012; de Mink et al., 2014).

The interactions during binary evolution can have a number of different impacts on the timing of core-collapse SNe. For example, adding or removing mass during mass transfer alters a star's evolutionary timescale and thus its time until core collapse, since more massive stars evolve more quickly. In the case of stellar mergers, two low mass stars that merge will reach core collapse much later than single stars of the same mass (De Donder & Vanbeveren, 2003; Zapartas et al., 2017). Rejuvenation of accretor stars can also have strong impacts on the timing of SNe, such that additional hydrogen is mixed into their cores, potentially altering their time to core collapse by ∼Myrs (e.g., Neo et al., 1977; Schneider et al., 2016; Renzo et al., 2023), though this effect may be lower for different assumptions regarding convective boundary mixing (e.g., Braun & Langer, 1995).

Binary interactions can also displace massive stars from their birth sites. First, the above changes in SN timing has a first-order effect on how close a SNe is to its birth place, where SNe delayed as a result of binary evolution will spend a longer time dispersing from their parent cluster (e.g., Aghakhanloo et al., 2017). Second, binary evolution can directly change the kinematics of individual stars. Secondary stars that were formed in binaries that are disrupted by a primary SN may travel large distances as runaway stars before a SN occurs (e.g., Blaauw, 1961; Boersma, 1961; Renzo et al., 2019).

Despite the fact that binary interactions are likely important for the timing and locations of massive star SNe, most hydrodynamical zoom-in simulations currently implicitly assume all stars evolve without a companion (e.g., Leitherer et al., 1999, 2014; Hopkins et al., 2018a, 2023a; Applebaum et al., 2021; Smith et al., 2021; Christensen et al., 2024). As such, late SN feedback is not accounted for and the feedback is always assumed to occur at the position of the star particle.

While the timing of SNe and the ejection velocities of their progenitors can be addressed entirely within the context of binary evolution models, understanding the subsequent impact of the resulting SNe requires embedding the evolving binaries within realistic galactic potentials. Earlier work has established the likelihood that runaway stars ejected from binaries may end their lives far from their birth locations (e.g., Eldridge et al., 2011; Renzo et al., 2019), but without considering the restraining forces provided by a galactic potential. Similarly, previous studies have considered the impact of binary physics on the timing of SNe and the delays that are possible (De Donder & Vanbeveren, 2003; Zapartas et al., 2017).

We build upon these works by leveraging the capabilities of the new open-source code `cogsworth` (Chapter 5), which provides a framework for performing self-consistent population synthesis and galactic dynamics simulations. In this way, one can evolve binary stars within their galactic context, integrating their orbits through the galaxy while accounting for the effects of SN natal kicks on a star's galactic trajectory. This enables us to track the precise time and galactic location of each SN in a galaxy. `cogsworth` also takes advantage of the advances in numerical galaxy simulations to draw realistic populations in age, position and local potential.

This paper is structured as follows. In Section 2, we outline our methods for simulating populations of SNe with `cogsworth` and the settings chosen for our fiducial model. We demonstrate and explain the impact of binary interactions on SN location and timing in Section 3. We explore how our results depend on our choice of initial conditions, binary physics and galaxy parameters in Section 4. In Section 5, we present analytic fits to our distributions for the timing and distance of SNe. We discuss the implications and limitations of our work in Section 6 and draw conclusions in Section 8. All code for producing our simulations, as well as the simulations themselves, is available on GitHub[1] and Zenodo[2].

---

[1] `https://github.com/TomWagg/supernova-feedback/`

[2] `https://doi.org/10.5281/zenodo.15273993`

2 Simulating binary stellar populations with `cogsworth`

We use `cogsworth` (Chapter 5) to quantify the impacts of binary interactions on the timing and location of SNe. `cogsworth` uses `COSMIC` (Breivik et al., 2020b) to rapidly synthesize populations of binary stars and `gala` (Price-Whelan, 2017; Price-Whelan et al., 2024) to integrate the subsequent orbits of the stars in model galactic potentials. A detailed description of `cogsworth` can be found in Chapter 5, but we give a brief overview of our approach here; further details are presented below in Section 2.2.

Here we use `cogsworth` to simulate the most recent 150 Myr of star formation in a galaxy. We do so by replacing newly-born star particles in hydrodynamical galaxy simulations with an equivalent mass cluster of binary stars. We then rapidly evolve the individual binary stars while simultaneously integrating their orbits through the galaxy, using a galactic potential fit to the galactic mass distribution in the hydrodynamical simulation. We then record the time and location of each core-collapse SN in the galaxy, evolving each binary until 200 Myr beyond present day, such that the oldest binaries are evolved for a total of 350 Myr; these timescales capture a reasonable fraction of the recent star formation whilst also allowing enough time for every star to reach core collapse. We evolve only binaries in which the initially more massive star has a mass of at least $4\,\mathrm{M}_\odot$, since less massive binaries are unlikely to produce a core-collapse SN, even through mergers.

2.1 Initial stellar distributions & orbital evolution

For our fiducial model, we use initial conditions from the `FIRE`-2 galaxy `m11h` (Hopkins et al., 2018a; El-Badry et al., 2018), which is a dwarf galaxy, similar in mass to the LMC ($M_\star = 4\times10^9\,\mathrm{M}_\odot$), with a strong disc component and a slightly super-solar typical metallicity of 0.017. The star formation history is reasonably uniform over the last 150 Myr, with $\sim$0.5 $\mathrm{M}_\odot$ formed per year. We focus our study on `m11h` since earlier work has shown the effect of spatially distributed runaway star feedback is more prominent in dwarf galaxies (e.g., Steinwandel et al., 2023). This simulation has a force resolution identical to its spatial resolution, which is fully adaptive and has a minimum length scale on the order 0.7 pc (Hopkins et al., 2018b). However, the spatial scale resolved at the star formation density threshold is approximately 7 pc. We identify the initial position and kinematics of every star particle by reverse integrating it through the galactic potential.

When generating binary populations, we assume that each of the $\sim$7000 $\mathrm{M}_\odot$ star particles in the simulation represents a stellar cluster of radius 3 pc, which is typical of open clusters and similar to the radius of Orion Nebular Cluster (e.g., Kroupa et al., 2018). We use this radius to adopt a Gaussian spread in initial position for each binary relative to the location of its parent star particle. This is similar to methods used in previous works (e.g., Sanderson et al., 2020).

We adopted an initial velocity dispersion of $1.7\,\mathrm{km\,s^{-1}}$ for the binaries. This corresponds to a cluster of $10^4\,\mathrm{M}_\odot$ (the mass of a star particle in `m11h`) with virial parameter of $\alpha_{\rm vir} = 1$, which means that the cluster is initially gravitationally bound (Bertoldi & McKee, 1992). This velocity dispersion approximately follows measurements of the velocity dispersion of the Orion Nebula Cluster (Da Rio et al., 2017; Kroupa et al., 2018; Kuhn et al., 2019).

Parameter	Symbol	Fiducial model	Variations	References
Initial conditions				
Initial mass function slope	$\alpha_{\rm IMF}$	-2.3	$[-1.9, -2.7]$	Kroupa (2001); Schneider et al. (2018b,a)
Mass ratio slope	κ	0	[-1, 1]	Mazeh et al. (1992); Sana et al. (2012)
Orbital period slope	π	-0.55	[-1, 0]	Sana et al. (2012); de Mink & Belczynski (2015)
Initial orbital period limit	$P_{0,\rm max}$	$10^{5.5}$ d	10^3 d	de Mink & Belczynski (2015)
Eccentricity slope	η	-0.45	-	Sana et al. (2012)
Binary fraction	$f_{\rm bin}$	1.0	0.0	
Metallicity	$\bar{Z}$	$\bar{Z}_{\rm m11h} \approx 1.2 Z_\odot$	$[0.5, 0.2, 0.1, 0.05]\bar{Z}_{\rm m11h}$	Wetzel et al. (2023)
Binary physics				
Mass transfer efficiency	β	$\beta \propto 10\tau_{\rm th,acc}$	[0, 0.5, 1]	Schneider et al. (2015)
Case B critical mass ratio	$q_{\rm crit,B}$	HW1987	$[0, \infty]$	Hjellming & Webbink (1987)
Common-envelope efficiency	$\alpha_{\rm CE}$	1.0	[0.1, 10.0]	Webbink (1984); de Kool (1990)
CCSN natal kicks	σ	$265\,{\rm km\,s^{-1}}$	$20\,{\rm km\,s^{-1}}$	Hobbs et al. (2005); Igoshev (2020)
ECSN/USSN natal kicks	$\sigma_{\rm low}$	$20\,{\rm km\,s^{-1}}$	$265\,{\rm km\,s^{-1}}$	Hobbs et al. (2005); Igoshev (2020)
Black hole kicks	-	Fallback limited	No rescaling	Fryer et al. (2012)
Galaxy settings				
Hydrodynamical simulation	-	`FIRE m11h`	`ChaNGa r442`	El-Badry et al. (2018); Keith et al. (2025)
Velocity dispersion	$v_{\rm disp}$	$v_{\rm disp} \propto \alpha_{\rm vir}$	$[0.5, 5]\,{\rm km\,s^{-1}}$	Bertoldi & McKee (1992)

Table 7.1: A summary of the fiducial choices (and variations) of the settings for the simulations in this work.

When replacing star particles with collections of individual stars, we use `cogsworth` to sample binary stellar populations that match the total mass, metallicity, and birth time of each star particle formed in the most recent 150 Myr of the simulation.

After formation, the positions of new stars are evolved within a galactic potential matching that of the original simulation. `cogsworth` calculates the potential using the self-consistent field method implemented in `gala` based on Hernquist & Ostriker (1992) and Lowing et al. (2011), which fits the galactic mass distribution (accounting for the contributions from stars, gas and dark matter) using a basis function expansion in spherical harmonics. For a more detailed explanation of the full implementation in `cogsworth`, see Chapter 5.

Within the potential, and using the initial stellar kinematics, we integrate the galactic orbit of each single or binary system with an adaptive integration scheme (Dormand & Prince, 1980), tracking individual stars if a binary is disrupted by a SN.

2.2 Fiducial binary model

In this Section we describe our fiducial settings, which are also summarised in the third column of Table 7.1. The fourth column lists the variations of these parameters, which we discuss in Section 4.

Initial population sampling — We draw masses of the primary (initially more massive) star in each binary, m_1, following the broken power law initial mass function (IMF) from Kroupa (2001), such that $p(m_1) \propto m_1^{\alpha_{\rm IMF}}$ and sampling only stars with $m_1 > 4\,{\rm M}_\odot$. The high mass slope (for stars with $m_1 > 1\,{\rm M}_\odot$) of this IMF has $\alpha_{\rm IMF} = -2.3$. We sample mass ratios, $q \equiv m_2/m_1$, uniformly in $[q_{\rm min}, 1]$, where $q_{\rm min}$ is set such that the pre-main sequence lifetime of the secondary is not longer than the full lifetime of the primary if it were to evolve as a single star (Mazeh et al., 1992; Goldberg & Mazeh, 1994). Eccentricities, e, are sampled following Sana et al. (2012) such that $p(e) \propto e^{-0.45}$ with $e \in [0, 0.9]$, where the upper limit is chosen to avoid Roche lobe overflow at pericenter. For orbital periods, P, we follow the distribution of de Mink & Belczynski (2015), which is an extrapolation of Sana et al. (2012) to lower masses, such that $p(P) \propto (\log_{10}(P/\,{\rm days}))^{-0.55}$ with $\log_{10}(P/\,{\rm days}) \in [0.15, 5.5]$. We use a fixed binary fraction of 100% since we simulate only massive stars which are almost all formed in binaries (e.g., Offner et al., 2023). Additionally, though we do not include truly single stars, our upper orbital period limit allows for a fraction of stars to be effectively single, such that they have no interactions with companions. Metallicities and birth times are set based on the star particles from the galaxy that is postprocessed.

Binary physics — `cogsworth` uses `COSMIC` for binary population synthesis. `COSMIC` is based on the `BSE` code (Hurley et al., 2000, 2002), which uses fitting formulae from Tout et al. (1997) to the single star models of Pols et al. (1998)), but with extensive modifications and updated prescriptions based on more recent work (see Section 3 of Breivik et al., 2020b). Our simulations make use of `COSMIC` v3.4.16 and use the default settings for that version. In particular, we assume the efficiency of mass transfer, $\beta \equiv \Delta M_{\rm acc}/\Delta M_{\rm don}$, is such that the amount of mass accreted during Roche-lobe overflow is limited to 10x the thermal rate of the accretor for main sequence, Hertzsprung gap and core helium burning stars and unlimited for giant branch stars (e.g., Kippenhahn & Weigert, 1967; Schneider et al., 2015). We assume that mass lost during Roche-lobe overflow is lost from the system via isotropic re-emission (as if it is a wind from the secondary) (e.g., Massevitch & Yungelson, 1975).

We determine the stability of mass transfer using the critical mass ratios defined in Hurley et al. (2002) for all stellar types except for giant branch stars, for which we use Hjellming & Webbink (1987). For unstable mass transfer we assume common-envelope events follow the α-λ prescription (Webbink, 1984; de Kool, 1990) and by default take $\alpha_{\rm CE} = 1$ and use the λ prescription from Claeys et al. (2014).

Asymmetries in SN explosions impart a natal kick on the compact object that is formed, which may unbind a binary orbit, thereby ejecting a secondary star. We assume SN natal kicks are distributed as a double Maxwellian distribution, with one component peaking at $\sigma_{\rm CC} = 265\,{\rm km\,s^{-1}}$ for core-collapse SNe and the other at $\sigma_{\rm low} = 20\,{\rm km\,s^{-1}}$ for electron-capture SNe (ECSN) and ultra-stripped SNe (USSN) (Hobbs et al., 2005; Igoshev, 2020). By default in `COSMIC` ECSN are assumed to occur if the helium core mass is between 1.6 and $2.25\,{\rm M_\odot}$, whilst USSN occur for helium stars that undergo a common-envelope with a compact object companion. The magnitude of SN natal kicks for black holes is modulated based on the fallback mass following Fryer et al. (2012).

3 Results I: Supernovae times and locations in the fiducial model

In this section we consider the behaviour of the fiducial model described above. Throughout, we separate the individual progenitor types into four groups: (1) effectively single stars that had no binary interaction prior to their SN, which we define as no Roche Lobe overflow and less than 5% mass accretion from the stellar winds of a companion; (2) primary stars that were initially the most massive star in a binary that did not undergo a merger; (3) secondary stars that were the less massive partner in a non-merging binary; and (4) merger products that result from two stars merging in a binary prior to core collapse.

In the subsections below we discuss the trends in both timing and distances, based on the plots shown in Figure 7.1 for our fiducial model. These plots show absolute (top row) and cumulative (bottom row) distributions for: the time between the formation of a stellar cluster and the eventual SNe of the member stars (left column); and the distance between the SN and the position that the original stellar cluster would have had at the time of the SN, had it not been replaced with a cluster of evolving binary stars (right column). Each distribution is plotted as a “sandpile diagram”, with contributions from different SN precursors stacked on top of each other. Throughout the paper, we colour-code contributions from effectively single stars in light blue, primary stars in dark blue, secondary stars in green, and merger products in purple.

3.1 Supernova timing

For a single star population, there is a tight relationship between initial mass and time to core collapse, with some small scatter from metallicity effects (e.g., Figure 5 of Hurley et al., 2000). Therefore, one would expect the SN rate to peak at early times and decline with a slope that is set by the convolution of the lifetime–mass relation and the initial mass function.

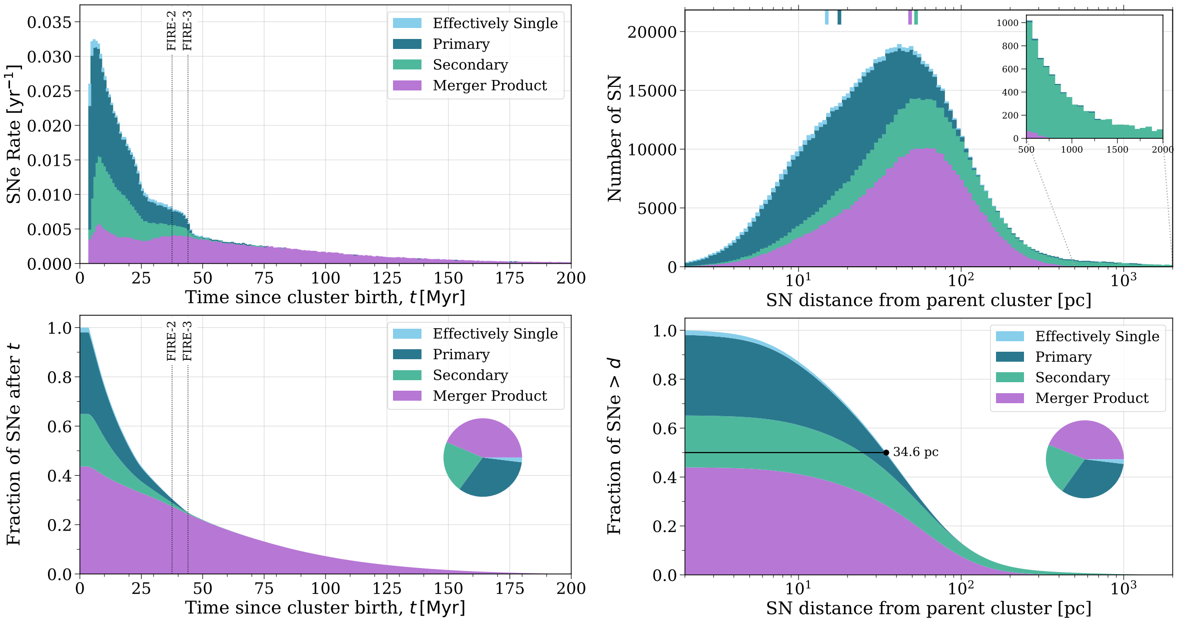

Figure 7.1: Binary interactions can result in delayed SNe and displace SNe far from their parent clusters and molecular clouds. **Top panels:** Stacked histograms separated by progenitor type. Effectively single stars had no binary interaction prior to SN (no Roche Lobe overflow and less than 5% wind mass accretion). Primary (secondary) stars were the initially more (less) massive star in a binary that did not undergo a merger. Merger products are the result of a merger between two stars prior to SN in a binary. **Bottom panels:** Fractional complementary cumulative distribution also separated by progenitor type. Pie charts shows the relative contribution from each progenitor type. **Left:** Histograms show the SN rate as a function of time since cluster birth. Dashed lines indicate time beyond which no SN occur in `FIRE` simulations. **Right:** Histograms show the distance between a SN and its parent cluster at the time of the SN. The inset panel highlights the tail of SNe at extremely long distances. Solid markers on the top axis of the top panel indicate the median value for each progenitor type. Black circle on the bottom panel shows the overall median distance. (Interactive figure available.)

For a binary stellar population, mass transfer and mergers complicate this relationship. The addition or removal of mass during mass transfer will alter a star's evolutionary timescale. In the vast majority of cases, stars that lose mass to their companions will evolve more slowly and so have delayed SNe, with the opposite being true for those that accrete mass (e.g., Pols, 1994). Alternatively, two stars that were initially not massive enough to reach core collapse may merge after their (relatively) slow evolution and later explode, resulting in later SNe (e.g., Zapartas et al., 2017).

In the left panels of Figure 7.1 we show the distribution of SN times relative to cluster birth, separated by SN progenitor types. For a single star population, we would expect core-collapse SNe to cease after the lowest mass star that could reach SNe concludes its evolution. At the average metallicity of `m11h`, this corresponds to a star of $\sim 7\,\mathrm{M}_\odot$ and a time of $\sim$44 Myr. A lower mass star cannot reach core collapse with the default `COSMIC` settings and hence this is the latest time a single star could result in a core-collapse SN in our simulations. We note that the limits of `FIRE`-2 and `FIRE`-3 (shown as dotted lines on Figure 7.1) reproduce this limiting time well.

Critically, binary interactions allow for SNe to occur at later times. Similar to single star evolution models, the distribution in Figure 7.1 shows that the earliest SNe occur around 3.7 Myr and the overall distribution quickly peaks around 6 Myr as the most massive stars reach core collapse. Unlike single star models, however, after $\sim$44 Myr there is a long tail of SNe that is almost entirely from stars that merged before core collapse (purple). These merger product SNe occur later because the two less massive stars that merged evolve on slower timescales that massive stars that can reach core collapse alone.

In addition to the late-time tail from mergers, the age distribution of SNe also shows a transition around $\sim$25 Myr (top left panel of Figure 7.1), which is an imprint of stable mass transfer on SN timing. The most massive stars expand significantly on the main sequence, and often transfer mass to their lower mass partner, delaying their eventual SNe. In contrast, lower mass stars rarely expand enough on their main sequence to initiate mass transfer (e.g., de Mink et al., 2008; Burt et al., 2025). This transition appears in our simulations at roughly $10\,\mathrm{M}_\odot$ where the prevalence and mechanism of mass transfer changes, as shown in Figure 7.2. The transition leads to a break in the distribution at the lifetime corresponding to this mass, approximately $\sim$25 Myr.

More specifically, this feature in the SN age distribution is driven by both the reduced likelihood of mass transfer at lower masses and the relative importance of case A, B, & C mass transfer mechanisms. A primary star that donates a significant amount of mass during its main sequence (case A mass transfer), before developing a significant helium core, lengthens its nuclear timescale, delaying its SN. If the mass transfer occurs later in the star's life (case B or C mass transfer) then it has a lesser effect on the SN timing. We find that primary star SN progenitors with a lifetime longer than around 25 Myr do not undergo case A mass transfer, leading to the feature in the age distribution. We note that the delays occurring as a result of case A mass transfer are strongly dependent on assumptions in population synthesis regarding core evolution and rejuvenation (see Section 6.3).

Due to these two effects — mergers and mass transfer — 25% of SNe in our fiducial simulation occur beyond the cutoff used in the `FIRE`-3 prescription for type II SNe (at 44 Myr), as shown in the lower left panel of Figure 7.1. The progenitors of these later SNe are primarily merger products, though with a small fraction of primaries and secondaries

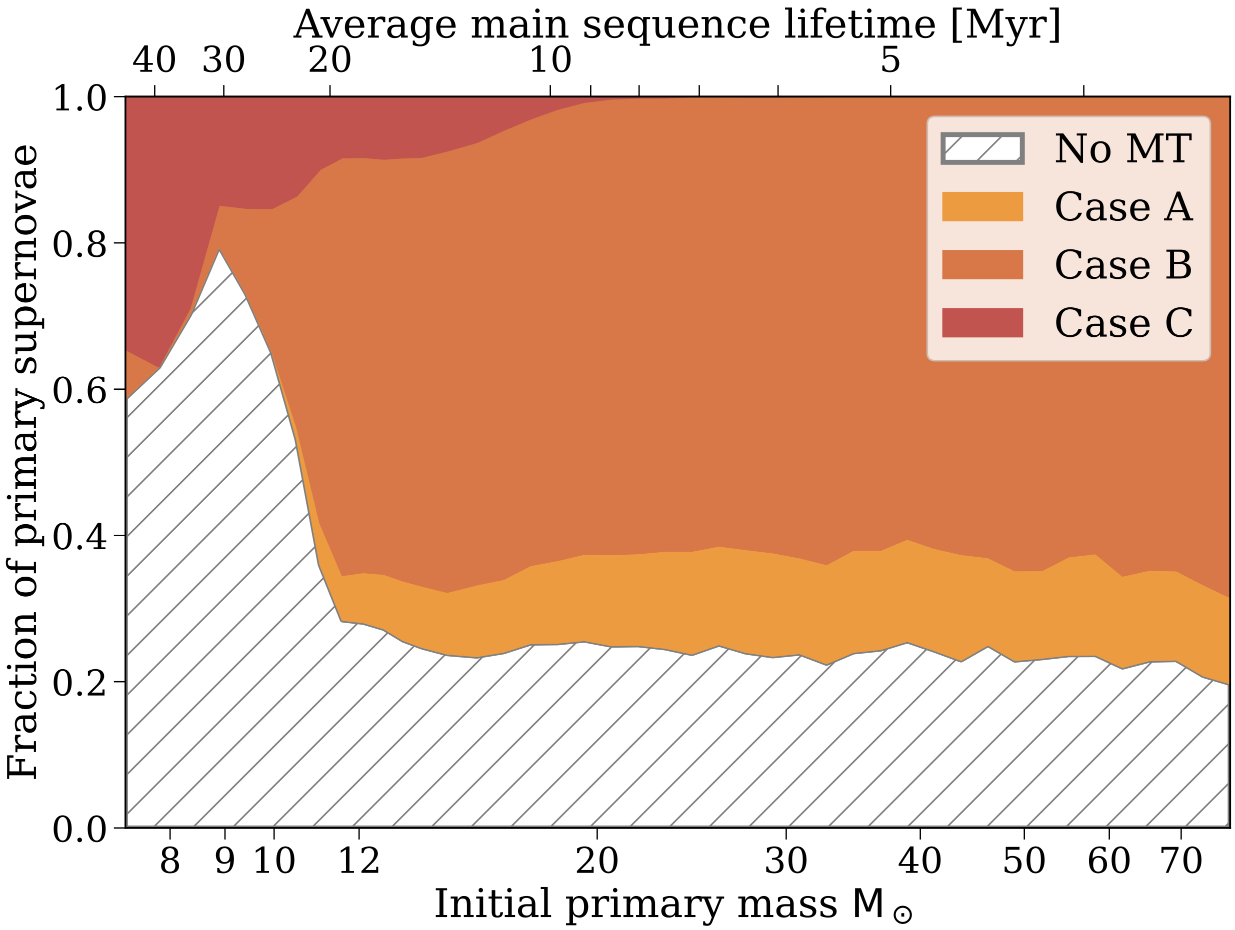

Figure 7.2: The type of mass transfer that a primary star first initiates prior to its core collapse is a function of its initial mass. The transition at $\sim 10\,\mathrm{M}_\odot$ leads to the knee in the distribution in the upper left panel of Figure 7.1. Mass transfer initiated on the donor star's main sequence is case A, while case B corresponds to expansion during the Hertzsprung gap, and case C occurs after the onset of core helium burning. Primary stars below $\sim 10\,\mathrm{M}_\odot$ mostly initiate case C mass transfer, whilst more massive stars initiate case A or B mass transfer. Top axis shows the main sequence lifetime corresponding to the initial primary mass for the average metallicity of `m11h` as calculated by `COSMIC`.

before 50 Myr.

3.2 Supernova separation from parent clusters

In our simulations, there are two mechanisms by which SNe can occur far from their birthplace. First, typical stars are born in stellar clusters, almost all of which rapidly dissolve into the field (e.g., Lada & Lada, 2003; Portegies Zwart et al., 2010) after their natal gas clouds are expelled. This imprints stars with an initial velocity dispersion, leading them to have slightly different orbits through the galaxy, and thus to diverge from the cluster centre-of-mass over time.

The second mechanism is through ejection from a binary. If a binary disrupts after the primary star reaches core collapse, the secondary star is ejected with approximately its orbital velocity (e.g., Wagg et al., 2025; Renzo et al., 2019). These secondaries can then travel large distances as runaway stars before reaching core collapse themselves.

In the right panels of Figure 7.1, we show the distance between each SN and its parent cluster (or simulation star particle)'s centre of mass at the moment of core collapse. SNe from effectively single stars and primary stars are typically located close to the parent cluster. This is expected given that they are short-lived with less time to disperse with their modest $\sim 2\,\mathrm{km\,s^{-1}}$ initial velocity dispersion, and are almost never ejected from binaries. On average, SNe from effectively single or primary stars occur ∼18 pc from their parent cluster and 90% occur within 50 pc.

In contrast, SNe from merger products, which constitute a third of all SNe in the fiducial simulation (e.g., Sana et al., 2012), are located further from their parent cluster at the time they explode. On average, they occur at a separation of ∼49 pc, with 17% occurring beyond 100 pc. This is a result of SNe from merger products happening later (see Section 3.1), giving them more time to disperse from their parent cluster as a result of their initial velocity dispersion (e.g., Aghakhanloo et al., 2017).

The spatial separations are even further extended for SNe from secondary stars, which constitute nearly a quarter of SNe in the fiducial case. These SNe have an additional contribution from their high ejection velocities, and are located at an average separation of ∼54 pc, with 25% occurring beyond 100 pc and 4% beyond 500 pc.

Overall, the average distance between a SNe and its parent cluster is ∼35 pc, which is within the resolution element of `FIRE` simulations at the average ISM density (`FIRE`'s spatial resolution is $\approx 85\,\mathrm{pc}$ for an ISM density of $1\,\mathrm{cm^{-3}}$). However, the distribution has a significant tail, with ∼13% of SNe occurring at a separation of more than 100 pc.

3.3 The joint distribution of SN time delay and separation, for single and for binary stars

In the above discussion of Figure 7.1, we treat SN timing and spatial distributions as independent. In practice, however, these two distributions are coupled; extended SN timescales allow more time for precursor stars to disperse, and specific binary evolutionary pathways can change both timing and velocities, due to mass transfer and kicks from evolving higher-mass partners.

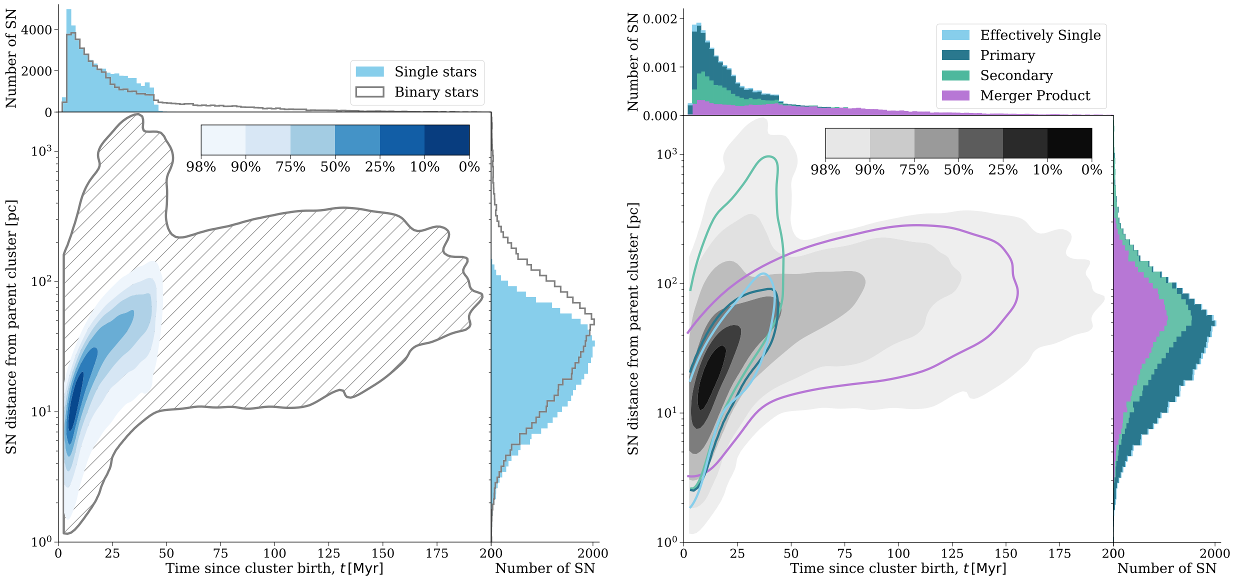

Figure 7.3: Left: Binary stars produce a significantly different distribution of SNe times and locations than single stars. Main panel shows a 2D kernel density estimation of the distribution of SNe for single star population in blue. The grey hatched contour shows the 98% region for a binary star population. Marginal histograms are shown in each side panel. **Right**: The timing and location of SNe relative to their parent cluster are strongly correlated and a function of their progenitor type. Grey filled contours show the full 2D density distribution, whilst coloured contours indicate the region within which 90% of each subpopulation is contained. Marginal histograms are the same distributions as the top panels of Figure 7.1.

In this section, we look at the joint distributions of SN timing and position, as a function of binary evolution channel. We also compare these distributions to expectations for purely single star populations. We generate this comparison set by using identical initial conditions to our fiducial model, except we set each binary's initial orbital period to effectively infinitely wide orbits, such that they never interact. This ensures a consistent normalisation of the two populations and that differences are entirely a result of binary interactions.

We show the resulting joint distributions in Figure 7.3. In the top panel, the blue distribution shows the joint distribution for single stars. As expected, there is a correlation between time and distance, which is entirely driven by the dissolution of stellar clusters, with individual progenitor stars' orbits diverging from the centre of mass after being perturbed by a velocity consistent with the assumed fiducial velocity dispersion (Table 7.1).

Also included in the upper panel is a cross-hatched contour that contains 98% of the fiducial binary population. This distribution is both more extended and more complex than the single star distribution. Binary interactions clearly result in significantly longer tails in both time and distance for core-collapse SN feedback when comparing to a population of single stars.

We note that in addition to the change in the joint distribution, the binary evolution channels also increase the overall number of SNe. As a result, binary interactions produce *additional* feedback compared to single star evolution, with ∼11% more SNe occurring in our fiducial model than in the single star model. This change is mainly a result of merger products allowing two less massive stars to reach core collapse.

In the lower panel of Figure 7.3, we explore in further detail the evolutionary pathways and progenitor types that lead to the structure in the binary population. To first order, the distribution of primary (dark blue) and effectively single (light blue) stars are quite similar to the distribution of single stars in the the upper panel. This similarity is expected, as these progenitor types are not displaced by companions and primaries experience only small delays as a result of stable mass transfer.

More significant differences are seen for the secondaries (green) and merger products (purple). The secondaries experience significant velocity kicks when their higher mass primary companions reach core collapse. These kicks lead the stars to eventually explode at much larger distances than single stars.

The merger products in Figure 7.3 are significantly extended in both time and distance. As discussed above, the long tail in ages is a result of the slower evolution of lower mass stars, which eventually merge and reach core collapse. This also means that the stars have much longer to disperse from their parent cluster. However, because the coalescence does not significantly change the kinematics of the system, the spatial distribution of the merger products continues the trend established by cluster dissolution, as seen for the single stars.

We also note that the joint distribution offers a way to help isolate different evolutionary pathways. Supernovae found at large separations from star-forming regions, but with young ages are very likely to be secondary stars, whilst older remnants are more likely to be merger products.

4 Results II: Robustness of results to model variations

In this Section, we assess the robustness of the predicted temporal and spatial distributions to variations in the initial conditions of our binary stellar population, binary physics assumptions, and galactic orbit evolution parameters. We focus these variations on settings that have the largest potential to influence the results of our fiducial model.

As discussed in Section 3.1, the timing of SNe is altered by mass transfer and stellar mergers. Therefore, we vary both the efficiency and stability of mass transfer. Moreover, changes to the initial population, such as the initial mass function and metallicity, also strongly influence this distribution by varying the number of massive stars and the minimum mass that can reach core collapse.

The locations of SNe are driven by a combination of cluster dissolution and stars ejected from their binaries with their orbital velocity. Therefore, we vary parameters that (i) change the rate of cluster dissolution and (ii) change the orbital velocity of the secondary at the moment of the primary SN. We also consider how using a different galactic potential may change these distances. We first outline the specific variations below and summarise them in the fourth column of Table 7.1. We then discuss the impacts of these variations in Section 4.

For variations that alter our assumptions regarding binary physics, we use an identical initial population, such that the same evolution of the same binary can be compared for different settings. This approach relies on `COSMIC`'s ability to record the precise initial conditions (and random seeds) of sampled populations, making them easily reproducible with different evolution settings (Breivik et al., 2020b). For variations that alter galactic orbits, we use populations with identical initial populations and stellar evolution, such that the direction, magnitude and timing of SN kicks are consistent across variations.

4.1 Variation descriptions

Common-envelope events — A common-envelope (CE) event occurs when mass transfer becomes dynamically unstable in a runaway process. In such an event, the donor continues to expand faster and further overflow its Roche lobe as it transfers mass. Eventually the cores of both the donor star and the accretor are engulfed within the envelope material of the donor. These events can significantly shrink the binary orbit of the massive cores, and in some cases lead to a stellar merger.

Given the uncertainty of these events, we vary the efficiency of common-envelope events significantly, from to $\alpha_{\rm CE} = 0.1$ and 10.0. This efficiency changes the fraction of orbital energy that is available to unbind the envelope (Webbink, 1984; de Kool, 1990). Thus $\alpha_{\rm CE}$ affects both the fraction of binaries that survive a CE without merging, and the post-CE orbital separations of those that survive as a binary, which later impacts the ejection velocity of the secondary if a disruption occurs.

We also vary our assumptions of the critical mass ratio, $q_{\rm crit}$, necessary for stable case B mass transfer, which is the most common type of mass transfer in our simulations (in agreement with van den Heuvel, 1969). The stability of mass transfer is determined by comparing the mass ratio of the binary at the moment of Roche-lobe overflow ($q_{\rm MT} \equiv m_{\rm donor}/m_{\rm accretor}$) to the critical mass ratio, with unstable mass transfer occurring when $q_{\rm MT} >$

$q_{\rm crit}$. We calculate the two possible extremes: $q_{\rm crit,caseB} = 0$ (all case B mass transfer is unstable) and $q_{\rm crit,caseB} = \infty$ (all case B mass transfer is stable).

Stable mass transfer — The efficiency of mass transfer, β, changes the fraction of mass transferred from the primary star that is accreted by a companion. For a larger value of β, more mass is accreted by a companion, which could lead to stronger rejuvenation and changes to the time at which an accretor's SN occurs. Additionally, decreasing β allows more angular momentum to leave the system and decrease orbital separations, thereby increasing the ejection velocities of ejected secondary stars.

In our fiducial model, our default assumption is that the efficiency is based on the thermal timescale of the accretor. This choice results in an average efficiency of $\beta \approx 0.77$ in our fiducial simulations. We vary β to three fixed values: 0.0 (non-conservative), 0.5 and 1.0 (conservative).

Supernova natal kicks — Binary orbits are frequently disrupted by SN natal kicks, resulting in the ejection of secondary stars (e.g., Renzo et al., 2019). In rare cases where a binary remains bound these SN kicks can also determine how far a secondary star travels before its own core collapse. We vary the strength of SN natal kicks in three different ways. By default, we assume a double maxwellian distribution for kick velocities, with a low velocity component from electron-capture and ultra-stripped SNe, and a higher velocity component from core-collapse SNe. Two of the variations change this distribution to a single maxwellian, the first assuming all remnants are in the low kick component of the fiducial double maxwellian and the other assuming all are in the high kick component. In the third variation we no longer reduce black hole kicks based on the amount of fallback mass.

Initial distributions — We vary the assumed power law distribution of initial primary masses, orbital periods and mass ratios. The initial mass function affects the fraction of stars that reach core collapse, as well as the number of lower mass stars available for mergers and later SNe. We vary the slope of the high mass end of the IMF from our fiducial assumption of $\alpha_{\rm IMF} = -2.3$ (Kroupa, 2001) to $\alpha_{\rm IMF} = -1.9$ based on the finding for the 30 Doradus starburst (Schneider et al., 2015, 2018a) and additionally $\alpha_{\rm IMF} = -2.7$ to bracket our fiducial assumption.

The initial orbital period can change the evolutionary phase during which mass transfer occurs (or whether it occurs entirely), as well as the orbital separation of the binary at the moment of the first SN. We vary the power law slope, π, to $\pi = -1$ and $\pi = 0$, bounding the fiducial assumption of -0.55, based on the uncertainties in Sana et al. (2012) and compatible with several observational constraints (e.g., Öpik, 1924; Kobulnicky & Fryer, 2007; Moe & Di Stefano, 2017). We additionally alter our assumed upper limit for the initial orbital period, changing our fiducial value of $10^{5.5}$ days to 10^3 days. The lower maximum initial period is often assumed in earlier works, based on the limitations of previous spectroscopic surveys. These surveys could only observe binaries wider than $\sim 10^3$ days, however interferometric studies have found wider binaries (e.g., Sana et al., 2014). de Mink & Belczynski (2015) introduced a higher limit of $10^{5.5}$ days based on these results and argued that there is no physical reason for the binary fraction to sharply drop at 10^3 days. For these reasons we adopt the same limit in our fiducial model.

The initial mass ratio can change the stability of mass transfer, because more unequal mass ratio binaries are more likely to experience unstable mass transfer, as well as the time between primary and secondary SNe. We vary the power law slope, κ, to $\kappa = -1$ and $\kappa = 1$, bounding the fiducial assumption of a uniform mass ratio distribution, based on the uncertainties in Sana et al. (2012) and compatible with several observational constraints (e.g., Kobulnicky & Fryer, 2007; Mason et al., 2009; Moe & Di Stefano, 2017).

Metallicity — Stars at a lower metallicity lose less mass via stellar winds and lower radial expansion (Leitherer et al., 1992; Brunish & Truran, 1982; Baraffe & El Eid, 1991). This means that they retain more mass and have different radii during their evolution and thus may experience mass transfer during different evolutionary phases. We lower the metallicity, Z, of each star in our initial population by constant factors of $0.5, 0.2, 0.1$, and 0.05. This preserves the distribution of metallicites within the binary population, but shifts the distribution to a systematically lower mean metallicity, $\bar{Z}$.

Cluster velocity dispersion — The initial velocity dispersion of a cluster determines how quickly its member stars will disperse, changing the distance between the eventual SNe and the centre of the parent cluster. Our fiducial assumption is a velocity dispersion of $1.7\,\mathrm{km\,s^{-1}}$ (based on a virial parameter of $\alpha_{\mathrm{vir}} = 1$), which approximately follows measurements of the velocity dispersion of the Orion Nebula Cluster (Da Rio et al., 2017; Kroupa et al., 2018; Kuhn et al., 2019). We consider the impact of initial velocity dispersion of the cluster in our simulation by using a lower choice of $v_{\mathrm{disp}} = 0.5\,\mathrm{km\,s^{-1}}$ and higher choice of $v_{\mathrm{disp}} = 5\,\mathrm{km\,s^{-1}}$, which bracket the range of dispersions found in Kuhn et al. (2019).

Hydrodynamical zoom-in simulation — Our fiducial model is calculated within the context of a specific simulated galaxy. Although the majority of our results are shaped by our choice of binary physics and cluster velocity dispersion, there may be additional effects that can be traced to the specifics of the simulated galaxy. In particular, how quickly cluster stars disperse can be shaped by the gravitational potential the cluster forms and evolves within. There may also be effects due to the metallicity distribution of the population of stellar clusters, which can change the binary evolution. To assess some of these possible effects, we repeat our calculations for the `ChaNGa` `r442` galaxy, to compare with the fiducial `FIRE-2` `m11h` galaxy, with and without controlling for metallicity.

4.2 Summary plots of parameter variations

In Figures 7.4–7.7 we show comparisons of the distribution of SN times and locations across an array of parameter variations. Rather than reproducing the full distributions in Figure 7.1 for every parameter value, we have chosen to make condensed summary plots that characterize the statistical properties of the distributions. We use vertical bars to indicate the interquartile ranges and medians (horizontal marks within the bars) of supernova timing (top panel) and distance from the parent cluster centre at the time of explosion (bottom panel) for the primary (blue), secondary (green), merger products (purple), and total (grey) populations.

We note that changes in the distributions can come from a combination of (a) a shifting of the times and locations of the original population in the fiducial model or (b) the creation or elimination of subpopulations which cause shifts. Therefore, we include bars at the top of each figure that indicate the total number of SNe that occur per $100\,\mathrm{M_\odot}$ of star formation.

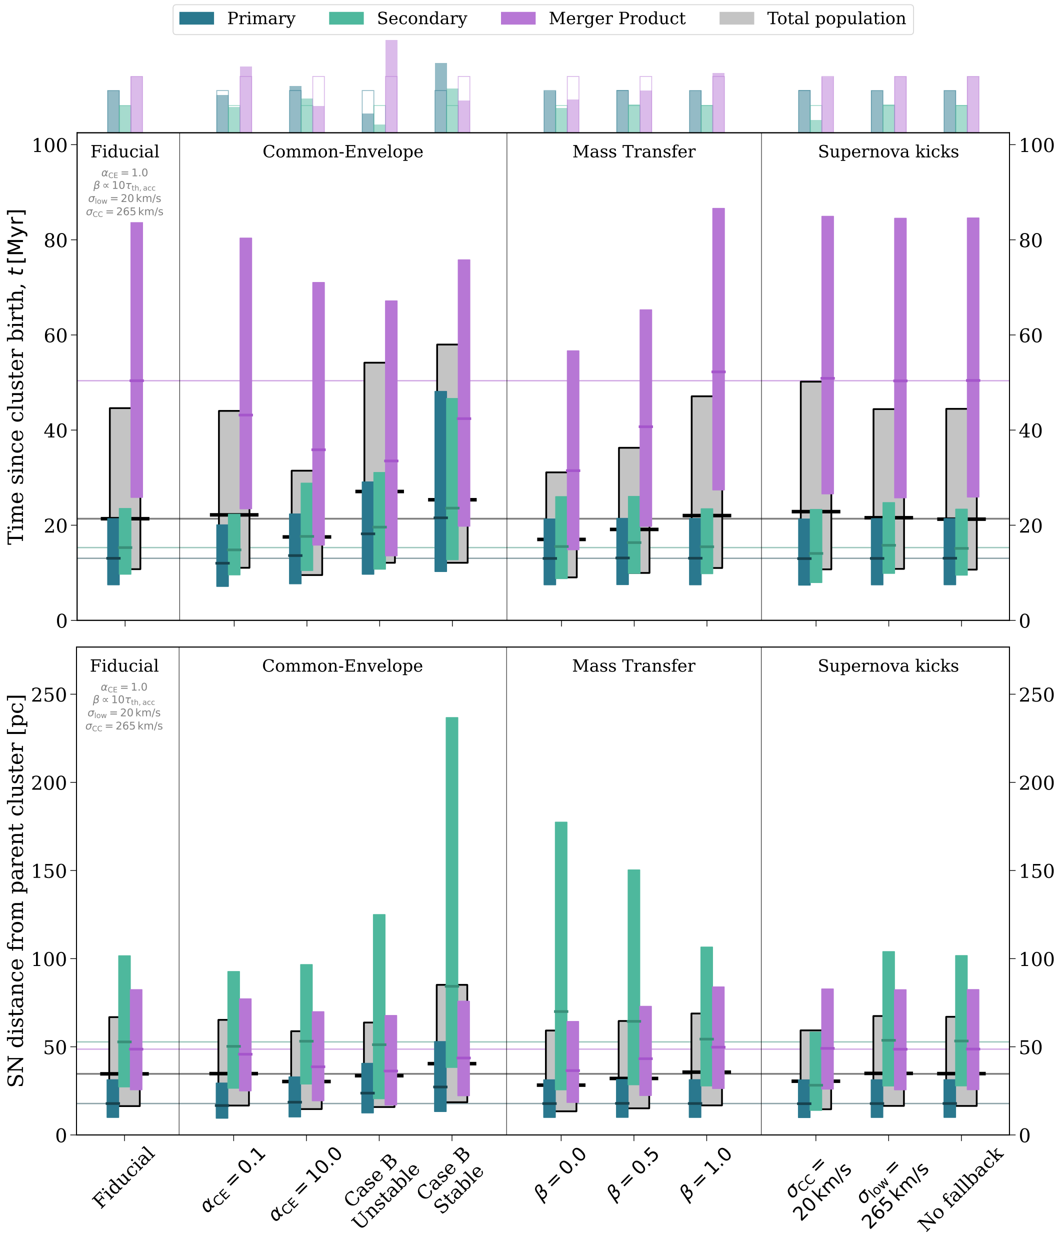

Figure 7.4: Comparison of the impact of binary physics variations on the timing (upper panel) and location (lower panel) of SNe. Each group of bars in the main panels corresponds to a different choice of binary physics. Coloured bars show the interquartile range for each subpopulation (labelled in legend) with darker lines at the median. The grey bars show the same information for the overall distribution. Full distributions for the fiducial model are shown in Figure 7.1. Thin horizontal lines indicate the fiducial median values for comparison. Fiducial parameter are annotated in the upper left corner. The small bars at the top indicate the number of SNe per 100 $\rm M_\odot$, with the fiducial values shown in a thin outline. Section 4.2 explains how to interpret this figure (and Figures 7.5–7.7) in further detail. (💻 Interactive figure available.)

These can be compared to the thin outlined bars, which indicate the fiducial values.

In every plot, the fiducial model is reproduced in the leftmost column for reference, with the relevant fiducial settings annotated at the top of the column. Its median values are also propagated throughout the plot as thin horizontal lines. The parameters that are being varied are shown along the bottom of the plot, and are grouped by the physics that the parameters control. These parameter variations may have significant impacts on the distributions of particular evolutionary pathways (i.e., indicated by moving the coloured bars up or down, and/or changing their relative widths), which can be a useful diagnostic of the controlling physics.

In addition to these plots, throughout the following sections we will use a series of summary statistics to show how variations change the tails of our distributions. These are summarised in Table 7.2 and displayed in Figure 7.13, where $f_{t>44\,\mathrm{Myr}}$ is the fraction of SNe that occur "late" (beyond 44 Myr, the `FIRE`-3 cutoff for type II SN feedback), and $f_{D>100\,\mathrm{pc}}$ and $f_{D>500\,\mathrm{pc}}$ are the fractions of SNe that occur more than 100 pc and 500 pc from their parent clusters respectively.

In Appendix 9.3, we include tables of data for reproducing the total population distributions from the comparison plots. Table 7.4 & 7.5 details the percentiles of both the timing and distance distributions, whilst Table 7.3 contains the total number of SNe in each subpopulation in each variation.

4.3 Binary physics variation trends

In Figure 7.4, we focus on the impact of varying parameters that directly control binary evolution – common-envelope evolution, the efficiency of mass transfer, and the strength of SNe kicks. As described in Section 4.1, we vary these parameters over a wide range of possible values, many of which are quite extreme.

4.3.1 Common-envelope efficiency

Approximately 35% of stars that reach core collapse experienced a common-envelope event during their evolution in our fiducial model. Therefore changing the efficiency of this phase can impact a large fraction of our population.

Figure 7.4 shows that decreasing the efficiency of common-envelopes (i.e., lower values of α_{CE}), results in more stellar mergers. A lower value of α_{CE} reduces the fraction of orbital energy that is available to unbind the envelope and so stars need to get much closer before releasing enough energy to unbind the common envelope, making a merger more likely. This change is also notable in the bar charts on top of Figure 7.4, which show that the number of mergers increase by 18% for $\alpha_{\mathrm{CE}} = 0.1$, such that their overall contribution to the total SNe population is 50%. A subset of these additional mergers come from stars that would have reached core collapse alone, and hence the number of primary and secondary SNe each decrease slightly by ∼5%.

The increase in stellar mergers at lower α_{CE} reduces the overall median of the SN time distribution for merger products. The median decreases because more higher mass stars, which would have exploded as primary or secondary SNe in the fiducial model, instead merge and evolve to core collapse faster than an average merger product. Nevertheless, the overall distribution of times is relatively unaffected, because merger products represent a larger fraction of the population. Therefore, though the average merger product SNe is earlier, the average SNe of *any* progenitor is relatively unchanged. Similarly, the total distance distribution is not significantly affected, though merger distances slightly decrease as a result of decreased times to core collapse giving them less time to disperse.

Following the same trend as lower $\alpha_{\rm CE}$, the variation with a higher value of $\alpha_{\rm CE}$ decreases the number of merger product SNe by 53%, such that merger products produce only 24% of all SNe in the $\alpha_{\rm CE} = 10.0$ variation. However, counterintuitively this variation *also* reduces the average merger product SNe time even more significantly than the $\alpha_{\rm CE} = 0.1$ variation. In this case, the increased efficiency allows more stars to avoid a merger during a common-envelope. However, unlike in the $\alpha_{\rm CE} = 0.1$, these stars don't shift to a different progenitor channel, but instead often avoid core collapse entirely. In this way a large fraction of the subpopulation that occupy the late SN tail is removed, shifting the distribution to smaller times, such that $f_{t > 44\,{\rm Myr}} = 14\%$ in this case. Additionally, the distance distribution decreases because there are fewer late SNe that have longer to disperse from their cluster.

4.3.2 Case B mass transfer stability

In our fiducial model, ∼85% of stars that reach core collapse have are involved in Roche-lobe overflow, either as a donor or an accretor. Of these stars, ∼66% experience case B mass transfer (i.e., the subset of Roche-lobe overflow events that are initiated during expansion on the Hertzsprung gap). Given that "case B" is the most common type of mass transfer, we explore how changing our criteria for its stability affects our results. These variations are shown in the common-envelope section of Figure 7.4 for the two extremes of case B behaviour.

At one extreme, when all case B mass transfer is unstable, ∼75% SN progenitors undergo a merger prior to core collapse. This increase in mergers is a result of failed common-envelope events in which the binary fails to eject its envelope and instead merges. Furthermore, the same process leads a greater fraction of merger product SNe to come from higher mass stars with shorter evolutionary timescales, which in turn decreases the typical SN times for merger products, compared to the fiducial model (in which these stars would have avoided merging).

While the evolution of merger products is accelerated with unstable case B mass transfer, the *overall* median SN delay time of the population as a whole has the opposite behavior, and significantly increases to longer times. Mergers become more common, making the tail of late-time ($> 44\,{\rm Myr}$) SNe that they produce even more prominent (see Figure 7.1). Thus, despite the lower average time of the merger product subpopulation, the median time of the *total* population increases by ∼30%, such that $f_{t > 44\,{\rm Myr}} = 31\%$, due to merged SNe progenitors representing a much greater fraction of the population. The secondary distance distribution also appears to show a slightly extended tail, but we caution that this distribution is very weakly populated (as shown in the bars at the top of the figure) and thus not well characterized or important to the overall behaviour of the population.

In the other extreme, we consider models that force all case B mass transfer to be stable. This change *also* increases the average core-collapse SN delay time, but for different reasons. In this variation, the number of merger product SNe decreases significantly because many binaries avoid common-envelopes entirely. These binaries instead reach core collapse as primary or secondary star progenitors. As a result, many low-mass stars are able to stably accrete enough matter from their companions to reach core collapse without requiring a merger. This means that the times of primary and secondary SNe are significantly later in this variation, shifting the overall SNe distribution to later times.

Another feature of the "always stable" case B variation is that it shows the strongest increase in SN distances among all the binary physics variations that we considered, particularly for the secondary SNe. The additional population of distant secondary SNe are a result of low-mass stars that could not have reached core collapse without accretion. The stable mass transfer variation allows these stars to avoid mergers and thus remain in binaries, enabling them to later be ejected at high velocities when the primary star explodes. Furthermore, these no-longer-merging stars are primarily lower mass, which means that they travel further before exploding and hence increase the tail of the distribution. Overall, this increases the fractions of distant SNe to $f_{D>100\,\mathrm{pc}} = 21.1\%$ and $f_{D>500\,\mathrm{pc}} = 2.9\%$, which are 1.6x and 3x the fiducial values respectively.

4.3.3 Mass transfer efficiency

The efficiency of mass transfer, $\beta \equiv \Delta M_{\rm acc}/\Delta M_{\rm donor}$, is defined as the fraction of material transferred by the donor ($\Delta M_{\rm donor}$) which is successfully accreted by the companion star ($\Delta M_{\rm acc}$), rather than being removed by winds. In our fiducial model, the efficiency is not set directly, but instead dictated by the thermal timescale of the accretor, such that the maximum accretion rate is given by

$$\dot{M}_{\rm acc,max} = \begin{cases} 10\frac{M_{\rm acc}}{\tau_{\rm KH}} & \text{Has radiative envelope,} \\ \infty & \text{else,} \end{cases} \tag{7.1}$$

where $M_{\rm acc}$ is the mass of the accretor, $\tau_{\rm KH}$ is the Kelvin-Helmholtz timescale of the accretor star and this limit is only applied for stars with radiative envelopes (Hurley et al., 2002). We calculate the fraction of mass that's accreted in each instance of stable mass transfer in our fiducial model and find that this criterion results in an average of $\beta \approx 0.77$. We show three variations using fixed values for β in Figure 7.4.

The distributions of primary stars are almost entirely unaffected by these variations. Changes in β only affect the amount of mass accreted, not the mass lost. Therefore, primary stars would only be affected by these variations in a rare case in which the accretion went the opposite way during a late phase of a binary's evolution, such that a primary star accreted a significant amount mass from an evolving companion before the primary's core mass is determined.

We do see overall changes in the timing of SN from variations in mass transfer efficiency. There is a clear trend that a lower value of β results in a weaker tail of SN at late times (see Figure 7.13), primarily because there are fewer late merger-product SNe. In our fiducial model, these late SNe most often follow a formation pathway in which a merger occurs after unstable mass transfer from a secondary star onto a primary star (e.g., Zapartas et al., 2017). A lower mass transfer efficiency results in a less massive secondary star and therefore fewer scenarios in which a secondary star expands sufficiently during its evolution to initiate a common-envelope at close separations. As a result, the fraction of merger product SNe from this pathway decreases from 50%, to 35% to 17% for $\beta = 1.0$ to 0.5 to 0.0, thus reducing the late SN rate.

The distance distributions show two main trends: an overall decrease in the median distance of SNe, but an increase in the distances that secondary stars travel. The overall decrease follows from the same reasoning as above, fewer late merger product SNe mean a larger fraction of the SNe happen close to their parent cluster. The increase in secondary distances for low β occurs for two reasons. First, lower β values mean secondary stars accrete less material, potentially delaying their SN explosions and thus giving them more time to travel after their primary explodes. Second, more angular momentum is lost from the system for lower β values, and thus binary separations tend to be tighter prior to their disruption, leading them to be ejected with higher velocities. Quantitatively, we find that the average ejection velocity for secondary stars changes from $14.6\,\mathrm{km\,s^{-1}}$ to $16.5\,\mathrm{km\,s^{-1}}$ to $22.1\,\mathrm{km\,s^{-1}}$ as β decreases from 1.0 to 0.5 to 0.0. At the same time, secondary stars on average live for an additional 1.5 Myr after ejection from a binary when $\beta = 0.0$ compared to $\beta = 1.0$. For these reasons, secondary stars tend to reach core collapse further from their parent clusters when β is lower.

4.3.4 Supernova natal kick magnitude

The magnitude of the SN natal kicks primarily alter our distributions by changing the fraction of binaries that are disrupted. Even in our fiducial model ∼85% of binaries are disrupted by their first SN, in agreement with earlier works (De Donder & Vanbeveren, 2003; Eldridge et al., 2011; Renzo et al., 2019). We emphasise that, for natal kicks that already disrupt a binary, increasing the magnitude of the kick has almost no effect on the ejection velocity of secondary stars (e.g., Wagg et al., 2025). The variations that increase the strength of the kicks only disrupt an additional 5% of binaries, hence they have negligible effects on our distributions, as shown in the supernova kicks section of Figure 7.4.

While the changes in the overall distributions do not change dramatically with variations in the SNe kick magnitudes, there are some subtle variations that are worth explaining. Most of these effects will be driven by the population of secondary stars; primary stars and merger products are unaffected as both cases represent the first SN in a binary system and hence their prior evolution has not been affected by SN kicks.

The first place one might expect to see the effects of SNe kicks are in the distance distributions, which are indeed skewed to lower values when using weaker core collapse SN kicks. However, the origin of this effect is not due to directly reducing the ejection velocities of secondary stars, as these velocities are entirely set by the pre-supernova orbital velocity (Wagg et al., 2025). Instead, the reduction in the typical cluster distance is due to having fewer ejected secondary stars. The number of secondary stars that reach core collapse in this variation decreases by a factor of 2 compared to the fiducial, such that the total number of SNe decreases by ∼12%. The decrease in secondary SNe is because the weaker primary SN kicks disrupt only ∼35% of binaries, compared to 85% in the fiducial model. These weaker kicks cannot unbind tight binaries, such that the only ejected secondaries come from the subset of weakly-bound binaries that have wide orbits with lower orbital velocities.

In contrast, SN delay times are slightly later in the variation in which core-collapse SN natal kicks are weaker. This again is the result of a much smaller fraction of binaries being disrupted. The surviving binaries can then often experience a subsequent unstable mass transfer event, which leads to a merger between the compact object and secondary star and prevents a subsequent core-collapse SN. The decrease in the number of secondary SNe means that the SN times are generally later because merger products represent a greater fraction of the population.

4.3.5 Summary of binary physics trends

The story told by Figure 7.4 is that the expected distributions of SNe time delays and distances from cluster centres are surprisingly robust to even extreme variations in binary physics. Focusing only on the medians of the distributions (dark horizontal lines across the grey bars), there is not much change across the entire plot, for either the timing or distance distributions; these vary by no more than 25% and 18%, respectively.

The effects of the variations are more visible in the tails (Figure 7.13), and in the number and type of SNe precursors. However, in all cases there is a long tail of late SNe from merger products, which is not currently accounted for in hydrodynamical simulations, and a long tail of distant SNe from secondary stars, which is particularly sensitive to changes to the stability of mass transfer.

4.4 Initial condition variation trends

The initial conditions of binary stellar populations, such as the initial primary mass, mass ratio and orbital period distributions, are not strongly constrained and may differ from our fiducial assumptions. In this section, we test the sensitivity of our results to variations in these distributions and display the results in Figure 7.5.

4.4.1 Initial mass function

We consider two variations of the slope, $\alpha_{\rm IMF}$, of the high mass end of the IMF (for primary masses of $m_1 > 1\,{\rm M}_\odot$). The $\alpha_{\rm IMF} = -1.9$ variation is a top-heavy IMF that favours more massive stars, while $\alpha_{\rm IMF} = -2.7$ is contrastingly a bottom-heavy IMF; the optimal value remains a major observational and theoretical uncertainty in any calculation of SNe rates. We keep the low mass ($m_1 < 1\,{\rm M}_\odot$) IMF the same as the Kroupa IMF, while ensuring continuity at m_1. The impact of these variations are shown in Figure 7.5.

As expected for any model (binary or otherwise), the overall total number of SNe is strongly affected by the IMF (de Mink & Belczynski, 2015; Zapartas et al., 2017). Varying the IMF slope has a first-order effect simply through changing the fraction of stars that have high enough masses to be candidates for eventual core collapse. Here, the top-heavy IMF variation produces significantly more massive stars, resulting in twice as many primary and secondary SNe in our simulations.

There is a more nuanced effect on the number of merger product SNe, which increases by a lesser extent, of only 40% rather than a factor of 2. Many of these merger product SNe are formed from binaries in which the primary star is not massive enough to reach core collapse alone, but, these binaries constitute a lower fraction of the population in a top-heavy IMF. Quantitatively, roughly half (52%) of the stars in our fiducial Kroupa IMF simulation have

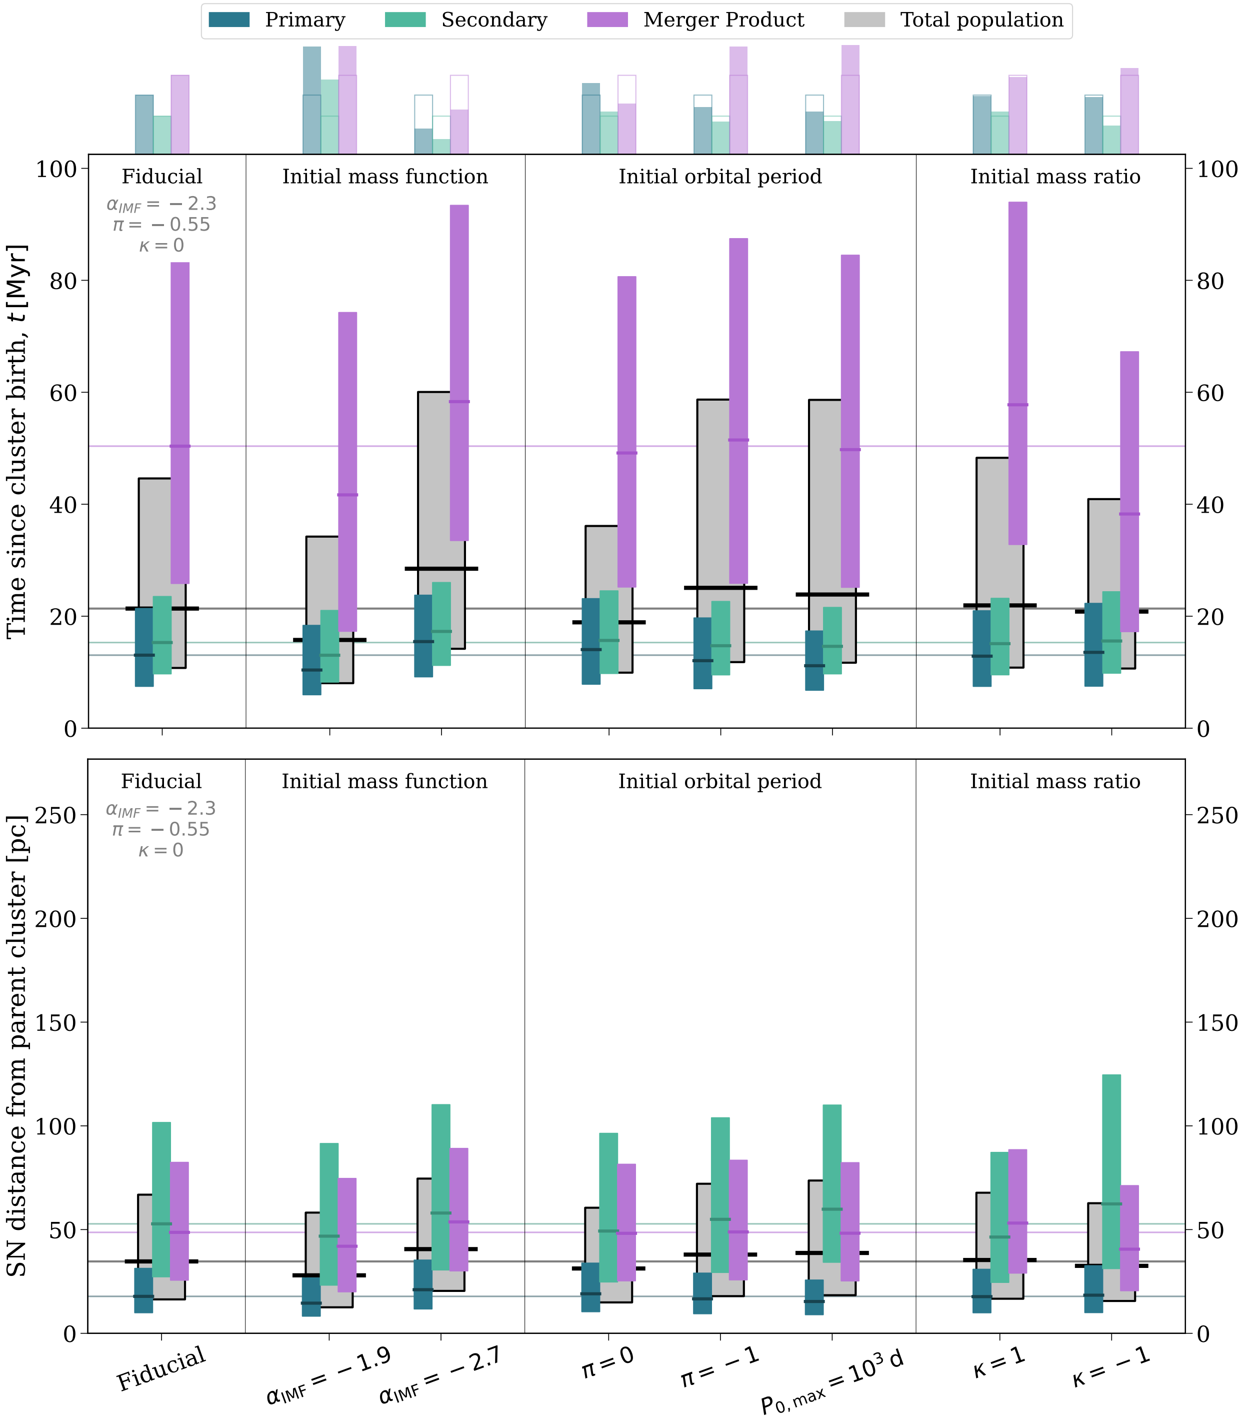

Figure 7.5: As Figure 7.4, except varying different choices of initial conditions. Variations of the initial mass function and initial orbital period distribution can affect SNe timing, but initial conditions do not strongly impact SN spatial locations. (💻 Interactive figure available.)

$m_1 < 7\,\mathrm{M}_\odot$, whilst for the top-heavy IMF the same fraction drops to only 35% (noting that we only simulate primary masses of $m_1 > 4\,\mathrm{M}_\odot$; see Section 2.2). This shift to higher masses reduces the fraction of the population that requires a merger in order to reach core collapse.

All of the trends above are reversed for the bottom-heavy IMF, with overall decreases in the total SNe rate that are more drastic for primary and secondary SNe (a 55% decrease) and less so for merger products (a 40% decrease).

Although a more top-heavy IMF produces a greater *number* of SNe, it decreases the fraction of these that occur at late times. The SN times of each progenitor type decrease for the top-heavy $\alpha_{\rm IMF} = -1.9$ case, because more stars form at high masses and reach core collapse in a shorter time. As such, only $f_{t > 44\,\rm Myr} = 17.9\%$ of SNe occur after 44 Myr in the case of a top-heavy IMF. However, though this variation has a lower *fraction* of the late SNe, there are still ∼17% more late SNe than the fiducial model because so many more stars reach core collapse. In contrast, the bottom-heavy IMF variation has the one of the largest fractions of late SNe of all variations we consider ($f_{t > 44\,\rm Myr} = 32.9\%$), but with a total absolute rate that is only 64% of the fiducial model.

These changes to the timing distribution lead to slight variations in the distances that stars travel before reaching core collapse. A top-heavy IMF results in a slightly more concentrated distribution of SNe, whilst a bottom-heavy IMF has a slight increase to all distances.

4.4.2 Initial orbital period

We explore the impact of changing our distribution of initial orbital periods in two variations from our fiducial model, which uses a power law slope of $\pi = -0.55$ (see Section 2.2). In the $\pi = 0$ simulation, binaries are generally placed at wider orbital periods, whilst in $\pi = -1$ they are concentrated on closer initial orbits. These wide- and tight-orbit variations are shown in the third section of Figure 7.5.

The strongest impact of these variations is to change the fraction of SNe from merger products. Binaries with tighter initial orbital periods are more likely to merge during their evolution. For the wide-binary case ($\pi = 0$) there are ∼54% fewer merger product SNe than the fiducial model with $\pi = -0.55$, whilst for $\pi = -1$ the total increases by ∼38%.

The merger fraction is the dominant driver of late-time SNe, and thus the increase in mergers for the tight-binary $\pi = -1$ case leads directly to a more prominent tail to late times. In fact, the $\pi = -1$ variation in particular produces the second-highest[3] fraction of late SNe of all of our variations, with $f_{t > 44\,\rm Myr} = 32.7\%$, in agreement with rates found by Zapartas et al. (2017)). In contrast, the overall distribution of SN times for the wide-binary $\pi = 0$ case decreases the late-time tail significantly.

Besides the change in merger frequencies, initially wider binaries result in slower ejection velocities and thus less distant SNe. In the wide-binary ($\pi = 0$) variation, orbital speeds are typically slower, reducing the average ejection velocity of secondary stars when the primary star explodes. The lower ejection velocities result in slightly fewer distant SNe in this variation ($f_{D > 100\,\rm pc} = 11.6\%$) compared to our fiducial model ($f_{D > 100\,\rm pc} = 13.4\%$) and the tight-binary $\pi = -1$ variation ($f_{D > 100\,\rm pc} = 14.7\%$).

---

[3] We note that the overall *rate* of late SNe is highest in this variation, because the bottom-heavy IMF suppresses SNe.

We additionally include a variation in which we reduce the upper limit on initial orbital periods to $P_{0,\mathrm{max}} = 10^3$ days. This lower limit suppresses the number of wide, effectively-single binaries. Although there is no reason for there to be a strict upper limit at this period (e.g., de Mink & Belczynski, 2015), this stricter limit is more typical of early work on binary populations. Similar to the $\pi = -1$ variation, this skews the distribution of initial binaries to tighter orbits. For the same reasons as $\pi = -1$, this produces a high fraction of late SNe ($f_{t>44\,\mathrm{Myr}} = 33.3\%$) and moderately high fraction of distant SNe ($f_{D>100\,\mathrm{pc}} = 15.3\%$).

4.4.3 Initial mass ratio

We change the power-law slope, κ, of our initial mass-ratio distribution in two variations. The $\kappa = -1$ variation produces more unequal-mass binaries, whilst $\kappa = 1$ favours more equal-mass binaries. These variations are shown in the final section of Figure 7.5.

Merger product SNe generally occur earlier in the variation with more unequal mass ratios ($\kappa = -1$). Generally, this timescale is set by the lifetime of the (more massive) primary star, for the subset of "forward mergers" — those occurring when the primary star overflows its Roche lobe). More unequal mass ratios tend to produce more forward mergers (73% of mergers for $\kappa = -1$, compared to 52% in our fiducial model), because the larger mass difference reduces the chance that both the primary and secondary will be evolving significantly at the same time. As a result, the times of merger product SNe are generally earlier in this variation. The inverse reasoning applies to the $\kappa = 1$ variation and explains its later merger times.

The earlier merger SNe times for more unequal mass ratios ($\kappa = -1$) also shifts the overall distance distribution to lower values, because those SNe have less time to disperse from the cluster. Since merger products represent the largest progenitor population (see histograms in the top row of Figure 7.5), the overall distance distribution is skewed to lower values.

While the changes in merger-progenitor SNe times shifts the overall distribution to shorter times and smaller separations when the initial mass-ratio is more uneven, there is an additional (albeit smaller) change driven by changes to the secondaries, which have larger distance separations on average. The increased distances stem from secondary stars in binaries with more unequal mass ratios taking longer to evolve than their primary star. This delay means that there is more time between the disruption of a binary (after the primary SN) and the eventual secondary SN, allowing the ejected secondary companion star to travel further before exploding. In particular, we find that secondary stars in disrupted binaries travel for 2.5 Myr longer before exploding in the $\kappa = -1$ variation than in the $\kappa = 1$ variation. This means that though $f_{D>100\,\mathrm{pc}}$ slightly decreases to 11.9% (from 13.4%) in the $\kappa = -1$ variation, $f_{D>500\,\mathrm{pc}}$ increases to 1.2% (from 0.9%).

4.5 Metallicity

In our fiducial model, there is an intrinsic distribution of stellar metallicities, set by the specific history of star formation and gas flows within a single model galaxy. However, metallicity is a well-known driver of stellar evolution, changing the internal structure, radius, and time-evolution of single stars, all of which will have downstream effects when those stars are evolving within binaries. We therefore now consider the impact of metallicity variations.

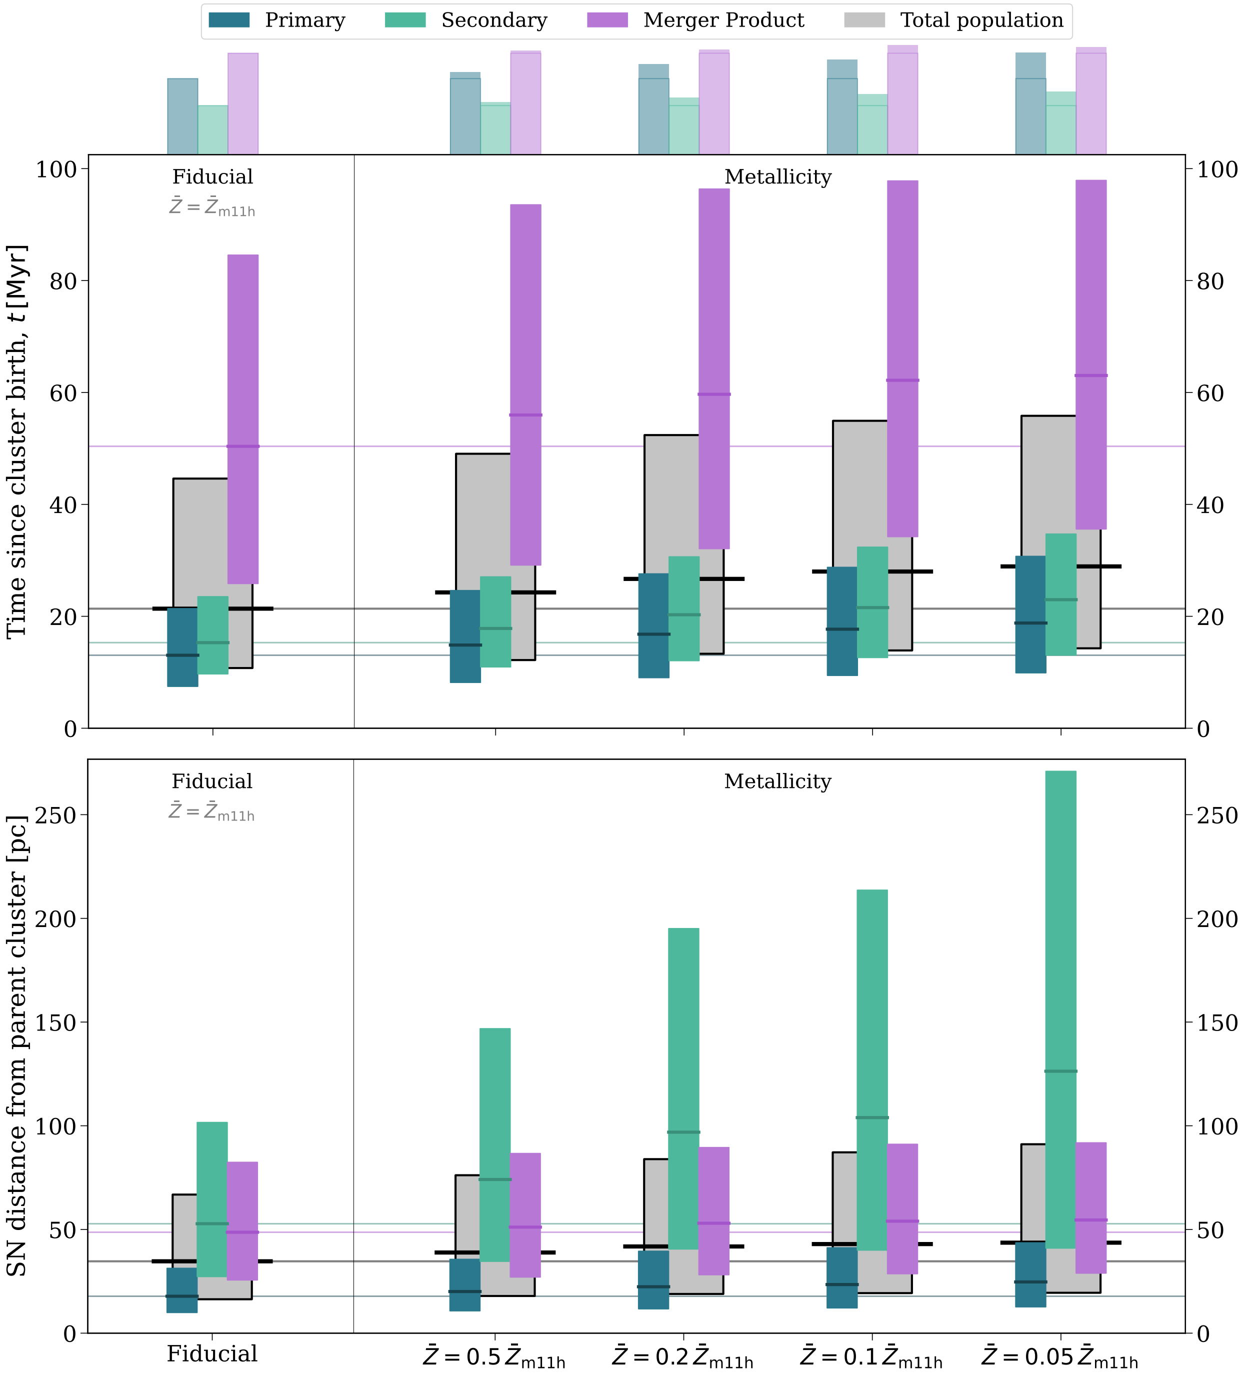

Figure 7.6: As Figure 7.4, except varying average population metallicity. Decreasing the metallicity of simulated binary stars leads to a systematic increase in both SNe times and locations. (💻 Interactive figure available.)

Our fiducial simulation has a relatively narrow, near-solar metallicity distribution ($\bar{Z}_{\rm m11h} = 0.017 \pm 0.001$, where the uncertainty indicates the interquartile range). We explore the impact of metallicity by decreasing the metallicity of each binary in the fiducial model simulation by a constant factor in a series of variations. This process leaves the overall shape of the metallicity distribution unchanged, while decreasing the average metallicity, $\bar{Z}$. We not explore higher metallicity variations as the highest metallicity that can be evolved in `COSMIC` is $Z = 0.03$ (Breivik et al., 2020b) and `m11h` already has star particles with $Z = 0.0223$, so any significant increase would require changing the overall metallicity distribution to avoid exceeding $Z = 0.03$. The resulting SNe distributions for these metallicity variations are shown in Figure 7.6, for a variety of increasingly sub-solar metallicities.

As metallicity decreases, the amount of feedback increases. In the $Z = 0.1 Z_{\rm m11h}$ variation, the total number of SNe increases by ∼17% compared to the fiducial near-solar model (as shown by the top histograms in Figure 7.6). At lower metallicities, opacity and the strength of wind mass loss is decreased. As a result, the lower limit for the initial mass of a single star that can reach core collapse is decreased (Pols et al., 1998). We find that the lowest mass star that accreted no mass from its companion and still reached core collapse is $\sim 7.0\,{\rm M}_\odot$ in our fiducial model, but $\sim 6.2\,{\rm M}_\odot$ in the $Z = 0.1 Z_{\rm m11h}$ variation. The ∼17% increase in SNe is therefore exactly as expected from integrating the initial mass function from this decreased lower mass limit.

The addition of this subpopulation of lower mass stars increases the times at which all subtypes of SNe occur. The main sequence lifetimes of these stars are longer than the rest of the population and thus their addition skews the average SN time to later values. For the lowest metallicity variation $f_{t > 44\,{\rm Myr}} = 34.4\%$, producing the largest fraction of late SNe across all variations we consider.

The addition of more late SNe with decreasing metallicity also increases the overall SN distance distribution for each subtype. While some fraction of this effect is simply allowing more time for cluster dissolution, a particularly dramatic change is in the behaviour of SNe from secondary stars.

Secondary stars have strongly metallicity-dependent evolution, due to changes in the radial expansion of stars, which is decreased at low metallicities (e.g., Brunish & Truran, 1982; Xin et al., 2022; Klencki et al., 2022). In our simulations, the median of primary stars' maximum radial extent decreases from $\sim 1200\,{\rm R}_\odot$ in our fiducial model, to $\sim 350\,{\rm R}_\odot$ in the lowest metallicity variation. The smaller stellar radius allows binaries to detach after Roche-lobe overflow at much smaller separations. As a result, binaries are typically tighter prior to the primary SN and thus the ejection velocities of secondary stars are increased, in agreement with Renzo et al. (2019). In our models the average ejection velocity of secondaries from disrupted binaries increases from $15\,{\rm km\,s^{-1}}$ in the fiducial model to $27\,{\rm km\,s^{-1}}$ at the lowest metallicity. As a result, the very low metallicity $\bar{Z} = 0.05 \bar{Z}_{\rm m11h}$ variation has one of the strongest distance SNe tails, with $f_{D > 100\,{\rm pc}} = 22.3\%$ and $f_{D > 500\,{\rm pc}} = 2.4\%$ (see Figure 7.13).

While the metallicity clearly has a significant effect on the tails of the distance and age distributions, the bulk of the distribution experiences actually much more modest changes. The median age and distance change by only 34% and 25%, respectively, at the most extreme variations. The overall numbers of SNe also only changes by no more than 20% (see Figure 7.14).

4.6 Galaxy parameter variation trends

There are two aspects of the work above that rest on assumptions we have made in connecting the `cogsworth` infrastructure to a simulated galaxy. The first is how individual star particles in the simulation are turned into a cluster of binary stars. The chosen statistics of the binary populations are discussed above, but the assumed velocity dispersion of the stellar population is also a factor in setting the eventual orbital diffusion of the stars. The second aspect to be considered is the choice of simulated galaxy. Different evolutionary pathways can lead to different metallicities, whose impacts are discussed in Section 4.5. However, there are other potential impacts on the orbital evolution of young stars due to varying the gravitational potential and the spatial distribution of young stars.

In this section we consider the impacts of both of these effects, by varying the assumed initial velocity dispersion of the stellar clusters that replace star particles, and by changing the galaxy simulation used for our calculations.

4.6.1 Cluster initial velocity dispersion

Observational evidence suggests that nearly all young stellar clusters become unbound after expulsion of their natal gas cloud (e.g., Hills, 1980; Krumholz et al., 2019; Hennebelle & Grudić, 2024, and references therein). Thus, even in the absence of binary evolution, the stars in young stellar clusters will naturally drift apart due to their initial velocity dispersion. The amplitude of this velocity dispersion will directly impact the average separation of member stars from the cluster's guiding centre, at the time the stars undergo core collapse.

There are currently no strong observational constraints on the appropriate value of the initial cluster velocity dispersion. We have made a "typical" choice for our fiducial model ($v_{\rm disp} = 1.7\,{\rm km\,s^{-1}}$), but given the uncertainty, we also vary the initial cluster velocity dispersion to fixed values of 0.5 and $5\,{\rm km\,s^{-1}}$. We plot the impact of these variations in Figure 7.7, but note that the timing distributions (top panel) are unchanged, because the velocity dispersion does not change any aspect of stellar evolution.

In contrast to the timing distributions, the initial velocity dispersion has a strong impact on the distances at which SNe occur (bottom panel). Indeed, these variations produce the strongest change in the median of the overall distance distribution among all of our variations.

While the means do change, the tail of distant SNe from secondary stars is relatively unchanged at lower velocity dispersions. For these secondaries, the distances are mostly driven by binary physics rather than initial velocities. Therefore even for a very low choice of the initial velocity dispersion, some SNe still occur far from their parent cluster in low density environments. For $v_{\rm disp} = 0.5\,{\rm km\,s^{-1}}$, the fraction of SNe beyond 100 pc decreases to $f_{D>100\,{\rm pc}} = 5.2\%$, but the most distant tail is almost unchanged from our fiducial model with $f_{D>500\,{\rm pc}} = 0.8\%$.

4.6.2 Hydrodynamical simulation codes

So far, all of the calculations presented here have assumed that a binary population is evolved from the recent star formation in a specific simulated galaxy (`m11h` from `FIRE`). This choice can affect the resulting distributions by changing the gravitational potentials and stellar kinematics of the young stars, as well as their metallicities (see Section 4.5 above).

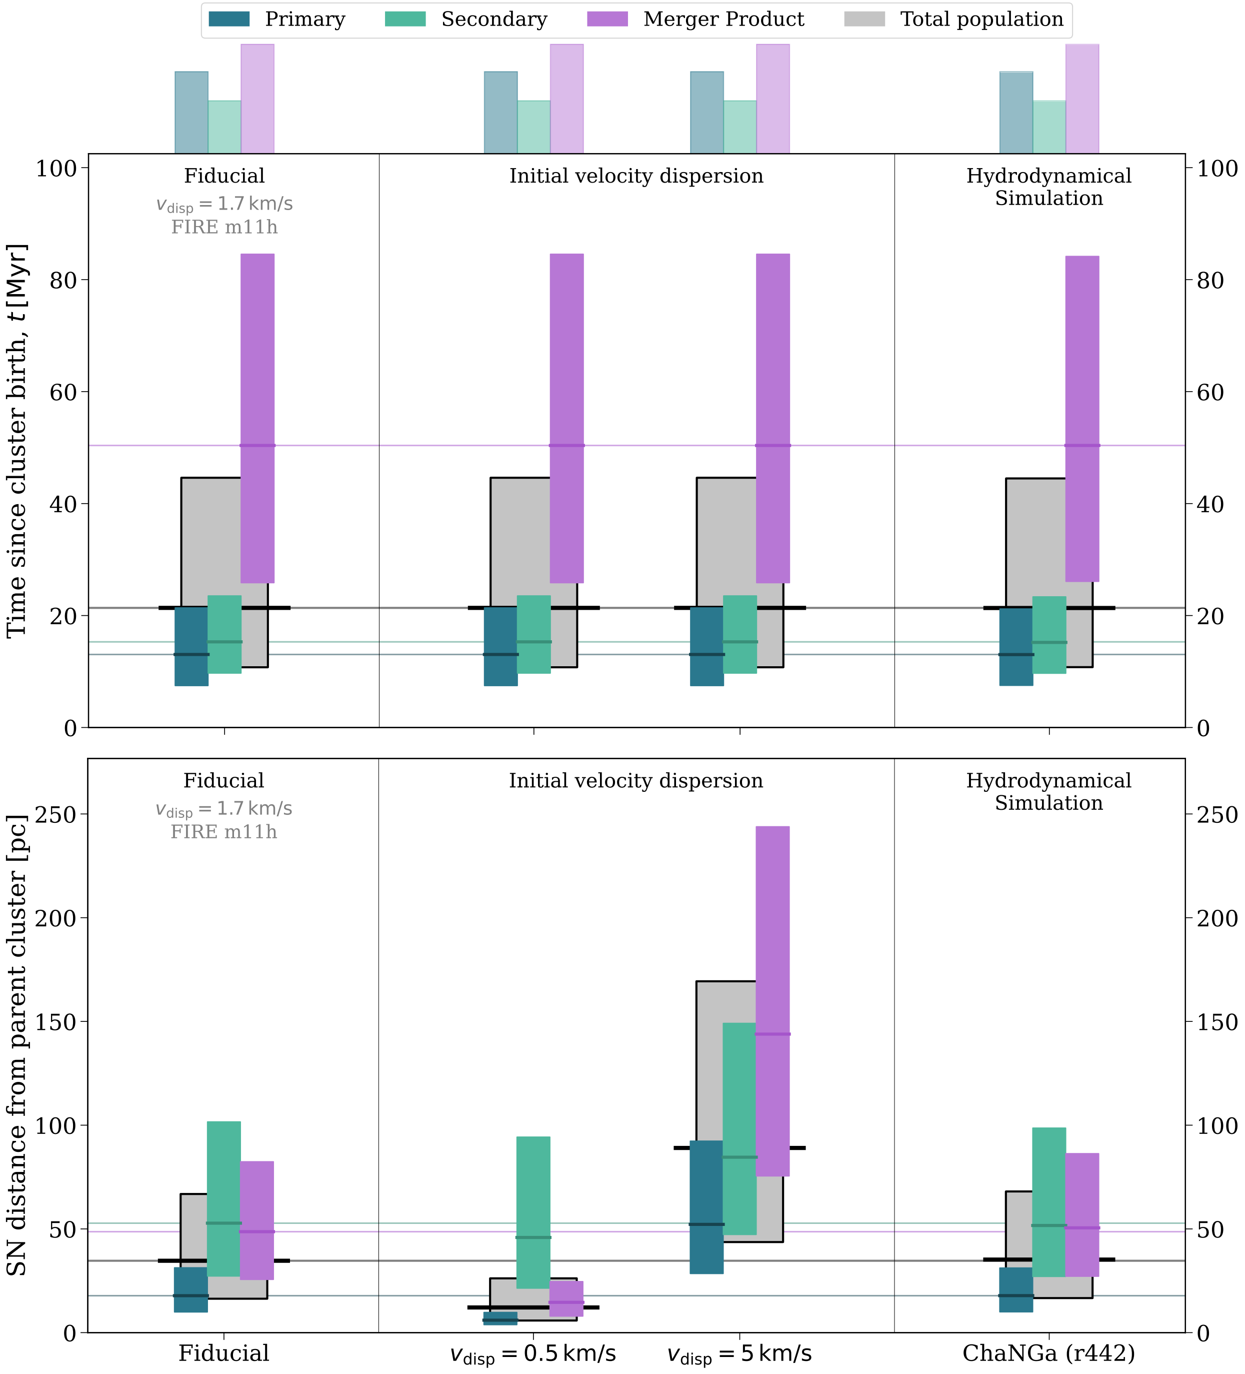

Figure 7.7: As Figure 7.4, except varying different choices of galaxy settings. SNe timing is unaffected by galaxy settings, but SNe location is strongly dependent on the choice of initial cluster velocity dispersion (though a tail of distant ejected secondary SNe still exists with a low dispersion).
(💻 Interactive figure available.)

We explore the possible impact of the choice of simulation here, by considering instead the galaxy `r442` from the Massive Dwarfs suite of simulations run with `ChaNGa` (Menon et al., 2015). `r442` is a similar (Keith et al., 2025; Ruan et al., 2025) but moderately larger galaxy, with a total mass approximately 1.5 times higher than `m11h`. Its recent star formation rate is similar, such that the total stellar mass formed in the most recent 150 Myr is only ∼6% higher than in the fiducial simulation `m11h`. Additionally, the disk of `r442` is approximately 20% less massive which is important for the galactic potential.

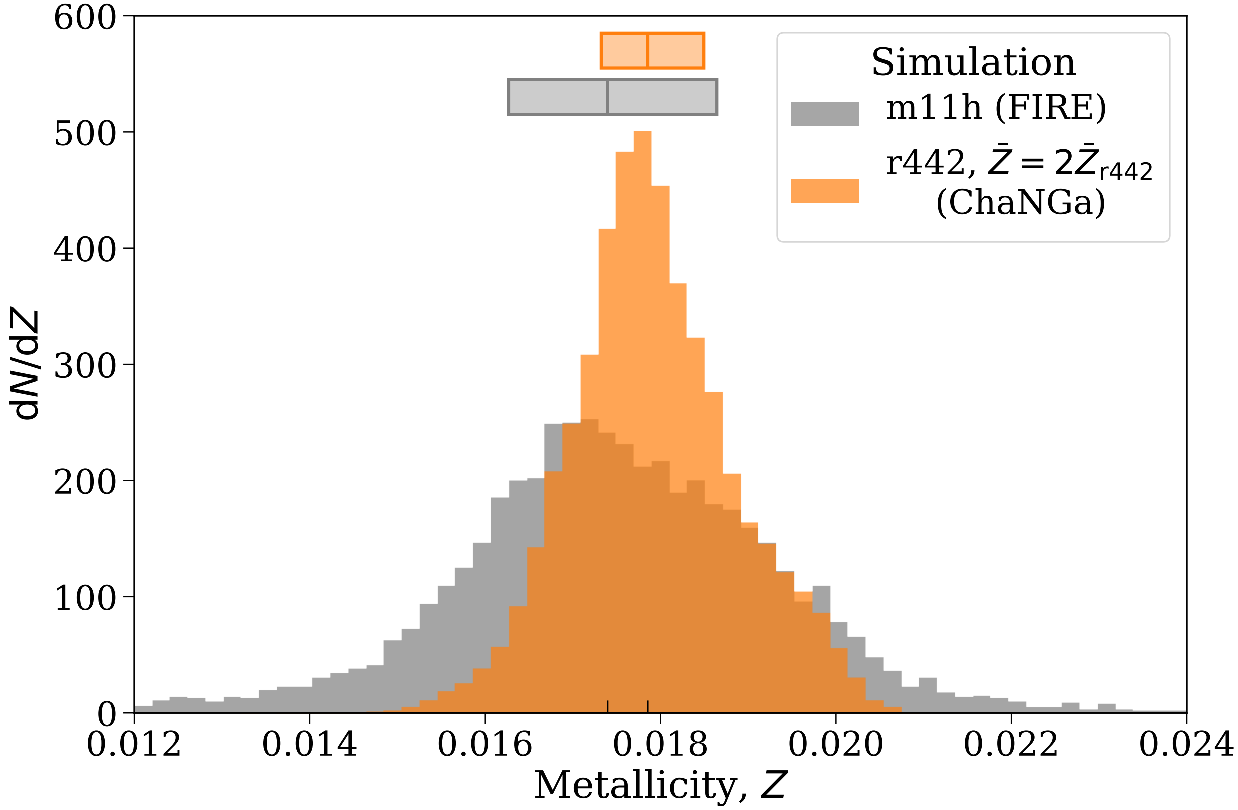

Figure 7.8: The distribution of metallicities of stars in our fiducial `m11h` simulation, compared to those in the `ChaNGa` `r442` simulation with a factor 2 increased metallicity. The boxes at the top indicate the interquartile range, with the median shown with a line in the box. As noted in Section 4.5, the average metallicity is very similar between the simulations, with a slight difference in the width of the distribution.

In addition to changes in the mass and internal structure, the average metallicity of `r442` is a factor of ∼2 lower than `m11h`, which as discussed in Section 4.5, will strongly affect the binary evolution. We therefore increase the metallicity of each star particle in `r442` by a factor of two, such that the median metallicity of the stellar populations seeded from each galaxy are now very similar. As a result we can better isolate the effects of changing galaxy simulations. We show the resulting metallicity distributions in Figure 7.8. We note that the width of the metallicity distribution for `r442` is still slightly wider in `m11h`, which may contribute to slight differences between the populations.

In the final panel of Figure 7.7 we show the timing and distance distribution of SNe for an `r442` simulation with a factor of 2 increase in metallicity. We find that the distributions of SN times between the two simulations are statistically indistinguishable, consistent with the lack of expected other impacts on binary populations once metallicity is controlled for.

The same similarity is found for the distance distributions. Overall, our results show no strong dependence on the hydrodynamical simulation code used to seed star formation. Moreover, our results indicate that a change in the mass of the gravitational potential by a factor of three do not strongly affect the spatial distribution of SNe.

5 Analytic model for feedback

The results in Section 4 suggest that there is enough consistency in the behaviour of SNe populations that models of stellar feedback should be able to include the effects of binary evolution without being overly sensitive to choice of parameters. Hydrodynamical simulations typically use "subgrid models" to implement stellar feedback. These models are currently based on single-star stellar evolution, and are implemented as simple parameterised analytic models (e.g., Hopkins et al., 2023a).

In this section, we develop an analytic model appropriate to evolving binary star populations. We fit both the rate of core-collapse SNe over time, and the velocities of SN progenitors, including the effects of metallicity. This model can be used in future hydrodynamical simulations to more consistently model the times and locations of core-collapse SN feedback.

We make a sampling routine for our analytic model available as a simple Python script. This routine can be used to rapidly generate an array of times and velocities for all SNe associated with a given star particle in a hydrodynamical simulation. The routine is available on the GitHub repository[4] for this paper.

Our general procedure for sampling SNe resulting from a starburst of a given metallicity, Z, is as follows. We first use our model to sample the time, $t_{\rm SN}$, at which each SN occurs (Eq. 7.2). Based on this time, we sample a progenitor velocity, $v_{\rm SN}$ (Eq. 7.19). The distribution for this velocity changes based on (a) whether the progenitor was ejected from its binary (Eq. 7.5) and, if so, (b) what type of mass transfer it experienced mass transfer before doing so (Eq. 7.8–7.11). This procedure is illustrated in a flowchart in Figure 7.16. In the following subsections we outline our model for each of these distributions and assess their goodness-of-fit to our simulations.

5.1 Core-collapse SN rate

In the top panel of Figure 7.9 we plot in red the distribution of SNe explosion times for our fiducial model. This distribution has a number of clear features, including a rapid rise after 3.5 Myr to a peak at 6 Myr, the distribution then declines, with a knee around 25 Myr, which is produced by the transition from binaries that experience case A or case B mass transfer to those that experience case C mass transfer (see Section 3.1 for a discussion of this feature). After this decline there is a sharp drop at 44 Myr associated with the maximum

---

[4] `https://github.com/TomWagg/supernova-feedback/blob/main/analytic_feedback_model.py`

age of a single-star that can still reach core-collapse, followed finally by a tail to long times due to the subset of SNe from merger products. Existing SNe prescriptions capture some of these features, but all lack the tail beyond 44 Myr, which contains 25% of the entire SNe population.

We develop a fitting formula for this distribution that captures these various transitions. We first define the rate of core-collapse SNe, $\mathcal{R}_{\mathrm{CSSN}}(t)$, as a piecewise power-law, with an exponential tail, as a function of time

$$\frac{\mathcal{R}_{\mathrm{CSSN}}(t/\,\mathrm{Myr})}{\mathrm{Gyr\,M_\odot}} \equiv \begin{cases} 0 & t < t_1, t > t_6, \\ a_i (t/t_i)^{\psi_i} & t_i \leq t < t_{i+1}, i < 5, \\ a_6 e^{-t/t_5} & t_5 \leq t < t_6 \end{cases} \tag{7.2}$$

where

$$\begin{aligned} a_i = [&0.38 + 0.13\,[\mathrm{Fe/H}], 0.47 + 0.05\,[\mathrm{Fe/H}], \\ &0.22 + 0.02\,[\mathrm{Fe/H}], 0.13, 0.1, 0.175 + 0.05\,[\mathrm{Fe/H}]], \end{aligned} \tag{7.3}$$

$$\begin{aligned} t_i = [&3.5, 6, 23 - 6.5\,[\mathrm{Fe/H}], \\ &28 - 6.5\,[\mathrm{Fe/H}], 45.5 - 16.5\,[\mathrm{Fe/H}], 200], \end{aligned} \tag{7.4}$$

the subscript i ranges from 1 to 6 and represents the different physically-motivated transition points that we discussed above, $\psi_i = \ln(a_{i+1}/a_i)/\ln(t_{i+1}/t_i)$, and $[\mathrm{Fe/H}] \approx \log_{10}(Z/Z_\odot)$ is the metallicity of the starburst relative to solar, where we assume $Z_\odot = 0.0142$.

Given that some of the features in the distribution in Figure 7.9 are metallicity-dependent (see arguments in Section 4.5), we include metallicity-dependent transition points. The transition from binaries that experience case A or case B mass transfer to those that experience case C mass transfer occurs around 25 Myr at the metallicity of `m11h` and moves to later times at lower metallicity (by 2 Myr for every factor of 2 decrease in metallicity). The final single star SN occurs 44 Myr at the metallicity of `m11h` and increases by 5 Myr for every factor of 2 decrease in metallicity.

The analytic model described above closely reproduces our simulated population, accounting for both metallicity and effects from binarity, shown as the black line in the top panel of Figure 7.9. The model reproduces the late SN rate to 0.5% and overall normalisation to 0.3%. This model produces approximately 1.2 SNe for every 100 solar masses of star formation, which is in good agreement with the current `FIRE`-3 model (Hopkins et al., 2023a). In the lower panel of Figure 7.9, we show the metallicity dependence of the model for each of the metallicities that we consider in the variations in Section 4.5. In each case the model reproduces the late SNe rates and normalisations of our lower metallicity simulations to within 1%.

5.2 Spatial distribution of core-collapse SNe

The variations presented in Section 4 show that the SN distance distribution is dependent on the assumed initial cluster velocity dispersion, and on the metallicity and galactic potential of the simulated galaxy. An analytic model of the distance distribution in our fiducial model

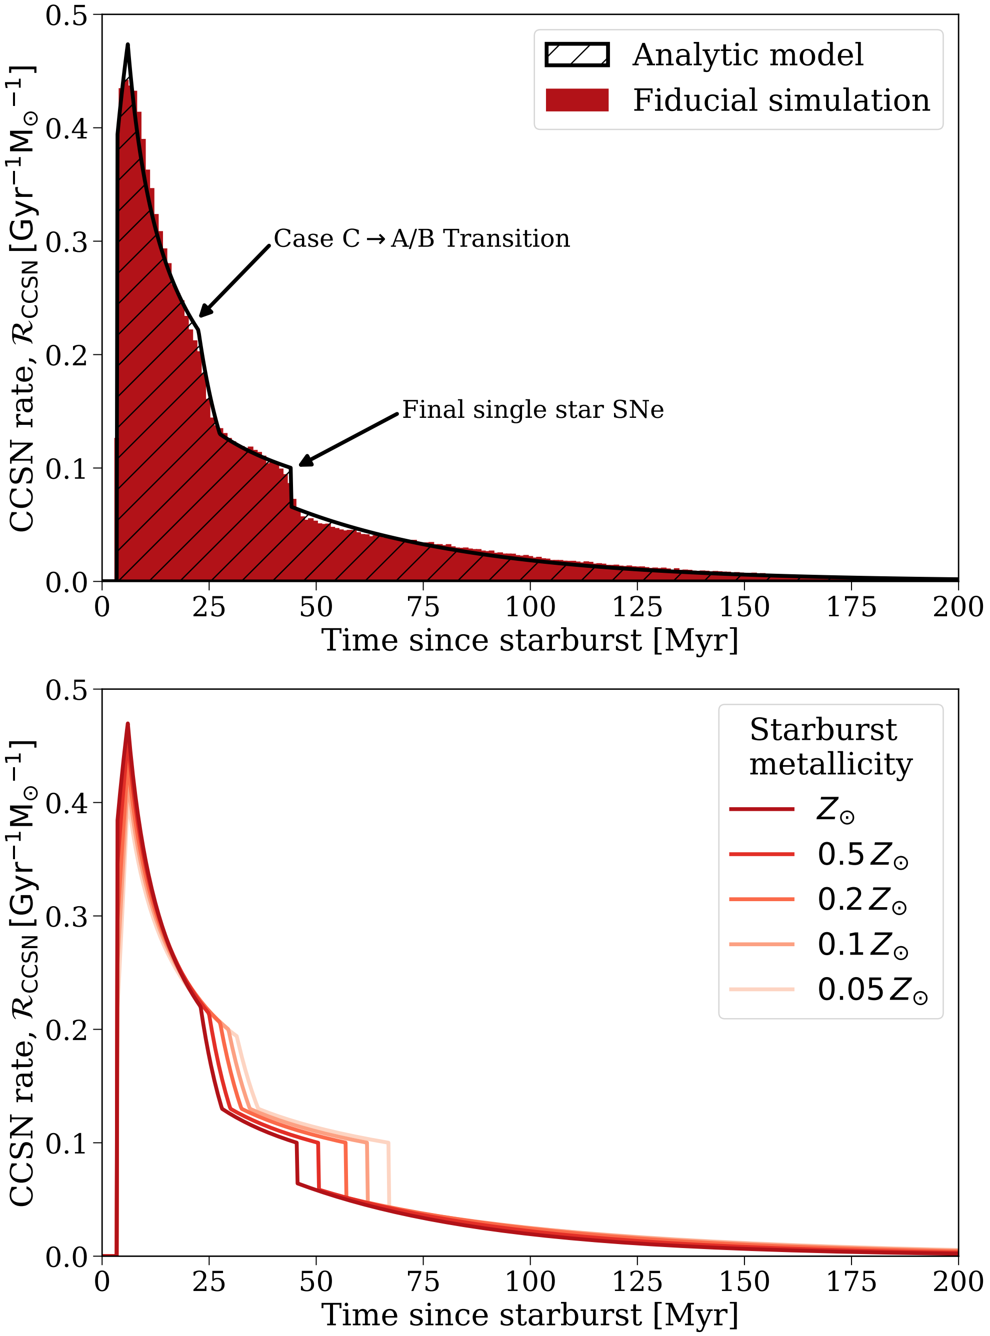

Figure 7.9: Top: A comparison of our analytic model (Eq. 7.2, shown in black) to the SNe delay time distribution to our fiducial simulation (shown in red). Transition points are annotated with the physical process driving them. The model reproduces the late SN rate and overall normalisation to within 0.5%. **Bottom:** The metallicity dependence of our analytic model to the SN time distribution. Lower metallicity star formation events produce more SNe that typically occur later. Our analytic model reproduces the late SNe rates and normalisations of our lower metallicity simulations to within 1%. (💻 Interactive figure available.)

may therefore not always be applicable other simulations. Therefore, we instead chose to model the velocity distribution at which supernova progenitors travel away from their parent cluster, which can more easily be adapted to other contexts.

There are two components to the velocity distribution of SNe precursors: (1) the velocities of primary stars, merger products, and unejected secondaries, which are all dominated by cluster dissolution through the initial cluster velocity distribution v_{disp}; and (2) the velocities of progenitors that are ejected from binaries after a companion's SN. Binary physics and metallicity will shape the second of these distributions, and set the fraction of SNe that go through the latter channel.

We adopt a definition of a progenitor being considered ejected as a walkaway or runaway star if the ejection velocity is greater than $5\,\mathrm{km\,s^{-1}}$ (following e.g., Eldridge et al., 2011; Renzo et al., 2019). Given this definition, we estimate that the fraction of core-collapse SNe that originate from an ejected progenitor is approximately,

$$f_{\mathrm{eject}}(t) = \begin{cases} 0.24 & 5 \leq t_{\mathrm{SN}}/\,\mathrm{Myr} < t_5(Z), \\ 0.1 & t_5(Z) \leq t_{\mathrm{SN}}/\,\mathrm{Myr} < 60, \\ 0.0 & \text{else} \end{cases} \tag{7.5}$$

where t_{SN} is the time at which the SN occurs and t_5 is the time of the last single-star SN, defined in Eq. 7.4. We find that this relation does not have a significant metallicity dependence and varies by less than 1% even when decreasing the metallicity by a factor of 20.

Unejected progenitor velocity distribution — For most SNe progenitors, the velocity at which they move away from the cluster is determined solely by the initial cluster velocity dispersion[5]. Under the assumption that each star moves isotropically away from the cluster following a 3D Gaussian with a dispersion velocity of v_{disp}, the probability that a SN progenitor moved away with a given 1D velocity follows a Maxwellian as

$$p_{\mathrm{unejected}}(v) = \frac{1}{\sigma^3}\sqrt{\frac{2}{\pi}}v^2 e^{-v^2/(2\sigma^2)}, \tag{7.6}$$

where v is the velocity in $\mathrm{km\,s^{-1}}$, $\sigma = v_{\mathrm{disp}}/\sqrt{3}$ and v_{disp} is the initial velocity dispersion of clusters.

Ejected progenitor velocity distribution: — For progenitors ejected from their binaries, there is an additional component to the velocity beyond the initial cluster velocity dispersion, which is shown in the red histogram in the top panel of Figure 7.10 for our fiducial model. The distribution peaks quickly around $8\,\mathrm{km\,s^{-1}}$ before declining gradually up to $100\,\mathrm{km\,s^{-1}}$.

We model the ejection velocity distribution as a mixture of four physically-motivated distributions that are distinguished by what type of mass transfer the star experienced prior to ejection; mass transfer significantly alters the pre-SN orbital separation of the system, and hence the orbital velocity and ejection speed of a companion. With this assumed mixture model, the probability distribution for the velocity of an ejected SN progenitor at a given

---

[5] The galactic potential can also alter these velocities, producing accelerations or decelerations based on the trajectory of an individual star, but this effect is generally secondary to the velocity dispersion.

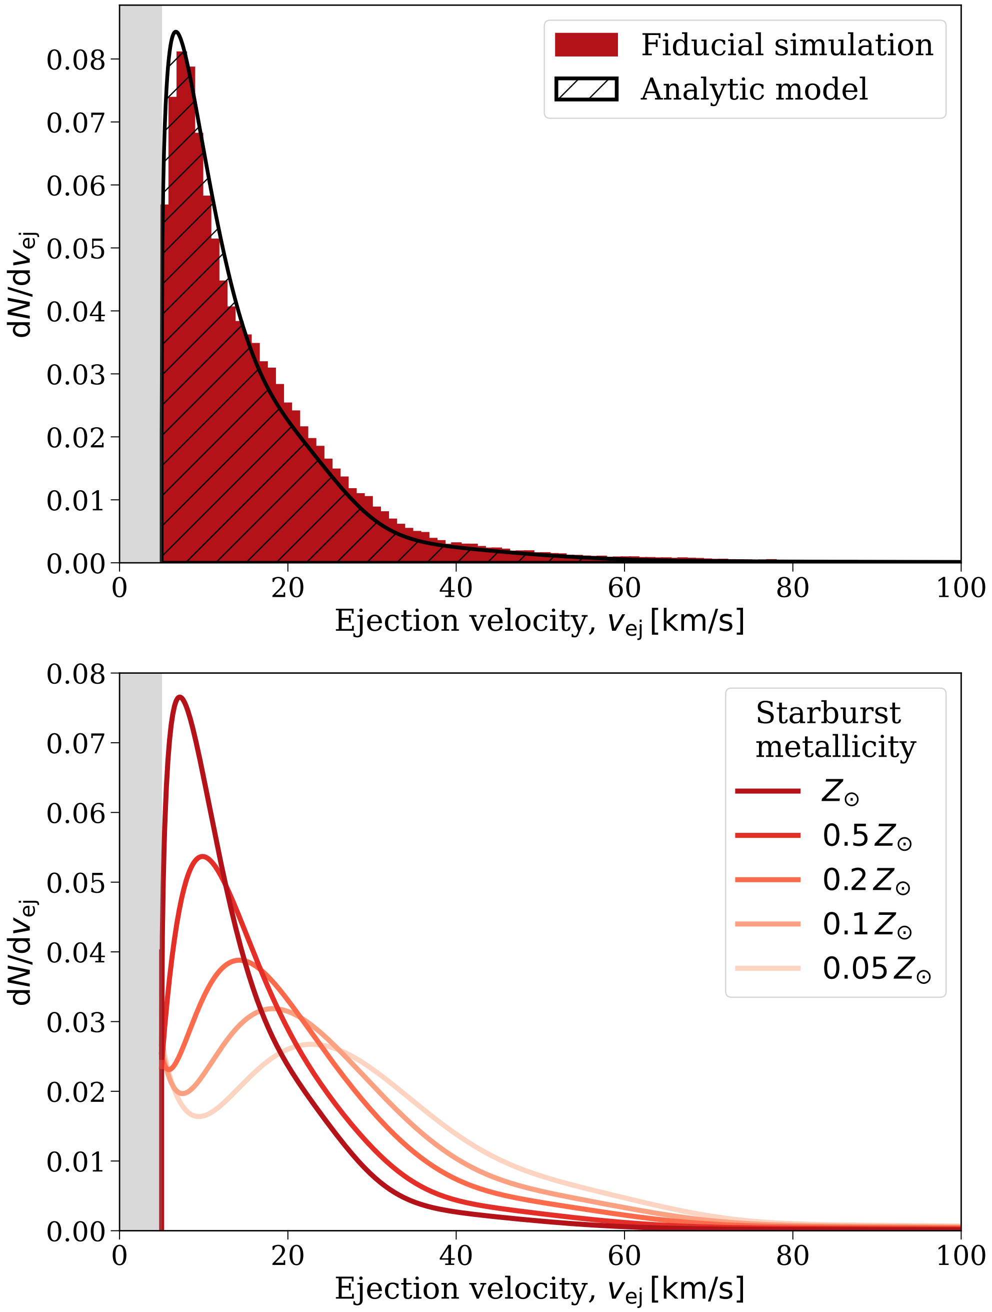

Figure 7.10: Top: A comparison of our analytic model for the distribution of ejection velocities (Eq. 7.7, shown in black) to our fiducial simulation for the subset of SN progenitors that are ejected with $v_{\rm ej} \geq 5\,{\rm km\,s^{-1}}$ (shown in red, with the lower velocity region shaded in grey). **Bottom:** The metallicity dependence of our analytic model for the ejection velocity distribution. Lower metallicity binaries typically produce higher ejection velocities. (💻 Interactive figure available.)

metallicity can be written as

$$\begin{aligned} p_{\rm eject}(v|Z) = {} & f_{\rm noMT}(Z)\, p_{\rm noMT}(v|Z) \\ & + f_{\rm MT,A}(Z)\, p_{\rm MT,A}(v|Z) \\ & + f_{\rm MT,B/C}(Z)\, p_{\rm MT,B/C}(v|Z) \\ & + f_{\rm CE}(Z)\, p_{\rm CE}(v|Z) \end{aligned} \tag{7.7}$$

where v is the ejection velocity in $\rm km\,s^{-1}$, Z is the metallicity, $f_{\rm noMT}$ and $p_{\rm noMT}(v|Z)$ are the fraction of ejected stars that had no mass transfer before ejection and the velocity distribution of those progenitors, respectively, while $f_{\rm MT,A}$ and $p_{\rm MT,A}(v|Z)$ are the same for stars that experienced case A mass transfer, $f_{\rm MT,B/C}$ and $p_{\rm MT,B/C}(v|Z)$ are the same for stars that experienced case B/C mass transfer, and $f_{\rm CE}$ and $p_{\rm CE}(v|Z)$ are the same for stars that experienced a common envelope. We fit the fractions of stars in each subpopulation to our five metallicity variations and find

$$f_{\rm noMT}(Z) = 0.14 - 0.12\,[{\rm Fe/H}], \tag{7.8}$$

$$f_{\rm MT,A}(Z) = 0.12 + 0.035\,[{\rm Fe/H}], \tag{7.9}$$

$$f_{\rm MT,B}(Z) = 0.67 + 0.12\,[{\rm Fe/H}], \tag{7.10}$$

$$f_{\rm CE} = 1 - f_{\rm noMT} - f_{\rm MT,A} - f_{\rm MT,B/C}. \tag{7.11}$$

The metallicity dependence in these fractions arises because the lower radial expansion of stars at low metallicity (e.g., Brunish & Truran, 1982; Xin et al., 2022; Klencki et al., 2022) alters when (and whether) Roche-lobe overflow occurs.

We develop fitting formulae for the distribution of ejection velocities from these subsets of binaries.

The distribution of velocities for stars that did not experience mass transfer before ejection, $p_{\rm noMT}$, follows a truncated power law distribution as

$$p_{\rm noMT}(v|Z) = \begin{cases} A v^{-1.8 + 0.5\sqrt{|\,[{\rm Fe/H}]|}} & v_{\rm min} \le v \le v_{\rm max} \\ 0 & \text{else} \end{cases} \tag{7.12}$$

where A is a normalisation constant to ensure the distribution integrates to unity. The use of a power law in this case follows naturally from the assumption that the initial orbit period distribution also follows a power law, because no mass transfer occurs to drastically alter the shape of the orbital period distribution before ejections.

For stars that experienced case A mass transfer, the distribution of velocities, $p_{\rm MT,A}$, is restricted to moderate values, and is well-fit by a normal distribution as

$$p_{\rm MT,A}(v|Z) = \mathcal{N}(22 - 8\,[{\rm Fe/H}], 6 - 3\,[{\rm Fe/H}]). \tag{7.13}$$

Case A mass transfer doesn't produce slow ejections since that would require wider binaries, which result in case B/C mass transfer. Similarly faster ejections don't occur since the shorter orbital period required most often leads to a merger instead.

The distribution of velocities for stars that experienced case B or case C mass transfer before ejection, $p_{\mathrm{MT,B/C}}(v|z)$, is well-fit by a scaled beta distribution (e.g., Beyer, 1987) that evolves with metallicity as

$$p_{\mathrm{MT,B/C}}(v|Z) = \frac{\beta(x, \alpha_{\mathrm{B/C}}, \beta_{\mathrm{B/C}})}{v_{\mathrm{max}}}, \tag{7.14}$$

where

$$x = \frac{v - v_{\mathrm{min}}}{v_{\mathrm{max}}}, \quad v_{\mathrm{min}} = 5, \quad v_{\mathrm{max}} = 100, \tag{7.15}$$

$$\alpha_{\mathrm{B/C}} = 1.5 - 1.5\,[\mathrm{Fe/H}], \quad \beta_{\mathrm{B/C}} = 18 + 5\,[\mathrm{Fe/H}]. \tag{7.16}$$

For stars that experiences a common-envelope, the velocity distribution also well-fit by a scaled beta distribution, but skewed to higher velocities as a result of how common-envelopes shrink binaries.

$$p_{\mathrm{CE}}(v|Z) = \frac{\beta(x, \alpha_{\mathrm{CE}}, \beta_{\mathrm{CE}})}{v_{\mathrm{max}}}, \tag{7.17}$$

where

$$\alpha_{\mathrm{CE}} = 5 - 4\,[\mathrm{Fe/H}], \quad \beta_{\mathrm{B/C}} = 10. \tag{7.18}$$

The analytic model for runaway star ejection velocities accurately reproduces the distribution we find in our fiducial simulation. In the top panel of Figure 7.10 we compare the model to the distribution of velocities for all stars that are ejected from their binary with a velocity greater than $5\,\mathrm{km\,s^{-1}}$. The fraction of stars with velocities greater than $10, 30$, and $50\,\mathrm{km\,s^{-1}}$ are each reproduced within 3%. In the bottom panel of Figure 7.10, we show how the model changes as a function of metallicity. Similarly, at lower metallicity, the fraction of stars with velocities greater than $10, 30$, and $50\,\mathrm{km\,s^{-1}}$ are each reproduced within 4%.

Overall, with these two distributions one can write the full SN progenitor velocity distribution as follows

$$p(v|Z,t) = f_{\mathrm{eject}}(t)p_{\mathrm{eject}}(v|Z) + (1 - f_{\mathrm{eject}}(t))p_{\mathrm{unejected}}(v). \tag{7.19}$$

5.3 Joint distribution

Our analytic model reproduces the joint distribution from the fiducial simulation well. In Figure 7.11 we compare the joint distribution of SN times and the linear distance an SN progenitor could travel (the product of the velocity and time). The top panel shows the distribution for the fiducial simulation, which is similar to Figure 7.3 but with a skew to higher distances since this proxy for distance does not account for how the potential may alter the progenitor's velocity. In the middle panel we show the same distribution for a random sample from our analytic model with the same number of SNe. From the bottom panel, one can note that the model closely reproduces the distribution from the fiducial simulation, particularly within the 98% contour of the fiducial simulation (shown with the black line).

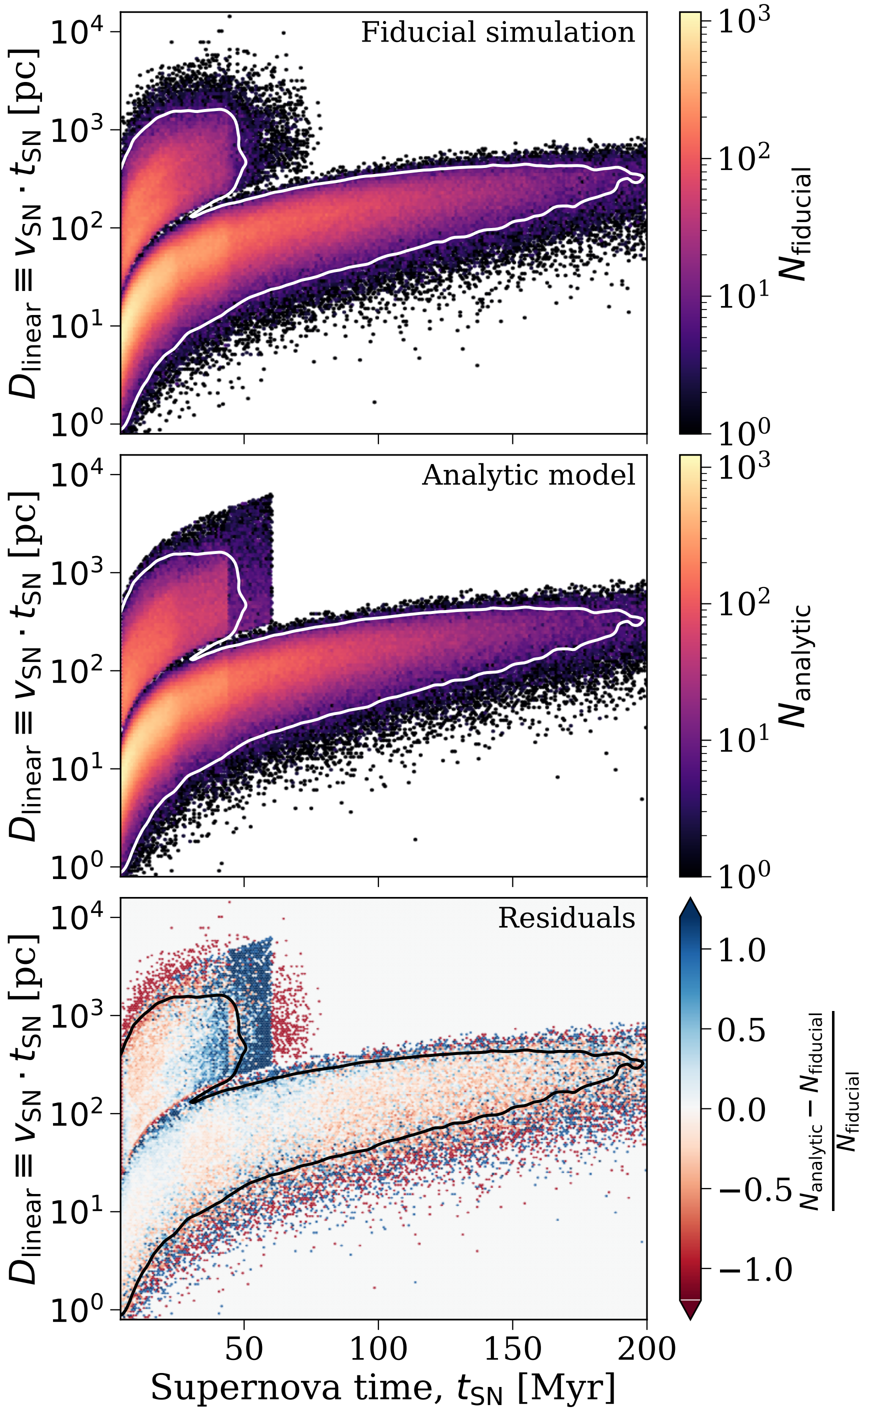

Figure 7.11: A comparison of the joint distribution of our analytic model to the fiducial simulation, which shows good agreement. Each panel shows a 2D histogram with the supernova time on the x-axis and the maximum distance travelled (i.e. the product of the SN time and velocity) on the y-axis. The top panel shows the distribution of the fiducial simulation, while the middle shows the analytic model and the bottom gives the fractional difference in each bin from the top two panels. The solid contour lines in each panel show the 98% regime for the fiducial simulation.

6 Discussion

In this Section, we summarise key takeaways (Section 6.1), consider the potential implications of our results for galaxy evolution (Section 6.2), and highlight potential limitations of our analysis (Section 6.3).

6.1 Key takeaways of our results

6.1.1 High-level differences between binary and single star feedback models

We have demonstrated that binary interactions significantly shift the timing and spatial distribution of SNe when compared to an equivalent single star population. For our fiducial model, we find that the median SN occurs 22 Myr after a star formation event and at a distance of 35 pc from its parent cluster. These quantities are 29% and 52% larger, respectively, than for a single star population (17 Myr and 23 pc respectively).

We also find that the total number of SN is ∼11% higher for the binary fiducial model than in the equivalent model including only single stars. Including binary interactions allows lower mass stars (i.e. those initially below the threshold for a SN) to accrete enough material, or merge, to attain the requisite mass to achieve core collapse.

In addition to lengthening the time and length scales for SNe energy injection, the inclusion of binary interactions also introduces long tails in the distributions to late times and long distances. For our fiducial model we find that $f_{t>44\,\mathrm{Myr}} = 25\%$ and $f_{D>100\,\mathrm{pc}} = 13\%$, compared to $f_{t>44\,\mathrm{Myr}} = 0\%$ and $f_{D>100\,\mathrm{pc}} = 1\%$ for the single star model.

These differences in feedback, particularly the increased spatial spread of SNe, have the potential to significantly affect the SNe energy injection in the ISM, which we will discuss in further detail in Section 6.2. Given that the vast majority of core-collapse progenitors are formed in interacting binary systems (e.g., Sana et al., 2012), these differences should be accounted for in models of feedback, motivating the analytic prescription provided above in Section 5.

6.1.2 The robustness of binary SN feedback models

Through varying many input parameters, we find that these predictions for SN timing and distance are remarkably robust to a wide range of extreme variations to binary physics, initial conditions and galaxy parameters. In Figure 7.12, we summarise how the median time and distance of SNe changes under the variations that we consider. The majority of variations retain a median value within 15% of the fiducial variation. The exceptions to this for timing are variations of metallicity, the initial mass function, and the stability of mass transfer. For SN distances, changes to the initial velocity dispersion can significantly shift the distribution.

In Figure 7.13, we summarise how the tails of the timing and distance distribution change for variations that we consider. The tails of the distance distribution are insensitive to most changes, only significantly shifting by more than 15% for variations for metallicity, the stability of mass transfer, and initial cluster velocity dispersion. With the exception of the

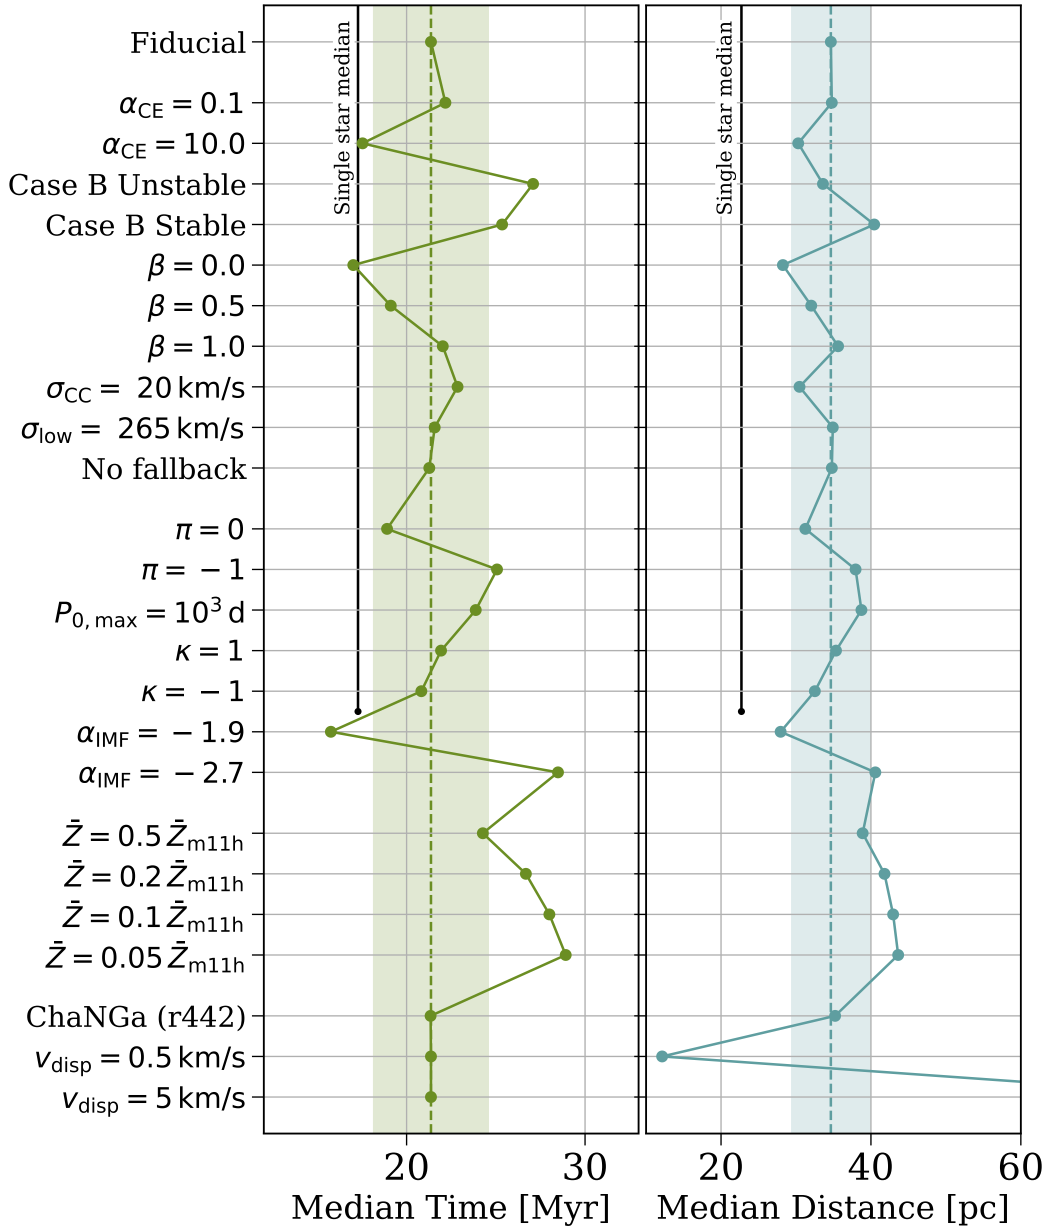

Figure 7.12: The median time and distance at which SNe occur for each variation (variations are outlined in Section 4.1). Dotted lines and shaded regions show the fiducial value and a ±15% region around this value. The solid black lines show the median value of a population of single stars equivalent to our fiducial model. The lines don't continue for all variations, since the later ones would alter a single star population as well. Given the extreme physics variations that we consider, the relatively small differences in these medians demonstrates an encouraging robustness of our results to uncertainties. The exact values for each variation are given in Table 7.4 & 7.5 (💻 Interactive figure available.)

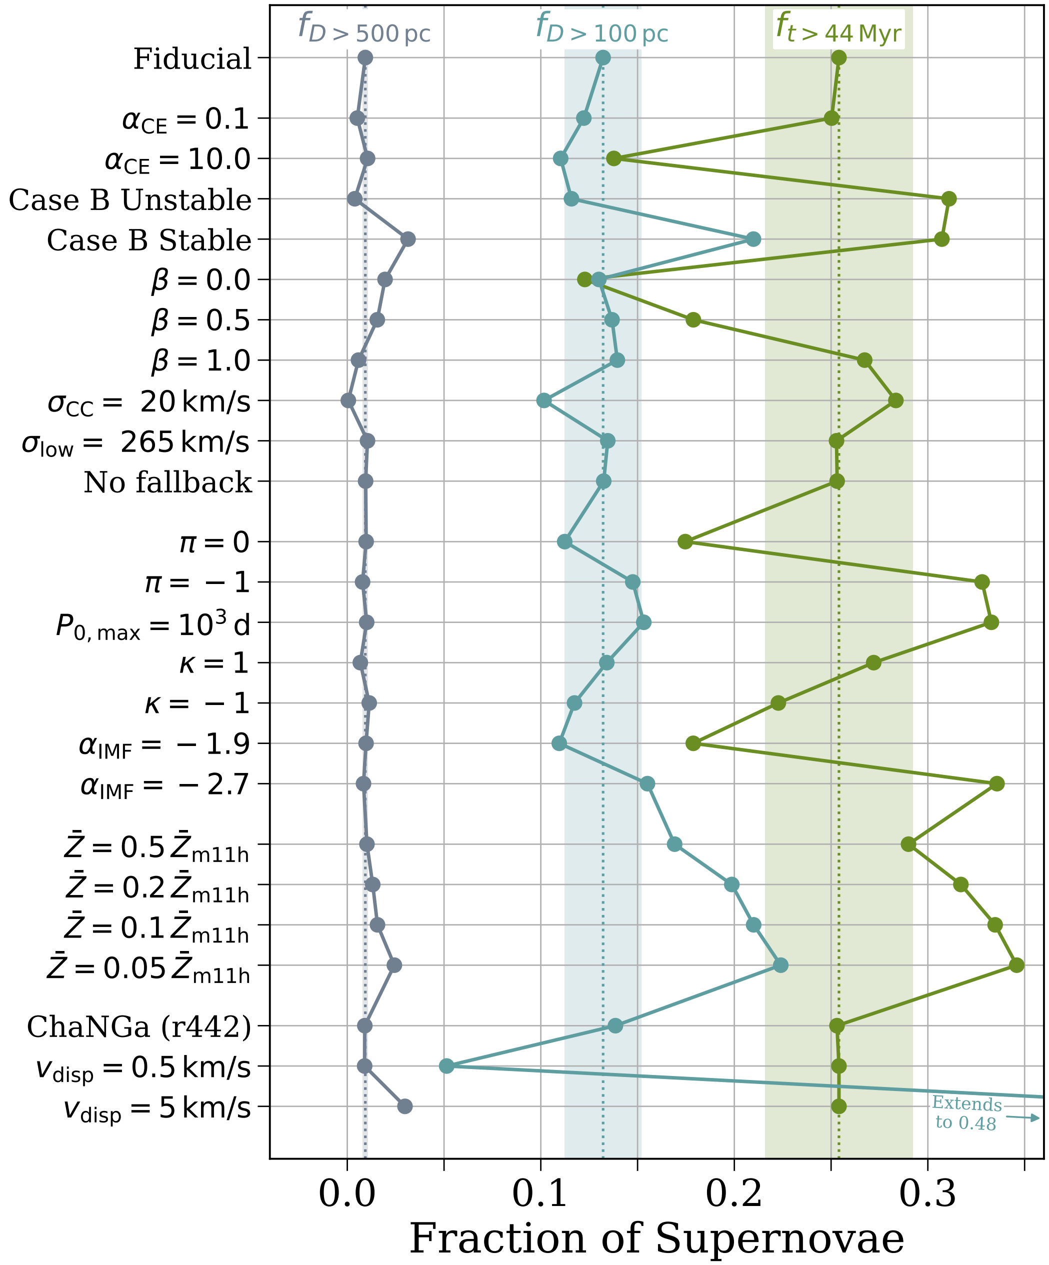

Figure 7.13: The fraction of the total population in the tails of the SNe timing and distance distributions for each variation (variations are outlined in Section 4.1). Scatter points indicate the $f_{t>44\,\mathrm{Myr}}$ (green), $f_{D>100\,\mathrm{pc}}$ (light blue), and $f_{D>500\,\mathrm{pc}}$ (slate blue) fractions. Note that 44 Myr is the assumed final core-collapse SN time in `FIRE`-3. Dotted lines and shaded regions show the fiducial value and a ±15% region around this value. The exact values for each variation are given in Table 7.2. (💻 Interactive figure available.)

low velocity dispersion variation, all variations predict that at least 10% of SNe occur more than 100 pc from their parent cluster.

In contrast, the timing tail does change by more than 15% with several variations to binary physics, initial conditions and galaxy parameters. The late time tail is most sensitive to changes in metallicity, initial masses, and orbital periods, as well as mass transfer physics. Yet even for the model predicting the fewest late SNe (in which mass transfer is fully non-conservative, $\beta = 0.0$), 13% of SNe still occur at times later than 44 Myr, the assumed limit from the single star evolution models used for feedback prescriptions in current hydrodynamical codes (e.g., Hopkins et al., 2023a).

Given the suprisingly mild variations seen in Figures 7.12 & 7.13, we expect that our analytic model for core-collapse SN feedback from binary progenitors (Section 5) should be robust to significant parameter uncertainties. The model reproduces the medians, tails and overall normalisation of the timing and velocity distribution of the fiducial model, and also accounts for possible variations in metallicity and initial velocity dispersion, which produce the most significant change in the medians of the timing and distance distributions.

A natural concern about implementing SNe feedback based on binary models is that there is much larger uncertainty compared to single star evolution. We show that the resulting time and spatial distributions of SNe are actually not nearly as variable as one might expect, especially when one uses fitting functions that take metallicity into account. Overall, given the robustness of our findings, our model for the timing and location of SN feedback represents a significant improvement upon the existing implicit assumption of single star evolution in hydrodynamical simulations.

6.2 Implications for galaxy evolution

6.2.1 Additional late core-collapse SN feedback

Binary evolution may lead to more significant impacts on feedback through producing an extended tail of late-time supernova (see Figure 7.1). This tail is not insignificant, with ∼25% of SNe happening after the oldest SNe expected for single star models.

There are a number of possible changes that could result from the significant increase in late-time SNe that binary evolution produces following a star formation event. First, the extended population of SNe will make the resulting feedback more gradual, and less impulsive. One can think of this difference as akin to changing from slamming on the brakes to slowing to a stop. This change could result in less short-term burstiness on small spatial scales, where more extended feedback could either interrupt a "hard-stop, fast-start" pattern of star formation, or, cause longer pauses before star formation can restart. Other works have previously explored the impact on "burstiness" from variations in stellar feedback's efficiency and energy/momentum injection and found that these variations have no strong effect (e.g., Orr et al., 2018; Chan et al., 2018; El-Badry et al., 2018; Hopkins et al., 2023b). However, adding in *extended* stellar feedback may well produce different results, and should be tested numerically.

A second possible impact of extended SNe would be to change the susceptibility of the ISM to feedback. The longer the timescale for SNe, the more likely it is that the local ISM will have already been changed by earlier SNe and ionization from massive stars, which

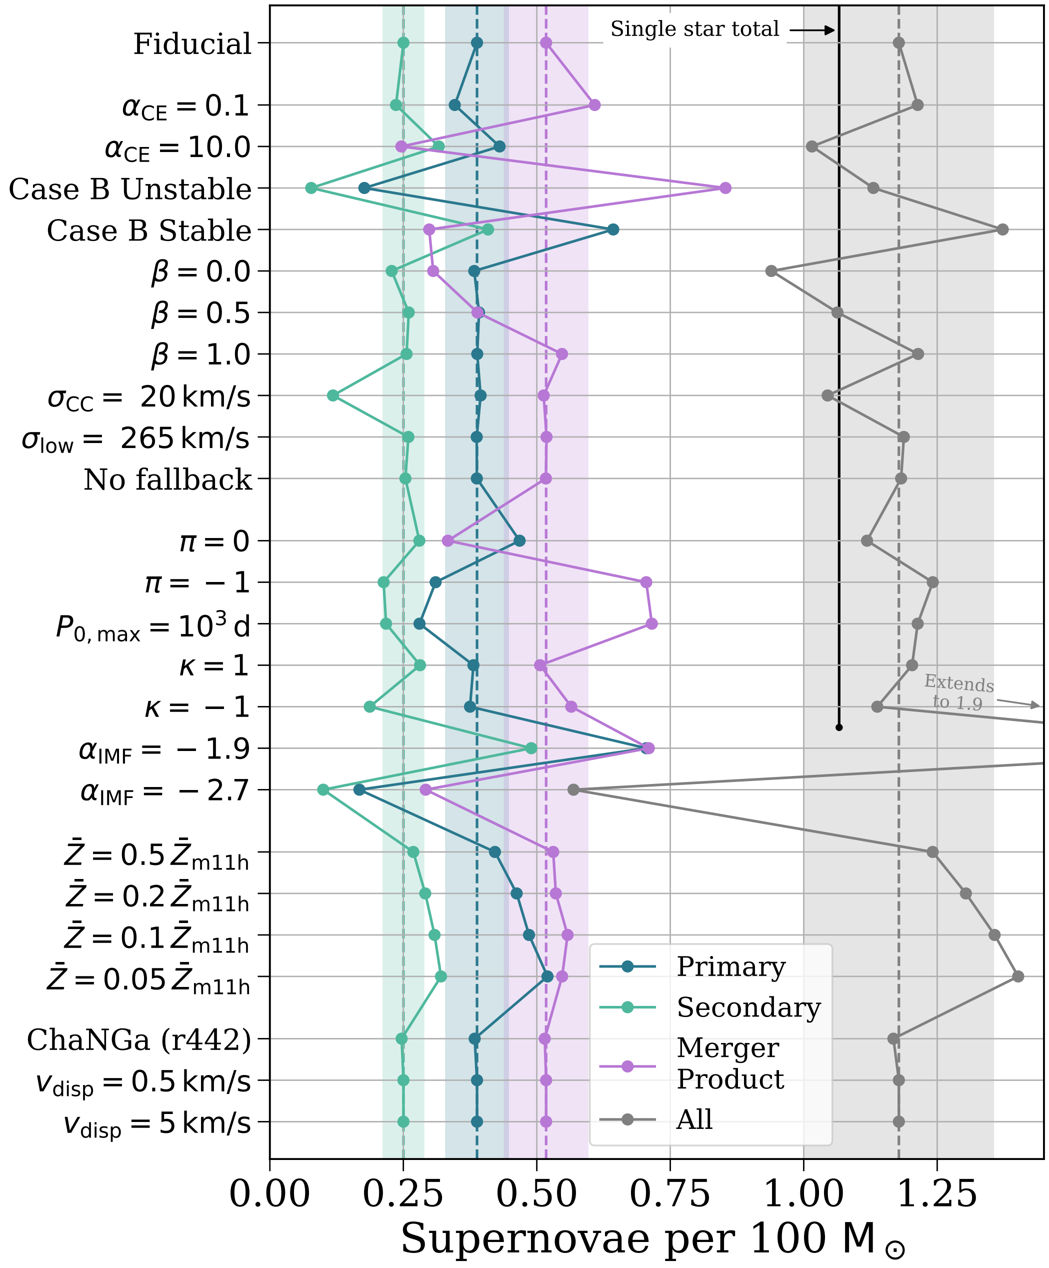

Figure 7.14: The total number of SNe that occur for each variation, separated by progenitor type (variations are outlined in Section 4.1). Dotted lines and shaded regions show the fiducial value and a ±15% region around this value. The exact values for each variation are given in Table 7.3. The solid black line shows the total for a population of single stars that is equivalent to our fiducial model, indicating that binary models typically increase the total SN feedback in a population. The line doesn't continue for all variations, since the later ones would alter a single star population as well. (💻 Interactive figure available.)

may make the impact of the late-time SNe qualitatively different (e.g., feedback into a dense molecular cloud versus feedback into an already expanding hot bubble). Even in the absence of immediate post-burst feedback, increasing the timescale of SNe allows more time for the ISM to "revert to the mean" through cooling and turbulent mixing, again changing the receptiveness of the local environment to feedback. As with assessing the impact on burstiness, full numerical tests are needed to understand the net impact.

Although we focus on core-collapse SNe in this work, we briefly consider Type Ia SNe, which are another source of binary-driven late stellar feedback. These SNe arise from slowly-evolving low-mass progenitors, have delay times that can range from ∼50 Myr to several billion years (e.g., Liu et al., 2023; Ruiter & Seitenzahl, 2025) and may contribute to the driving of winds in dwarf galaxies (e.g., Hu, 2019). The long tail of core-collapse SNe that we find may coincide with the earliest type Ia SNe, and therefore they could have a cumulative effect on the clustering of SNe and the driving of superbubbles.

6.2.2 Spatially-extended feedback

Binary stellar evolution changes where SNe occur, both through ejecting runaway stars via binary interactions, and increasing the impact of cluster dissolution through extending SNe progenitor lifetimes. Both of these effects can shift feedback away from the higher-density ISM that hosts star formation, and into low-density environments. In all of our models, at least 12−15% of all SNe occur more than 0.1 kpc from the centre of the clustered star formation (Figure 7.13. This separation is more than the size of even quite massive molecular clouds (e.g., Chevance et al., 2020), and should be sufficient to move a significant fraction of SNe to lower density environments.

Past numerical simulations have found that more distributed, low-density feedback — comparable to what is expected here — may increase the likelihood of driving galactic outflows (e.g., Ceverino & Klypin, 2009; Ceverino et al., 2014; Zolotov et al., 2015; Hu et al., 2017; Andersson et al., 2020; Steinwandel et al., 2023). However, other works have argued that the effect on outflows is weaker (Kim & Ostriker, 2017; Andersson et al., 2023) and that earlier results may have stemmed from insufficient resolution of SNe in dense gas (Kim et al., 2020). Our updated models with a more detailed characterisation of the temporal and spatial location of SNe may affect these results and thus further numerical simulations are necessary to test this.

Spatially-distributed SNe in dwarfs and at high-redshift — An underappreciated aspect of having more spatially-distributed SNe is their effect on small galaxies. The characteristic velocities at which runaways are released and the increase in SNe progenitor lifetimes are set by processes internal to the binary population, and as such, any characteristic length scale becomes increasingly important when considering physically smaller systems such as dwarf galaxies.

As an example, for galaxies with small effective stellar or gas radii (say, $R_{\rm eff} < 1$ kpc), the $\sim 12-15\%$ of SNe that take place more than 0.1 kpc from their birth cluster will presumably travel a substantial fraction of the effective radius, making any SNe feedback automatically a galaxy-wide event.

We note several reasons beyond "size" why the possibilities for distributed or even escaping SNe may become more important for dwarfs. First, our fiducial simulation used a galaxy with a stellar mass of $4 \times 10^9\,\mathrm{M}_\odot$, which is technically a dwarf, but not particularly low mass compared to the entire galaxy population. Compared to this fiducial simulation, binaries evolving in galaxies with even shallower potentials may lead to SNe at even larger distances, and may lead to complete escape of SNe from the galaxy.

The second possible increased impact is due to the strong metallicity dependence of the fraction of large-distance SNe. Most dwarf galaxies are also low metallicity, which can nearly double the number of stars in the $D > 100\,\mathrm{pc}$ tail (Figure 7.13), increasing to over 20% of stars with $D > 0.1\,\mathrm{kpc}$ for $Z \leq 0.1 Z_\odot$.

This effect may be particularly pronounced at high redshifts, where galaxies are systematically far more compact, intrinsically lower metallicity, and (more speculatively) likely have high characteristic velocity dispersions in star forming regions. While we cannot know the exact statistical properties of binary populations at these redshifts, it seems unlikely that they would conspire to counteract these trends, given the robustness of the distance distributions (Figures 7.12 & 7.13).

In support of the speculations above, we note that the potentially enhanced role of runaways in dwarfs has already been identified in existing galaxy simulations, (e.g., Steinwandel et al., 2023), which have shown that runaway stars have the potential to drive galactic outflow rates (e.g., Hu et al., 2017), impact the overall star formation rate, and boost the energy loading factor, (e.g., Steinwandel et al., 2023), albeit under the unrealistic assumption of fixed ejection velocities for all runaways. These methods do not fully capture the more complex, long-tailed velocity distributions, or their metallicity dependence. Therefore, our models motivate future work to investigate the effect of altering the spatial distribution of the feedback in hydrodynamical simulations in this way.

6.2.3 Stellar wind feedback

We have previously limited our analysis to considering how binary interactions impact core-collapse SN feedback. However, binary interactions likely also have a significant impact on the feedback from stellar winds. In particular, binary interactions produce more stripped stars, which contribute a large amount of ionising emission (e.g., Götberg et al., 2019). While a quantitative analysis of the impact of stripped stars and wind-driven mass loss is outside the scope of the paper, we briefly comment on some qualitative expectations here.

In practice, we think that the spatial distribution of wind-driven feedback may not be as strongly affected by binarity as SNe feedback. The stripped stars produced via binary interactions are most likely to be primary, rather than secondary stars. These strong contributors to the feedback would therefore rarely be ejected from their parent cluster by the earlier evolution of a more massive companion. Moreover, any self-stripping from strong stellar winds is most prevalent in the most massive stars, which also have the shortest lifetimes and thus much less time to travel far from their parent cluster.

In total, the main effect of binary interactions on stellar wind feedback is likely on its magnitude (increased relative to single stars), rather than its spatial distribution. Nevertheless, future work should model the impact of these winds in detail.

6.3 Limitations

Although we have taken pains throughout this paper to widely explore the possible parameter space, there are some limitations that remain. We briefly discuss some of the most significant of these here, but note that many of these limitations are currently inherent in any analysis of this type, and may affect single-star evolution models as well.

Population synthesis treatment of binary physics — Our results are computed using rapid population synthesis, which relies on parametric prescriptions for stellar evolution that mimic more detailed simulations (Hurley et al., 2000, 2002; Breivik et al., 2020b). In this work, we explicitly consider the impact of varying our assumptions in these parameter prescriptions for a variety of binary physics (see Section 4), though we only vary one parameter at a time whilst some choices may be correlated. Moreover, this does not account for the fact that the parametric prescriptions themselves may not adequately describe the underlying binary physics. For example, the α-λ common-envelope prescription used in population synthesis is a relatively simplistic energy based prescription for a complex three-dimensional mass transfer event, which likely fails to capture all of the details of the process (Webbink, 1984; de Kool, 1990; Ivanova et al., 2013, 2020; Röpke & De Marco, 2023). This also highlights that population synthesis does not model processes on a thermal timescale and as such requires prescriptions for systems that are out of thermal equilibrium. Further observational constraints of massive stars and these rapid evolutionary phases are necessary to better model the binary physics and improve population synthesis models.

Reliance on pre-computed stellar tracks — Rapid population synthesis models such as `COSMIC` rely heavily on the choices made in the underlying stellar evolution models. These choices are then embedded into any binary population synthesis model that uses them and limit the possible variations one can consider.

For example, the original stellar tracks used by `COSMIC` (Pols et al., 1998) assumed one specific model for convective boundary mixing and overshooting, which are processes that can impact the final core mass and pre-supernova structure of the star (e.g., Ugliano et al., 2012; Kaiser et al., 2020).

Some of these limitations may be particularly pronounced for aspects of stellar rotation. The underlying stellar tracks used by `COSMIC` do not account for rotation, however rotation can significantly influence the evolution of massive stars (e.g., Ekström et al., 2012). For example, the efficiency of mass transfer can play an important role in the timing and location of SNe, and this may be limited by rotation-enhanced mass loss, which is particularly important for the critically-rotating accretors (e.g., Langer, 1998; Petrovic et al., 2005; Renzo & Götberg, 2021, though it remains debated whether this is physical, e.g., Paczynski 1991; Popham & Narayan 1991) that appear to be common (e.g., Bastian et al., 2017). A fuller treatment of stellar rotation could therefore impact our results.

Main sequence core evolution and rejuvenation — Most `BSE`-based population synthesis codes (including `COSMIC`, which we use in this work) rely upon single star evolution models for computing the core structure of stars (Hurley et al., 2000, 2002). However, many works have shown that binary interactions and accretion can alter the core structure and lead to rejuvenation of accretors (Hellings, 1983; Braun & Langer, 1995; Cantiello et al., 2007; Staritsin, 2019; Renzo et al., 2023; Lau et al., 2024; Wagg et al., 2024). As a result of the implicit assumption of single star evolution in core structures, the core of an accretor star that experienced case A mass transfer will have a core size commensurate with its final mass, without accounting for the core growth prior to mass transfer. Thefore, our simulations may overpredict the delay to the star's core collapse time. However, this is primarily relevant for case A mass transfer, which we find only 9% of stars in our fiducial simulation experience, and therefore we expect the impact of this effect will be small.

An additional consideration is that previous rejuvenation after accreting matter from a companion can change the binding energy of a star's envelope and affect its chance of surviving a common-envelope event (e.g., Renzo et al., 2023; Landri et al., 2025). Thus the number of merger progenitors for SNe that we predict could be affected by including this effect.

Future updates to population synthesis prescriptions for main sequence core evolution and rejuvenation are necessary to improve upon this analysis.

Explodability of massive stars — In our population synthesis models we use the Fryer et al. (2012) remnant mass prescription to infer the final core structure of the star and assume that all stars with a non-zero ejecta mass explode as SNe. However, it is possible that many of these stars may instead implode, failing to produce a SN as a result. Many works have considered the specific requirements for a massive star to end its life in a SN (e.g., O'Connor & Ott, 2011; Sukhbold & Woosley, 2014; Sukhbold et al., 2016; Ertl et al., 2016, 2020; Laplace et al., 2021; **?**; Ugolini et al., 2025). In particular, recent work has shown that the explodability depends on a variety of factors beyond CO core mass, including composition and rotation. (e.g., Patton & Sukhbold, 2020). Yet it is also important to consider that cores that fail to explode may still produce ejecta and impact their galactic surroundings (e.g., Piro, 2013; Lovegrove & Woosley, 2013; Antoni & Quataert, 2022, 2023), which, for example, may be achieved through disk outflows from massive collapsars (Siegel et al., 2022). Future work should consider how applying updated prescriptions for pre-supernova core structure in population synthesis could impact the total SN feedback produced by massive stars.

Dynamical cluster ejections — During the early evolution of clusters, stars can be ejected at high velocities as a result of dynamical N-body encounters (Poveda et al., 1967). Our simulations do not account for dynamical ejections, focusing only on creation of runaway stars via binary ejections (Blaauw, 1961; Boersma, 1961). This means that the distance distributions for SNe that we present are a lower limit on the true distribution. In reality, a fraction of these stars will receive additional dynamical kicks, which would result in more extended tails. This would additionally mean that primary star progenitors could also reach core collapse far from star forming regions, though we expect that the progenitors of the most distant SNe would remain secondary accretor stars. In principle, given mass-dependent (of the star and cluster) ejection velocity distributions from N-body simulations (e.g., Oh & Kroupa, 2016), these additional velocity kicks could be applied in our simulations at a minimal extra cost.

Higher-order stellar multiples — Progenitors of core-collapse SNe are expected to often be found in higher order stellar multiples. Indeed it has been predicted that the *majority* of O-type stars are formed in triples and higher-order stellar multiples (Offner et al., 2023). These systems could produce similar effects to dynamical cluster ejections (see above) by ejecting companions, and hence they should be considered in future work. Our simulations cannot evolve higher-order stellar multiples but there are several codes that address this problem, such as `TRES`, `MSE` and `TSE` (Toonen et al., 2016; Hamers et al., 2021; Stegmann & Antonini, 2024; Preece et al., 2024). These codes could be coupled to galactic dynamics (as `cogsworth` does with `COSMIC` and `Gala`) and applied in future work.

Self-consistency of post-processing hydrodynamical simulations — In this work we have used the star formation history of `FIRE m11h` (and the `ChaNGa` galaxy `r442` as a variation) to seed star formation in `cogsworth` simulations. However, we did not adjust the feedback in the underlying hydrodynamical simulations to account for how binary physics may affect it and therefore change the star formation history of the galaxy. Nevertheless, given that the effects on feedback that we discuss are dependent on stellar evolution and binary interactions we do not expect that changes to the overall star formation rate would impact our results significantly.

7 Comparison to previous work

This paper presents a thorough exploration of the impacts of binary evolution models on key aspects of SNe feedback. However, this work builds upon a large number of papers that have also explored some of these same issues, using a variety of techniques and physical assumptions. Here, we place our results in the context of this existing body of work, exploring points of agreement and seeking insight from areas of potential tension.

7.1 Core-collapse SN timing

The impact of binaries on the delay time distribution of core-collapse SNe has been considered in several previous works (e.g., De Donder & Vanbeveren, 2003; Xiao & Eldridge, 2015; Zapartas et al., 2017). These works adopt a series of different assumptions regarding initial conditions and binary physics, as well as use different simulation codes.

The earliest work exploring these effects systematically is De Donder & Vanbeveren 2003 (hereafter DD03), which used the population synthesis code described in Vanbeveren et al. (1998) to explore the galactic evolution of SN rates. Similar to our work, they make predictions for the rate of SNe after a fixed starburst at solar metallicity.

DD03 generally find a significantly higher rate of late SNe than our work. We use the data from their plots, which was digitized by Zapartas et al. (2017), to compute that for a 100% binary population, 52% of their SNe occur after 44 Myr. This fraction is double the rate of late SNe that we find in our fiducial model (∼25%).

The discrepancy between our results can be somewhat explained by differences in our choices regarding initial conditions. In their simulations, DD03 assume a bottom-heavy IMF with $\alpha_{\rm IMF} = -2.7$, thus their result is more comparable to our variation in which we find $f_{t>44\,{\rm Myr}} = 33\%$ (see Figure 7.13). Moreover, they assume an upper orbital period limit of 10^3 days. For this limit we also find $f_{t>44\,{\rm Myr}} = 33\%$. It is likely that a combination of this upper limit and IMF produces a much higher fraction of late SNe. Their assumed orbital period distribution is different as well, further complicating a precise comparison with work we present here. An additional consideration is that modern prescriptions of stellar wind mass loss are generally more conservative, which may lead more of our massive stars to reach core collapse on shorter timescales. Overall, the difference between our results are understandable given the various updates to binary evolution models in the past two decades.

Zapartas et al. 2017 (hereafter Z17) is a more recent work that investigated the delay time distribution of core-collapse SNe, exploring a series of variations similar to ours. Their physical assumptions are closer to those we adopt as well, making a comparison more straightforward than with DD03. They do use a different binary synthesis code, `binary_c` (Izzard et al., 2004, 2006, 2009, 2018; Izzard & Jermyn, 2023), but it is based on the same underlying fitting formulae as `COSMIC` (Pols et al., 1998; Hurley et al., 2000, 2002; Breivik et al., 2020b).

Our results for the rate of late SNe are close to agreement with Z17. In their model assuming a binary fraction of 100%, they find that 20% of core-collapse SNe occur at least 50 Myr after a star formation event. For most binaries in their simulation (those with $m_1 < 15\,\mathrm{M}_\odot$) they assume a power law slope for their initial orbital periods of $\pi = 0$ following Öpik (1924). For our $\pi = 0$ simulation, we find that 15% of core-collapse SNe occur after 50 Myr. The remaining discrepancy in our results can be understood through a difference in the upper period limit of binaries. In particular, we assume that upper orbital period limit is $10^{5.5}$ days, whilst Z17 assumes a limit of $10^{3.5}$ days. As such, we will produce more wide, non-interacting binaries, which skew our results to having fewer late SNe (see Section 4.4.2).

We are in broad agreement with Z17 on the trends associated with variations in initial conditions and binary physics. This agreement is evident from a comparison of our Table 7.2 to Table 2 of Z17. We note that the absolute values are shifted given the difference in our assumptions regarding the fiducial binary fraction and upper orbital period limit.

One point of apparent difference is that we find an opposite trend with metallicity. We find that the rate of late SNe increases with decreasing metallicity, whilst Z17 find the opposite trend. We have traced this difference to different choices for the definition of “late”. For our models, we use the `FIRE`-3 limit of 44 Myr, while Z17 use the time of the final SN from a single star at each different metallicity. These alternate definitions lead to our outwardly conflicting claims regarding the trend with metallicity. We explain this reasoning in more detail, and confirm that using the same criteria in our simulations results in an agreement with Z17, in Appendix 9.1.

7.2 Spatial distribution of core-collapse SNe

Earlier works have considered the rates, ejection velocities of massive runaway stars and the impact on the spatial distribution of core-collapse SNe (e.g., De Donder et al., 1997; Eldridge et al., 2011; Boubert & Evans, 2018; Renzo et al., 2019). We compare directly to a subset of these works which considered the spatial distribution of massive stars relative to star forming regions.

Eldridge et al. (2011) (hereafter E11) investigates the spatial distribution of different types of SNe. They simulate binary populations with the Cambridge `STARS` code, a detailed 1D stellar evolution code (Eggleton, 1971; Pols et al., 1995; Eldridge & Tout, 2004; Eldridge et al., 2008) and explore the fraction of binaries that are disrupted and the ejection velocities of secondary stars. They make predictions for the distances that stars will travel before exploding, but they do not account for the impact of a gravitational potential, instead multiplying the ejection velocity of the star by its remaining lifetime after ejection.

We find a higher fraction of binaries are disrupted, likely as a result of our initial orbital period distribution. E11 find that 70% of binaries are disrupted after the first SN, whilst we find a fraction of 85%. They use an upper initial orbital separation limit of $10^4\,\mathrm{R_\odot}$, while our limit is closer to $10^{5.1}\,\mathrm{R_\odot}$. This means that our simulations contain more weakly bound binaries, which are easier to disrupt and may account for the difference in our results.

Our results for the fraction of SNe at moderate distances are in reasonable agreement, but E11 predict a factor of 4 more SNe at large distances. They find that $f_{D>100\,\mathrm{pc}} = 14\%$, and $f_{D>500\,\mathrm{pc}} = 4\%$ travel more than 500 pc. Their results are most comparable to our $\pi = 0$ variation in which we assume a flat-in-log distribution for the initial orbital period. In this variation, we find $f_{D>100\,\mathrm{pc}} = 11\%$ and $f_{D>500\,\mathrm{pc}} = 1\%$.

The significantly higher distant SN fraction is likely due to the lack of a galactic gravitational potential in their models. In their work, they assume stars travel in a straight line without the influence of a potential, which allows them to travel much further. It is also worth nothing that their lower initial orbital separation limit (which reduces the number of disruptions on average) and higher initial mass limit of $5\,\mathrm{M_\odot}$ compared to our $4\,\mathrm{M_\odot}$ (which translates to shorter lifetimes after ejection) both mean that their distant SN fractions in an equivalent simulation would likely be even higher.

Renzo et al. (2019) (hereafter R19) also investigates massive runaway stars and their spatial distribution using assumptions that are closer to our own, using `binary_c` (Izzard et al., 2004, 2006, 2009, 2018; Izzard & Jermyn, 2023), which as we noted above is based on the same underlying fitting formulae as `COSMIC` (Pols et al., 1998; Hurley et al., 2000, 2002; Breivik et al., 2020b).

We are in strong agreement on the fraction of binaries that are disrupted. R19 finds that $86^{+10}_{-22}\%$ of binaries are disrupted after the first SN. The uncertainties on this estimate come from several population synthesis parameter variations, which include extreme choices for SN natal kicks. Our fiducial finding of 85% is in excellent agreement with this result. We additionally concur that the most impactful parameters for the disruption fraction are the natal kick magnitude and metallicity.

We find a distribution of ejection velocities that is similar to that of R19's fiducial model, with a slightly extended tail to higher velocities. Their distribution peaks at a velocity of $6\,\mathrm{km\,s^{-1}}$, in close agreement to our finding of $7\,\mathrm{km\,s^{-1}}$. We find that 4.3% of secondaries (regardless of mass) are ejected at more than $60\,\mathrm{km\,s^{-1}}$, which corresponds to 0.1% of secondaries with a mass larger than $15\,M_\odot$ at the moment of ejection, in good agreement with the 0.2% reported by R19.

Despite the similarity in our distribution of ejection velocities, we find a significantly lower tail of long distances for runaway star SNe. R19 calculates the *maximum* distance travelled as the product of the ejection velocity and stellar lifetime, neglecting the effect of a galactic potential. They find that on average main-sequence runaway stars travel 126 pc, which is very close to our finding of 120 pc. However, they find that ∼35% of these stars travel more than 100 pc, in contrast to the ∼11% we find in our fiducial model. These comparisons indicate that neglecting a galactic potential may be acceptable for reproducing the bulk properties of the distribution but fails to accurately predict the long distance tails.

8 Conclusions

We present predictions for the impact of binary interactions on the timing and location of core-collapse SNe feedback. We used self-consistent population synthesis and galactic dynamics simulations to trace the time and location of each SN from the recent star formation in the `FIRE` `m11h` galaxy. We compared our results to an equivalent simulation that considers only single stellar evolution. We additionally repeated our simulations for a wide-range of variations in initial conditions, binary physics and galactic settings to demonstrate the robustness of our results. Based on these simulations, we designed an analytic model for core-collapse SN feedback that accounts for binary interactions. We additionally consider how the impact of binary interactions on SN feedback could affect galaxy evolution. Lastly, we compared our findings to earlier works that have previously explored the impact of binary interactions on producing late SNe and runaway stars. Our main findings can be summarised as follows:

1. **Binary interactions can produce late core-collapse SN feedback**
   We predict that ∼25% of core-collapse SN feedback occurs long after a star formation event, beyond the 44 Myr cutoff typically adopted in hydrodynamical simulations. The progenitors of these SNe are primarily the products of stellar mergers.

2. **Supernovae from binaries are displaced from their parent clusters**
   We find that ∼14% of core-collapse feedback occurs at least 100 pc from the progenitor's parent cluster. The progenitors of these SNe are primarily secondary stars that were ejected from their binary after the primary star's SN.

3. **The distributions of times and locations are robust to changes in model parameters**
   We demonstrate that the overall distributions are surprisingly robust to variations across binary physics, initial conditions, and galactic settings. The median of the distribution of SN delay times remains within 7 Myr of the fiducial model across all variations. Similarly the distance median remains within 25 pc of the fiducial model.

4. **An analytic model accurately reproduces joint distributions of SN times and progenitor velocities, allowing binary evolution to be seamlessly included in SNe feedback models in hydrodynamical simulations**
   We develop a metallicity-dependent analytic model for the timing of SNe and the velocities at which their progenitors travel. This physically-motivated model reproduces the timing distribution to within 0.5% and the runaway ejection velocity distribution to within 4% (Figures 7.9–7.10).

5. **The tails of the spatial distribution are sensitive to changes in binary mass transfer**
   The ejection velocities of secondary stars depend on their pre-supernova orbital velocity and thus are sensitive to binary mass transfer variations which alter orbital parameters. In particular, the fraction of feedback beyond 100 pc varies from 10–21% and the fraction beyond 500 pc varies from 0–3% (Figure 7.13).

6. **The fraction of late and displaced SNe increases at low metallicity**
   Decreasing metallicity by a factor of 10 increases $f_{t>44\,\mathrm{Myr}}$ to 34% and $f_{D>100\,\mathrm{pc}}$ to 21%, as well as increasing the overall number of SNe by 17%. These trends are a result of the metallicity dependence of the assumed minimum core mass to reach core-collapse, in addition to the reduced radial expansion at low metallicity (Figure 7.6 and Section 4.5).

Overall, we have shown that core-collapse SN feedback is significantly different for binary star progenitors, yet is robust to many of the uncertainties in binary physics. Our analytic model for feedback enables future simulators to move beyond single star models at a low cost and include the more realistic impact of binaries on galaxy evolution.

We expect that the most substantial differences will be evident in low-metallicity, high-redshift environments, in which the extended time and spatial distributions of feedback from binaries are most prominent. Future investigations of binaries should be particularly directed towards better understanding the initial mass function, orbital period distribution and binary mass transfer physics, since these drive most of the remaining uncertainties in our models for core-collapse SN feedback.

9 Appendices

9.1 Dependence of late SN fraction on metallicity

In Section 7, we highlighted that our results for the trend of the late SN fraction as a function of metallicity seem to disagree with Zapartas et al. (2017) upon immediate inspection. In this Appendix, we demonstrate that this discrepancy is in fact due to a difference in our definitions of "late" SNe and that, when using the same definition, our results are in agreement.

Zapartas et al. (2017) found that a decrease in metallicity also leads to a decrease in the fraction of late SNe, which the authors point out is a result of lower metallicity stars being more compact and therefore interacting later or avoiding interaction entirely. Later interactions (or a lack of interactions) will have a lesser effect on the core mass of the star, which determines its time until core collapse. In contrast, we find the opposite trend, that decreasing metallicity increases $f_{t>44\,\mathrm{Myr}}$. We argue this is primarily a results of lower metallicity decreasing the mass required to reach core collapse, which results in more (slowly evolving) low mass stars being able to reach core collapse.

The difference in our results is a result of our choices of the time beyond which SNe are considered "late". We use the `FIRE`-3 limit of 44 Myr, while Zapartas et al. (2017) uses the time of the final SN from a single star population at each different metallicity.

We apply both limit choices to our simulations in Figure 7.15. Our choice of a fixed limit results in a negative correlation of the late SNe fraction with metallicity, while the variable limit used by Zapartas et al. (2017) gives a positive correlation. This is because the variable limit results in the additional SNe from lower mass stars reaching core collapse no longer being counted towards the "late" SNe. Instead, as Zapartas et al. (2017) explains, the more compact stars at low metallicity have later and fewer interactions, which reduces the number of delayed SNe and hence the late SN fraction. We note that the time of the last single star reaching core collapse is dependent upon the underlying stellar tracks used in population synthesis (Pols et al., 1998).

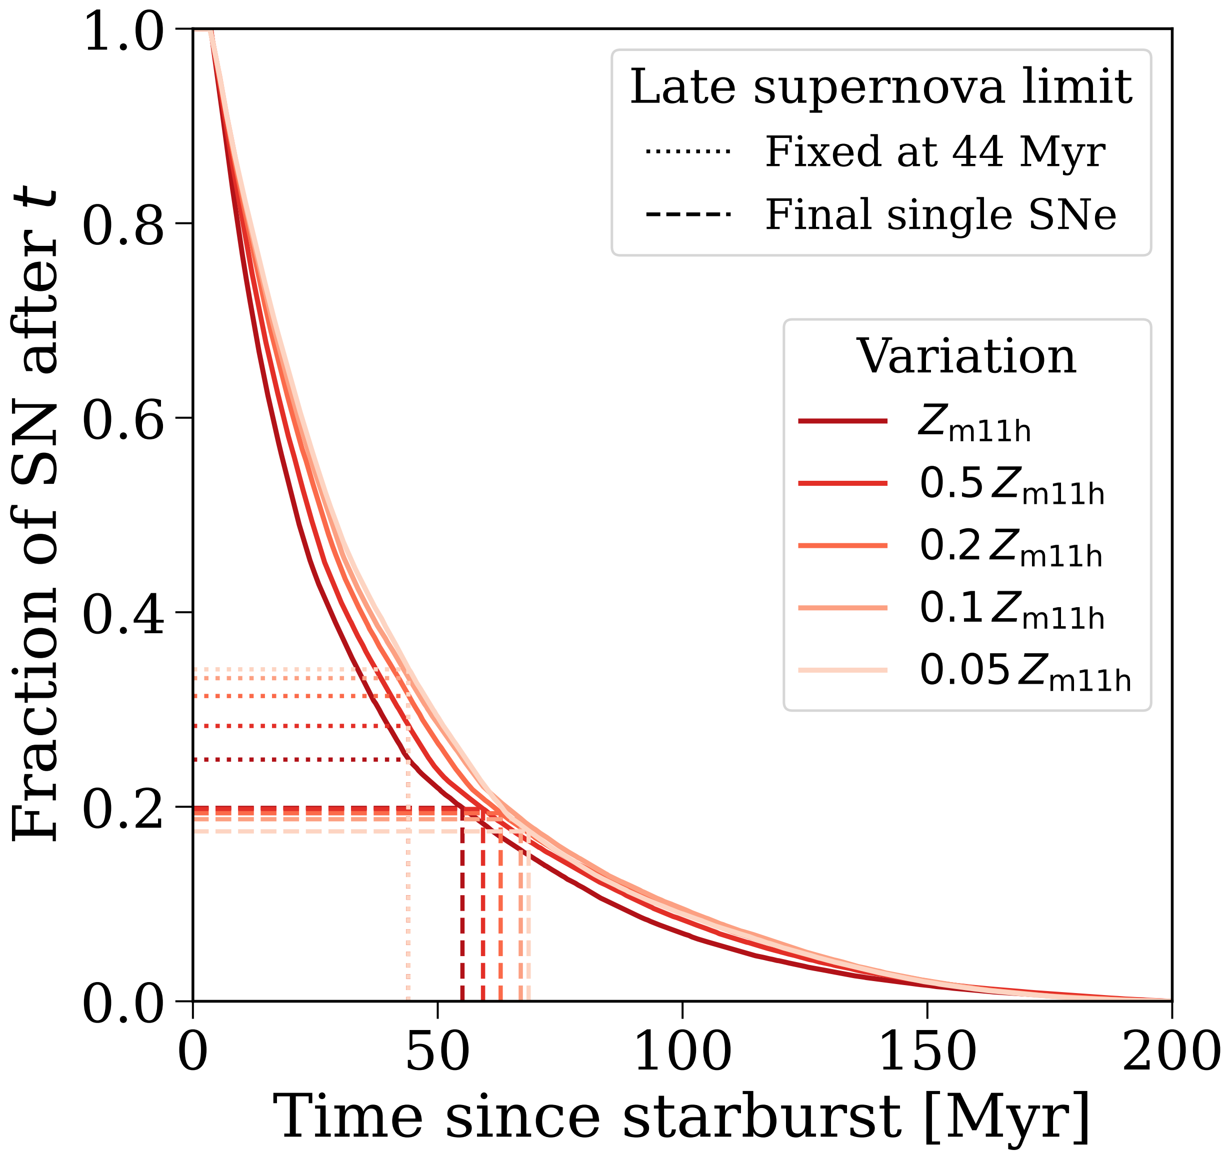

Figure 7.15: The fraction of SNe that occur after a given time for the difference metallicity variations we consider are shown in coloured lines. The dotted lines indicate the fraction of late SNe when using a fixed 44 Myr limit. The dashed lines show the same fraction when using a limit set by the time of the last single star SN. The fixed limit results in a negative correlation of the late SNe fraction with metallicity, while the variable limit gives a positive correlation.

Overall, our results are therefore in agreement with Zapartas et al. (2017) when using the same definition of "late" SN. Our fixed limit corresponds to the limit currently applied in `FIRE`-3 subgrid feedback models and hence is most applicable for the purposes of this work.

9.2 Sampling routine flowchart

This appendix contains Figure 7.16, which illustrates the sampling routine for the analytic model that we present in Section 5.

9.3 SN rate and distribution data

This section contains supplementary data for this work. Table 7.3 list the total number of SNe that occur in each simulation per $100\,\mathrm{M}_{\odot}$, separated by progenitor type, Table 7.2 does the same for the fractions of SNe in the tails of the timing and distance distributions. Table 7.4 & 7.5 details the data for the total distributions and SN times and distances.

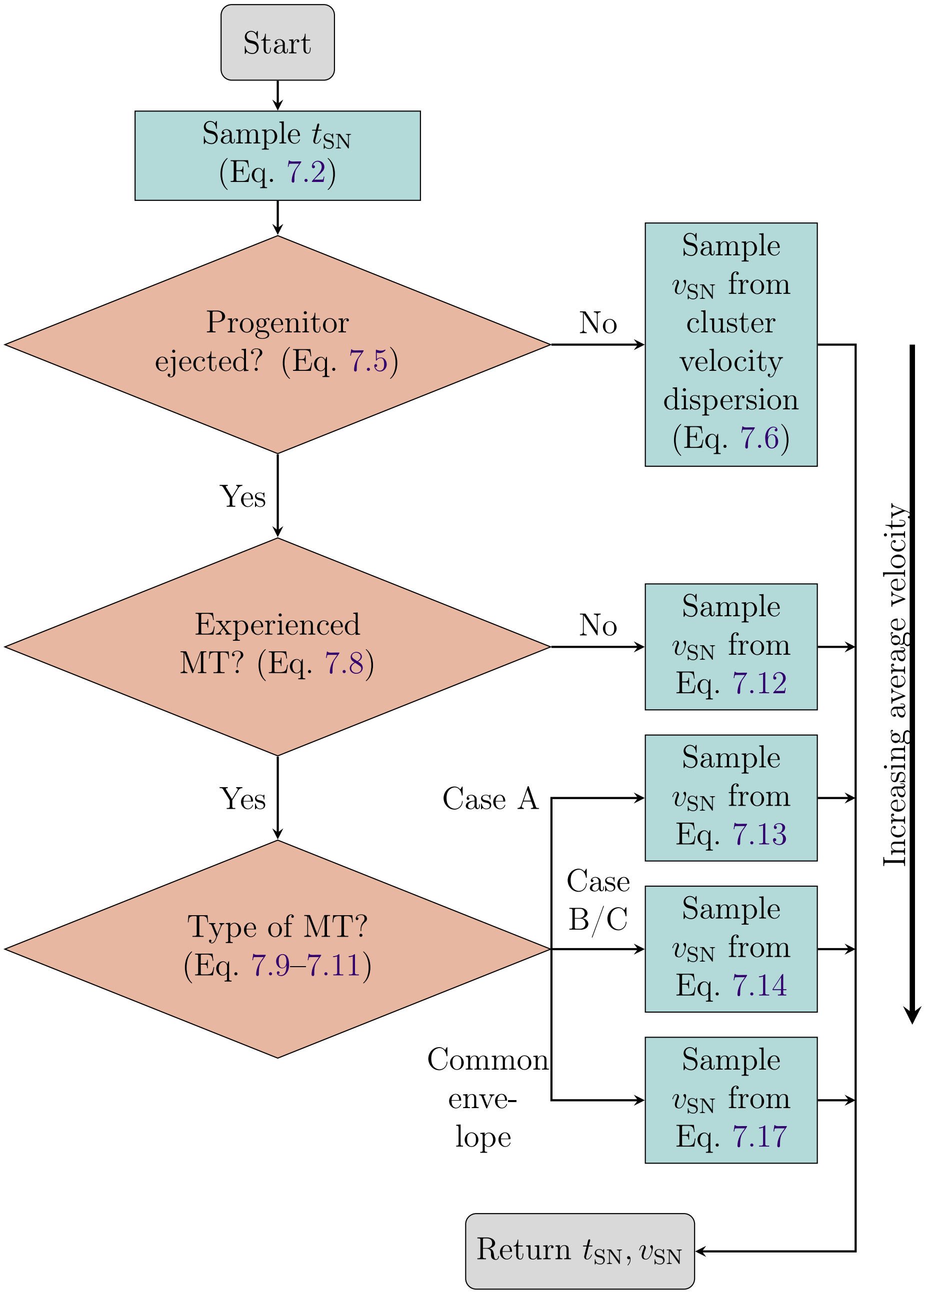

Figure 7.16: A flowchart of the sampling routine for our analytic model of SN feedback, where t_{SN} is the time of the SN and v_{SN} is the velocity at which its progenitor moves away from its parent cluster. Based on the sampled time, we draw a progenitor velocity, v_{SN}. The distribution for this velocity changes based on (a) whether the progenitor was ejected from its binary and, if so, (b) whether it experienced mass transfer before doing so, and, if so, (c) what type of mass transfer it experienced. The velocity distributions are sorted in this flowchart from top to bottom as slowest to fastest (where unejected progenitors proceed slowest on average, whilst stars ejected after a common-envelope are typically fastest).

Model variation	Distribution tails		
	$f_{t>44\,\mathrm{Myr}}$	$f_{D>100\,\mathrm{pc}}$	$f_{D>500\,\mathrm{pc}}$
Fiducial	25.4%	13.2%	0.9%
Binary physics			
Common envelope, $\alpha_{\mathrm{CE}} = 0.1$	25.0%	12.2%	0.5%
Common envelope, $\alpha_{\mathrm{CE}} = 10.0$	13.8%	11.0%	1.0%
Case B Unstable	31.1%	11.6%	0.4%
Case B Stable	30.7%	21.0%	3.1%
Mass transfer efficiency, $\beta = 0.0$	12.3%	13.0%	2.0%
Mass transfer efficiency, $\beta = 0.5$	17.9%	13.7%	1.6%
Mass transfer efficiency, $\beta = 1.0$	26.7%	14.0%	0.6%
Supernova kicks, $\sigma_{\mathrm{CC}} = 20\,\mathrm{km/s}$	28.3%	10.2%	0.1%
Supernova kicks, $\sigma_{\mathrm{low}} = 265\,\mathrm{km/s}$	25.3%	13.5%	1.0%
Supernova kicks, No fallback	25.3%	13.3%	0.9%
Initial conditions			
Singles, $f_{\mathrm{bin}} = 0.0$	1.3%	0.9%	0.0%
Initial mass function slope, $\alpha_{\mathrm{IMF}} = -1.9$	17.9%	11.0%	1.0%
Initial mass function slope, $\alpha_{\mathrm{IMF}} = -2.7$	33.6%	15.5%	0.8%
Orbital period slope, $\pi = 0$	17.5%	11.2%	1.0%
Orbital period slope, $\pi = -1$	32.8%	14.8%	0.8%
Initial upper orbital period limit, $P_{0,\mathrm{max}} = 10^3\,\mathrm{d}$	33.3%	15.3%	1.0%
Mass ratio slope, $\kappa = 1$	27.2%	13.4%	0.7%
Mass ratio slope, $\kappa = -1$	22.3%	11.7%	1.1%
Metallicity			
Metallicity, $\bar{Z} = 0.5\,\bar{Z}_{\mathrm{m11h}}$	29.0%	16.9%	1.0%
Metallicity, $\bar{Z} = 0.2\,\bar{Z}_{\mathrm{m11h}}$	31.7%	19.9%	1.3%
Metallicity, $\bar{Z} = 0.1\,\bar{Z}_{\mathrm{m11h}}$	33.5%	21.0%	1.6%
Metallicity, $\bar{Z} = 0.05\,\bar{Z}_{\mathrm{m11h}}$	34.6%	22.4%	2.4%
Galaxy settings			
Velocity dispersion, $v_{\mathrm{disp}} = 0.5$ km/s	25.4%	5.1%	0.9%
Velocity dispersion, $v_{\mathrm{disp}} = 5$ km/s	25.4%	45.3%	3.0%
ChaNGa (r442)	25.3%	13.9%	0.9%

Table 7.2: The summary statistics for the tails of the timing and distance distributions, for each model variation in our simulations. Column 1 indicates the variation, column 2 is the fraction of SNe that occur after 44 Myr, columns 3 and 4 are the fractions of SNe that occur beyond 100pc and 500pc respectively. Each column gives the value for a different subpopulation and the total number.

Model variation	Supernovae per $100\,\mathrm{M}_\odot$			
	All	P	S	MP
Fiducial	1.18	0.39	0.25	0.52
Binary physics				
Common envelope, $\alpha_{\mathrm{CE}} = 0.1$	1.21	0.35	0.24	0.61
Common envelope, $\alpha_{\mathrm{CE}} = 10.0$	1.02	0.43	0.32	0.25
Case B Unstable	1.13	0.18	0.08	0.85
Case B Stable	1.37	0.64	0.41	0.30
Mass transfer efficiency, $\beta = 0.0$	0.94	0.38	0.23	0.31
Mass transfer efficiency, $\beta = 0.5$	1.06	0.39	0.26	0.39
Mass transfer efficiency, $\beta = 1.0$	1.21	0.39	0.26	0.55
Supernova kicks, $\sigma_{\mathrm{CC}} = 20\,\mathrm{km/s}$	1.04	0.39	0.12	0.51
Supernova kicks, $\sigma_{\mathrm{low}} = 265\,\mathrm{km/s}$	1.19	0.39	0.26	0.52
Supernova kicks, No fallback	1.18	0.39	0.25	0.52
Initial conditions				
Singles, $f_{\mathrm{bin}} = 0.0$	1.07	0.00	0.00	0.00
Initial mass function slope, $\alpha_{\mathrm{IMF}} = -1.9$	1.95	0.71	0.49	0.71
Initial mass function slope, $\alpha_{\mathrm{IMF}} = -2.7$	0.57	0.17	0.10	0.29
Orbital period slope, $\pi = 0$	1.12	0.47	0.28	0.33
Orbital period slope, $\pi = -1$	1.24	0.31	0.21	0.71
Initial upper orbital period limit, $P_{0,\mathrm{max}} = 10^3\,\mathrm{d}$	1.21	0.28	0.22	0.72
Mass ratio slope, $\kappa = 1$	1.20	0.38	0.28	0.51
Mass ratio slope, $\kappa = -1$	1.14	0.37	0.19	0.56
Metallicity				
Metallicity, $\bar{Z} = 0.5\,\bar{Z}_{\mathrm{m11h}}$	1.24	0.42	0.27	0.53
Metallicity, $\bar{Z} = 0.2\,\bar{Z}_{\mathrm{m11h}}$	1.30	0.46	0.29	0.54
Metallicity, $\bar{Z} = 0.1\,\bar{Z}_{\mathrm{m11h}}$	1.36	0.49	0.31	0.56
Metallicity, $\bar{Z} = 0.05\,\bar{Z}_{\mathrm{m11h}}$	1.40	0.52	0.32	0.55
Galaxy settings				
Velocity dispersion, $v_{\mathrm{disp}} = 0.5$ km/s	1.18	0.39	0.25	0.52
Velocity dispersion, $v_{\mathrm{disp}} = 5$ km/s	1.18	0.39	0.25	0.52
ChaNGa (r442)	1.17	0.38	0.25	0.51

Table 7.3: The total numbers of SNe per $100\,\mathrm{M}_\odot$ for different subpopulations, for each model variation in our simulations. Column 1 indicates the variation, column 2-5 are the total SNe per $100\,\mathrm{M}_\odot$ for the total population, primary stars (P), secondary stars (S) and merger products (MP) respectively.

Model variation	Supernova time [Myr]				
	2.5	25	50	75	97.5
Fiducial	4.4	10.8	21.4	44.6	143.5
Binary physics					
Common envelope, $\alpha_{\rm CE} = 0.1$	4.4	11.0	22.2	44.0	139.9
Common envelope, $\alpha_{\rm CE} = 10.0$	4.3	9.5	17.5	31.5	105.6
Case B Unstable	4.3	12.1	27.1	54.2	142.5
Case B Stable	4.5	12.1	25.4	58.0	185.5
Mass transfer efficiency, $\beta = 0.0$	4.2	9.0	17.0	31.1	92.2
Mass transfer efficiency, $\beta = 0.5$	4.3	10.0	19.1	36.3	99.3
Mass transfer efficiency, $\beta = 1.0$	4.4	11.0	22.0	47.1	143.3
Supernova kicks, $\sigma_{\rm CC} = 20\,{\rm km/s}$	4.3	10.7	22.9	50.2	147.7
Supernova kicks, $\sigma_{\rm low} = 265\,{\rm km/s}$	4.4	10.8	21.6	44.4	143.3
Supernova kicks, No fallback	4.4	10.7	21.3	44.5	143.4
Initial conditions					
Singles, $f_{\rm bin} = 0.0$	4.2	8.9	17.3	29.1	43.0
Initial mass function slope, $\alpha_{\rm IMF} = -1.9$	4.0	8.0	15.8	34.2	125.1
Initial mass function slope, $\alpha_{\rm IMF} = -2.7$	5.1	14.2	28.5	60.1	157.4
Orbital period slope, $\pi = 0$	4.3	9.9	18.9	36.1	122.4
Orbital period slope, $\pi = -1$	4.5	11.8	25.1	58.7	157.5
Initial upper orbital period limit, $P_{0,\rm max} = 10^3\,{\rm d}$	4.6	11.7	23.9	58.6	153.0
Mass ratio slope, $\kappa = 1$	4.4	10.8	21.9	48.3	151.4
Mass ratio slope, $\kappa = -1$	4.4	10.7	20.8	40.9	125.1
Metallicity					
Metallicity, $\bar{Z} = 0.5\,\bar{Z}_{\rm m11h}$	4.5	12.2	24.3	49.0	154.1
Metallicity, $\bar{Z} = 0.2\,\bar{Z}_{\rm m11h}$	4.6	13.3	26.7	52.4	148.1
Metallicity, $\bar{Z} = 0.1\,\bar{Z}_{\rm m11h}$	4.6	13.9	28.0	54.9	146.9
Metallicity, $\bar{Z} = 0.05\,\bar{Z}_{\rm m11h}$	4.6	14.3	28.9	55.8	146.3
Galaxy settings					
Velocity dispersion, $v_{\rm disp} = 0.5$ km/s	4.4	10.8	21.4	44.6	143.5
Velocity dispersion, $v_{\rm disp} = 5$ km/s	4.4	10.8	21.4	44.6	143.5
ChaNGa (r442)	4.4	10.8	21.3	44.5	142.4

Table 7.4: The distributions for the times at which SNe occur, for each model variation in our simulations. Each column gives the value for a different percentile, which is listed in the heading.

Model variation	Distance from parent cluster [pc]				
	2.5	25	50	75	97.5
Fiducial	4.6	16.4	34.7	66.8	242.1
Binary physics					
Common envelope, $\alpha_{\mathrm{CE}} = 0.1$	4.6	16.7	34.8	65.3	208.7
Common envelope, $\alpha_{\mathrm{CE}} = 10.0$	4.4	14.7	30.3	58.8	250.5
Case B Unstable	4.4	15.8	33.6	63.8	197.8
Case B Stable	4.9	18.5	40.4	85.1	556.7
Mass transfer efficiency, $\beta = 0.0$	4.1	13.4	28.3	59.2	412.8
Mass transfer efficiency, $\beta = 0.5$	4.4	15.1	32.0	64.6	345.7
Mass transfer efficiency, $\beta = 1.0$	4.6	16.8	35.6	68.9	237.5
Supernova kicks, $\sigma_{\mathrm{CC}} = 20\,\mathrm{km/s}$	4.3	14.6	30.5	59.3	177.4
Supernova kicks, $\sigma_{\mathrm{low}} = 265\,\mathrm{km/s}$	4.6	16.5	34.9	67.4	253.8
Supernova kicks, No fallback	4.6	16.4	34.8	67.0	244.1
Initial conditions					
Singles, $f_{\mathrm{bin}} = 0.0$	3.9	11.8	22.7	40.1	83.7
Initial mass function slope, $\alpha_{\mathrm{IMF}} = -1.9$	3.8	12.5	28.0	58.2	230.0
Initial mass function slope, $\alpha_{\mathrm{IMF}} = -2.7$	5.5	20.4	40.6	74.6	249.0
Orbital period slope, $\pi = 0$	4.3	14.9	31.2	60.5	237.4
Orbital period slope, $\pi = -1$	4.8	17.9	37.9	72.1	237.7
Initial upper orbital period limit, $P_{0,\mathrm{max}} = 10^3\,\mathrm{d}$	4.9	18.4	38.7	73.6	250.7
Mass ratio slope, $\kappa = 1$	4.6	16.7	35.3	67.8	229.9
Mass ratio slope, $\kappa = -1$	4.5	15.6	32.5	62.7	241.5
Metallicity					
Metallicity, $\bar{Z} = 0.5\,\bar{Z}_{\mathrm{m11h}}$	4.7	17.9	38.9	76.2	289.9
Metallicity, $\bar{Z} = 0.2\,\bar{Z}_{\mathrm{m11h}}$	4.9	18.9	41.8	83.9	352.7
Metallicity, $\bar{Z} = 0.1\,\bar{Z}_{\mathrm{m11h}}$	4.9	19.3	43.0	87.2	385.7
Metallicity, $\bar{Z} = 0.05\,\bar{Z}_{\mathrm{m11h}}$	4.9	19.5	43.6	91.1	492.3
Galaxy settings					
Velocity dispersion, $v_{\mathrm{disp}} = 0.5$ km/s	1.9	5.8	12.1	26.2	181.2
Velocity dispersion, $v_{\mathrm{disp}} = 5$ km/s	12.5	43.7	89.0	169.4	537.4
ChaNGa (r442)	4.6	16.7	35.2	68.1	251.6

Table 7.5: The distributions for the distances at which they occur relative to their parent cluster, for each model variation in our simulations. Each column gives the value for a different percentile, which is listed in the heading.

CHAPTER 8

FUTURE DIRECTIONS AND APPLICATIONS

The most important step a man can take. It's not the first one, is it? It's the next one. Always the next step.

– Dalinar Kholin, in *Oathbringer*

1 Overview

In this Chapter we ponder the potential future applications of the research in this thesis. Much of the work presented thus far in this thesis has been theoretical in nature, focusing on making predictions and simulations. With the upcoming release of several large-scale observed populations from surveys such as *Gaia* and the *Vera C. Rubin Observatory*, in addition to the trove of archival imaging from *HST*, there is a wealth of observational data available to constrain `cogsworth` models. These datasets could provide critical constraints on several uncertain aspects of binary physics. With this in mind, in the sections below we consider some potential directions in which this research could be taken, with a focus on directly comparing simulated results to observations. We sort these approximately chronologically based on how soon they could be accomplished based on the availability of data and further development necessary.

2 Spatial correlations of massive stars and the ISM

Observationally there is an expanding wealth of extraordinarily high resolution data on cool and cold gas in nearby galaxies, with increasingly complete catalogues of stellar tracers of high mass stars (Neugent & Massey, 2011; Leroy et al., 2021; Lamb et al., 2016). Massive stars, the majority of which are found in binaries, are tracers of the high mass end of galaxy formation and play a critical role in galaxy feedback processes (e.g. Dekel & Silk, 1986; Hopkins et al., 2012; Somerville & Davé, 2015). We are now at a point at which we have complete censuses of astrophysical populations within tens of parsecs of star-forming gas in the Local Group, with data on spatial and velocity distributions of both stars and gas (Neugent & Massey, 2023; Massey et al., 2021; White et al., 2019). This presents the community with an opportunity to confront binary models and models of stars formation with increasingly constraining catalogue of ground-based data from observatories, such as *ALMA*, *GBT* and the *VLA* (Sarbadhicary et al., 2023). With the development of `cogsworth`, we now have the infrastructure necessary to perform end-to-end tests of the relationship between massive stars, gas and binary formation channels.

The key observational constraints provided by this data are the relative spatial distributions of gas and massive stars in different phases of binary evolution, as well as their relative numbers and kinematics. Each of these quantities is extremely sensitive to the exact parameters controlling binary evolution. For example, if supernova natal kicks are typically strong then one would find massive stars systematically far from their birth sites (depending on the mass and age of the binary) since even the most tight binaries can be disrupted and transfer the high orbital velocity into the runaway velocity of the newly unbound massive star. Conversely, if supernova natal kicks are weak, we would see highly clustered massive stars, subject only to standard dynamical dissolution of star forming regions. Alternatively, if some massive star formation takes place at very low column densities then one could observe a population of truly isolated massive stars, but this population would have different properties than kicked populations, such as kinematics and age. Simultaneously, the same supernova kick assumptions and star formation distributions that match these quantities must *also* produce results that remain consistent with observations from the *LIGO-Virgo-*

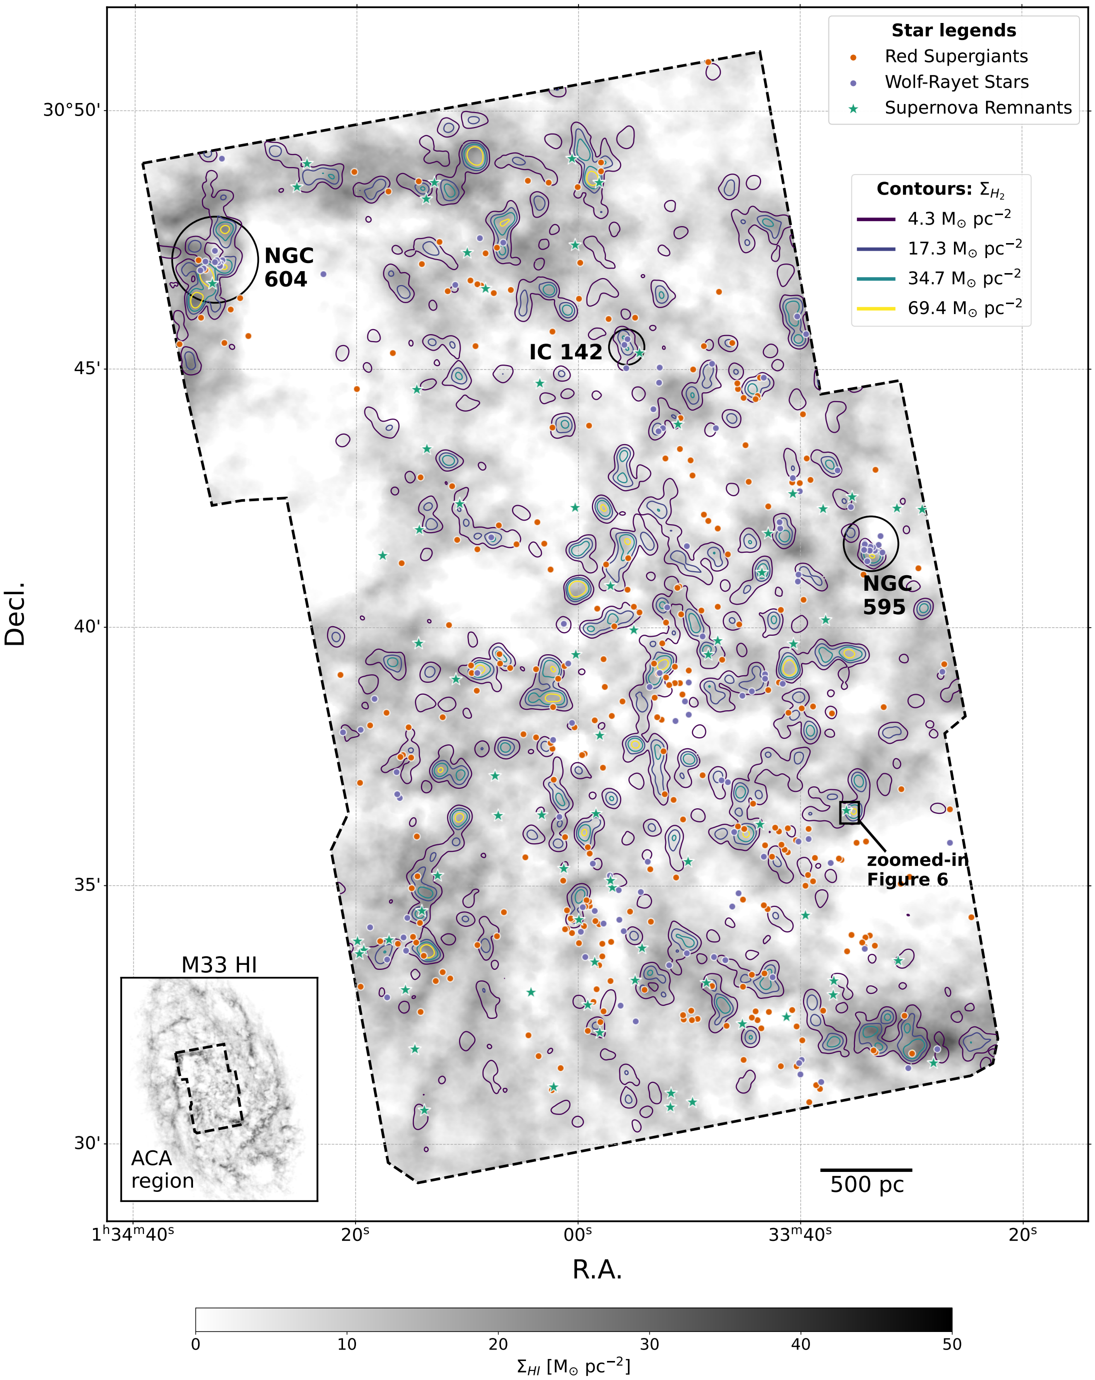

Figure 8.1: Observations of M33 reveal correlations between the cold ISM and massive stars at different evolutionary stages (Figure 1 from Sarbadhicary et al., 2023). Coloured points show massive stellar tracers (red supergiants, Wolf-Rayet stars and supernova remnants), while the density map and contours show HI and H2 gas in the M33 ACA survey region. Note that Wolf-Rayet stars show strong correlations, whereas supernova remnants are more uniformly spaced.

KAGRA collaboration at the endpoints of these channels, the census of electromagnetic transients and supernova remnants.

Distinguishing among these different scenarios using the rich library of information that now exists is now possible with the formalism that `cogsworth` offers. `cogsworth`'s ability to trace cluster dissolution, orbital evolution, and binary stellar evolution simultaneously allows the quantitative interpretation of observers' detectable signals (e.g., Sarbadhicary et al., 2023) in the context of a physical model.

In future work, we will make predictions for the correlations between massive stars and the ISM using `cogsworth` simulations that are linked directly to public hydrodynamical zoom-in simulations. These simulations will analyse the dependence of these correlations on aspects of binary physics, the boundedness of clusters and the galaxy model. Thus we will provide the underlying theoretical models necessary to qualitatively interpret the wealth of recent observations (e.g. Mayker Chen et al., 2023; Sarbadhicary et al., 2023). Figure 8.1 shows recent observations of M33 which detail the spatial correlations between the ISM and massive stars, against which `cogsworth` simulations could be compared.

Specifically, we will use the `FIRE`-2 simulations (Hopkins et al., 2018b) and simulations from `ChaNGa`, such as the MARVEL-ous dwarfs and DC Justice League simulations (Applebaum et al., 2021; Christensen et al., 2023) to set the initial conditions of the galaxy and define the galactic potential. These simulations track star particles (each representing $\sim$10000 $\mathrm{M}_{\odot}$) from which we can draw a population of binaries and single stars using `cogsworth`, with masses, metallicities and kinematics that are consistent with the hydrodynamical simulations. Additionally, we will transform the distribution of dark matter, star and gas particles in these simulations into galactic potentials using basis function expansions, following the self consistent field method (Hernquist & Ostriker, 1992; Lowing et al., 2011).

We will run a grid of models varying binary physics, cluster boundedness and hydrodynamical galaxy models. For binary physics we will explore variations in supernova natal kicks, common envelope efficiency and mass transfer efficiency, since these parameters have been shown to have a strong impact on the kinematics of binary systems (Renzo et al., 2019). We will vary the virial parameter of sampled clusters in order to assess how changing the boundedness of each star cluster affects the large scale spatial distribution (Bertoldi & McKee, 1992). Additionally we will estimate the dependence of our results on the underlying galaxy model by using different models from both `FIRE` and `ChaNGa`.

For each of these simulations we will make predictions for the spatial correlation between massive stars and the ISM (tracked using gas particles from the hydrodynamical zoom-in simulations) with 2-point correlation functions and explore the dependence of these correlations on the various parameters we vary in our grid of models. These predictions will be essential for future work to qualitative interpret and understand observed spatial correlations between massive stars and the ISM in the Milky Way and nearby galaxies.

3 Runaway star kinematics

3.1 Eliminating interloping runaway stars by tracing their orbits

Runaway stars are stars moving at high velocities within the galaxy, potentially as a result of ejection from a binary after the primary star goes supernova and disrupts the orbit (e.g., Renzo et al., 2019). The positions and kinematics of these stars are therefore imprinted with information about their pre-supernova orbits and thus could offer constraints on binary evolution.

The current method for identifying runaway stars is to measure whether their velocity relative to the local standard of rest (LSR) at their current position is above a particular threshold (e.g., Carretero-Castrillo et al., 2023; Guo et al., 2024). However, in some preliminary investigations we have recently shown that this method can lead to significant overestimation of the observed runaway population, since a runaway star's velocity will change over time, as will its location and thus the LSR (Wagg et al. (in prep.)). This means that the current population of observed runaway stars, and hence the constraints on binary stellar evolution they provide, are unreliable. In anticipation of the plethora of new observations in the *Gaia* DR4 data release, we will develop a robust method for the identification of runaway stars.

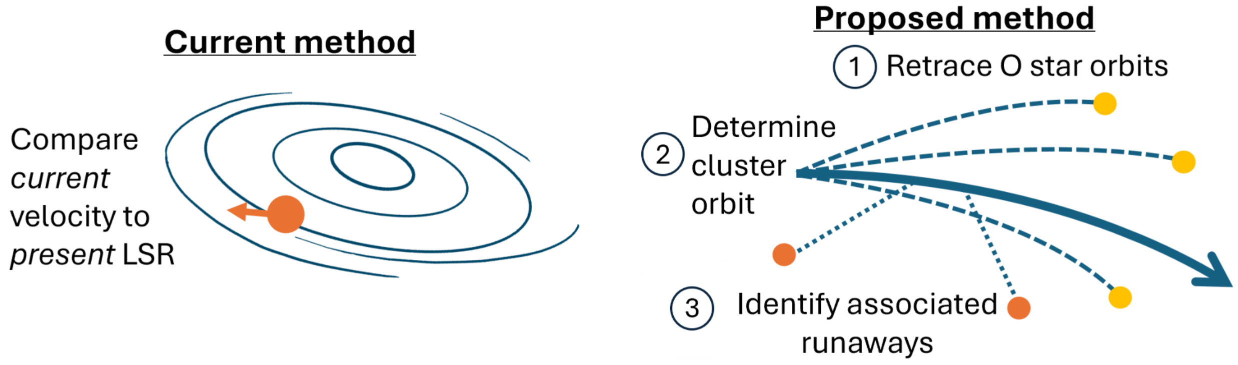

Figure 8.2: The current method for identifying observed runaways is inconsistent with theoretical definitions due to changes in position and velocity. A potential future direction could be implementing a new method that traces candidate runaway stars to their parent clusters, allowing one to calculate their original ejection velocities.

Massive stars have relatively short lives ($\lesssim 50\,\mathrm{Myr}$) and so their galactic orbits are usually unperturbed. Therefore, we can use `cogsworth` to trace back the full orbits of massive O-type stars detected in *Gaia*, finding points of intersection of multiple stars which represent the birth location of their parent star clusters. We will then compare the orbits of candidate runaway stars to these parent clusters, determining the time at which the candidates would have been ejected. Based on this time and the cluster trajectory, we will calculate the original ejection velocities of candidates, which we will apply as a more reliable selection threshold than present-day velocities (Figure 8.2). We will test the efficacy and robustness of this method from the perspective of an observer using `cogsworth` simulations of clustered star formation in the Milky Way. We designed `cogsworth` to enable the application of galactic dust maps (to account for dust extinction) (Green et al., 2019; Green, 2018), MIST bolometric correction functions (to convert to specific filters; Dotter, 2016), and the *Gaia* empirical selection function (Cantat-Gaudin et al., 2023), such that we are ready to make direct comparisons between my simulations and *Gaia* observations the day that the data is released and uncover a new, purer population of runaway stars.

3.2 Constraints on binary evolution from runaway stars

The ejection velocity of a runaway star is directly determined by its binary orbit prior to its ejection, which is strongly impacted by mass transfer during binary evolution (see Chapter 6). Additionally, the strength of the supernova kick will determine whether a binary disrupts and produces a runaway star (e.g., Blaauw, 1961). Thus, I will be able to use the kinematics of the newly discovered *Gaia* runaway population from project in Section 3.1 to place critical constraints on binary mass transfer and supernova natal kicks. Moreover, the VFTS sample of runaway stars in the 30 Doradus region of the LMC (Sana et al., 2022) provides the perfect laboratory for testing models of binary stellar evolution in a more controlled environment. This is because its environment, such as its metallicity and star formation history, is well constrained (Schneider et al., 2018a,b), which will allow me to more clearly infer the impact of binary interactions. The smaller VFTS sample is an ideal control to better characterise the large population of *Gaia* Milky Way runaway stars. An additional consideration is that observed runaways consist not only of those formed via binary evolution, but also via dynamical encounters within stellar clusters, which will not be affected by changes in binary physics (Poveda et al., 1967). The exact fraction of runaways formed through either channel is still debated (Hoogerwerf et al., 2001; Sana et al., 2022; Carretero-Castrillo et al., 2023), but comparisons of simulations and observations could prove key to constraining it.

We will perform `cogsworth` simulations of runaway stars in 30 Doradus region of the LMC and in the Milky Way. We will enhance `cogsworth` to use N-body dynamics for the initial evolution of star clusters, thus accounting for runaways formed via dynamical encounters. For this suite of simulations, we will vary the strength of supernova explosions and the efficiency of mass transfer. Using these simulations, we will predict the positions and kinematics of runaway stars, as well as the fraction of runaway stars formed via binary evolution. We will quantify the sensitivity of these predictions to the parameters we vary and, by comparing the results to observations, place additional constraints on our understanding of binary mass transfer and supernova explosions.

4 Time-evolving galactic potentials and the offset of short gamma-ray bursts from host galaxies

For many sources such as massive stars, it is reasonable to assume that the galactic potential is static on the timescale of their relatively short evolution. This assumption is also often applied to long-lived populations, such as short-gamma ray burst (sGRB) progenitors (e.g., Wiggins et al., 2018). Yet it is clear that traditional models using static galactic potentials are not capable of describing the dynamically complex evolutionary history of galaxies, and can lead to misleading results (e.g., Arora et al., 2022). This is particularly relevant for

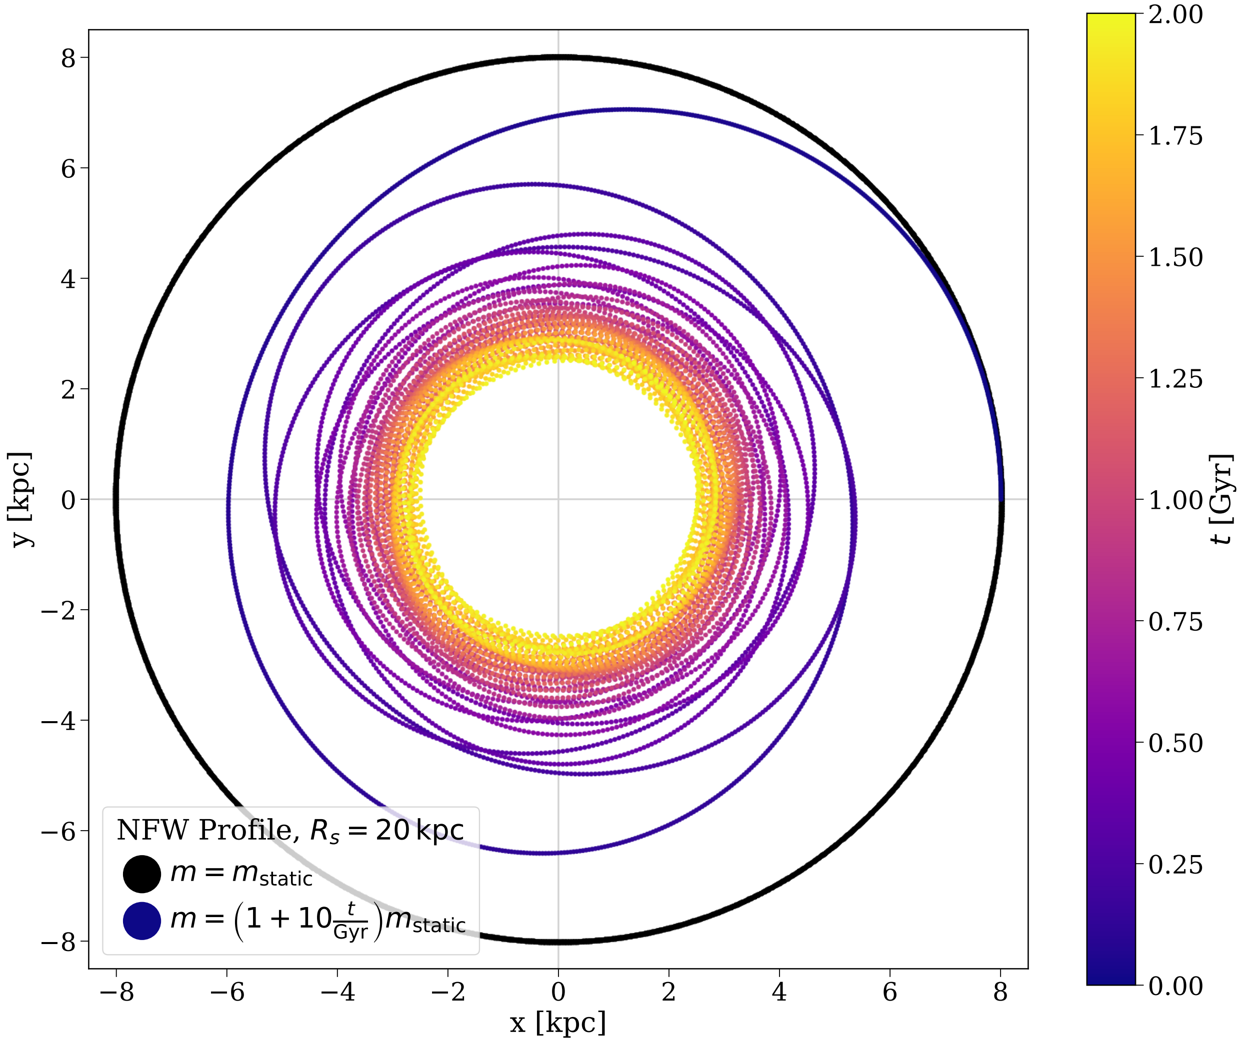

Figure 8.3: A preliminary test of a time-evolving galactic potential, compared to a static potential. Scatter points show the orbit of a test particle evolved for 2 Gyr in two different potentials, starting 8 kpc from the galactic centre with the circular velocity at that location. Black points correspond to a static Navarro et al. (1997) profile, with a scale radius of 20 kpc and a mass parameter of $m_{\mathrm{static}} = 1.77 \times 10^{12}\,\mathrm{M}_{\odot}$. The coloured scatter points show the orbit for a potential with an evolving mass parameter that grows the mass of the galaxy rapidly, such that $m = (1 + 10\,t/\,\mathrm{Gyr})m_{\mathrm{static}}$.

older tracers of binary endpoints such as double compact objects, which often have delay times of billions of years. These long-lived sources can provide critical, and complementary, constraints on supernova natal kicks compared to those provided by massive stars. It is therefore essential that we account for time-evolving potentials in our models.

At present `cogsworth` can only account for a static gravitational potential. We will develop an option for a time-evolving galactic potential that accounts for the mass growth of the galaxy. `cogsworth` is already equipped to postprocess a hydrodynamical simulation snapshots as input for the star formation history and potential of a galaxy. We will extend this functionality to consider *multiple* snapshots from a hydrodynamical simulation. In each case we will use the self-consistent field method (Hernquist & Ostriker, 1992) to model the galactic potential by fitting the galactic mass distribution using a basis function expansion in spherical harmonics.

As a proof of principle, we have tested this functionality from the upcoming `galax` code. In Figure 8.3, we show that `galax` can effectively trace the orbits of particles through a time-evolving galactic potential. The computation of these orbits are very similar in runtime despite the added complexity. We evolved a particle through a static Navarro et al. (1997) profile, with a scale radius of 20 kpc and a mass parameter of $m_{\rm static} = 1.77 \times 10^{12}\,{\rm M_\odot}$, which is shown with black scatter points. This orbit remains circular for the duration of the integration (2 Gyr). In contrast, the coloured points show the orbit for a potential with an evolving mass parameter that grows the mass of the galaxy rapidly, such that $m = (1 + 10\,t/\,{\rm Gyr}) m_{\rm static}$. As such, the particle is drawn further towards the centre of the galaxy as the simulation progresses. We will build upon this preliminary stage to fully connect this functionality to a `cogsworth` model with potentials fit by basis function expansions.

We will use this feature to predict the offset of binary neutron star (BNS) mergers, the progenitors of sGRBs, from their host galaxies and compare to prior work that assumes a static galactic potential. Over the relevant durations, galactic potentials may vary by tens of percent, which impacts escape speeds relevant for binary neutron stars (e.g., Diemer et al., 2013). Therefore the offset distributions and hostless fraction of sGRBs will likely be strongly impacted. Moreover, with `cogsworth` we can explore the full range of binary physics variations beyond the smaller subset considered in the literature currently (e.g., Behroozi et al., 2014), specifically focusing on supernova natal kicks and common-envelope efficiency, constraining these parameters with observations. This is also timely given the wealth of sGRB data available for comparisons to our models (Fong et al., 2022), particularly new sources associated with faint dwarf galaxies (Nugent et al., 2024).

5 Sky localisations of *LISA* gravitational-wave sources

As we discussed in Chapters 2 & 3, *LISA* is expected to detect tens of thousands of galactic gravitational wave sources (Nelemans et al., 2001; Korol et al., 2017), of which hundreds will be double compact objects (DCOs) involving a neutron star or black hole (Lamberts et al., 2018; Wagg et al., 2022c). These more massive DCOs are an exciting population of interest for detection with *LISA* given that there is no detection of a galactic binary black hole (BHBH) or black hole neutron star binary (BHNS) to date. This population is complementary to the one to which LIGO and other ground based detectors are sensitive.

Firstly, there is no bias towards higher masses, meaning that we will be able to gain insights about black hole mass function and potentially probe the observed lower mass gap (Shao & Li, 2021; Wagg et al., 2022c). Moreover, *LISA* will detect these binaries during the inspiral phase and will be able to measure the eccentricities of a large fraction of the observed population, unlike ground-based detections where the vast majority have circularised. This provides an additional probe into binary evolution.

It is therefore essential that we accurately model the predicted population of *LISA* detectable sources such that we can better constrain our models. However, to date no predictions for massive DCOs detectable with *LISA* (including those in chapters 2 & 3) have accounted for the effect of a galactic potential or supernova kicks (e.g. Nelemans et al., 2001; Lamberts et al., 2018; Lau et al., 2020; Wagg et al., 2022c). It has been argued that this effect may be neglected since *LISA* is not strongly distance-limited for this population of sources (e.g. Lau et al., 2020). Yet we expect that a sizeable fraction of compact objects may be displaced around the galaxy or even escape entirely, and this may impact the detection rate (Sweeney et al., 2022).

Moreover, as we discussed in Chapter 3, distinguishing more massive DCOs from the prominent double white dwarf (WDWD) foreground will prove difficult due to uncertain chirp mass measurements, with around 60% of the population remaining indistinguishable when considering mass alone (Wagg et al., 2022c). We seek to increase the sample of galactic BHBH, BHNS and NSNSs by expanding the population of massive DCOs in LISA that distinguishable from WDWDs. Supernova natal kicks are likely to eject binaries from the galactic plane, whereas finding a circular WDWD out of the galactic plane is highly unlikely. For this reason, localising a *LISA* source away from the galactic plane with any measurable eccentricity rules it out as a WDWD. We will therefore investigate the potential for using sky localisation as an additional discriminator between more massive DCOs and WDWDs. We expect that this will significantly increase the fraction of systems that are distinguishable from the WDWD foreground, and thus the sample of galactic black hole and neutron star binaries.

6 Outlook

Overall, the work in this thesis has enabled several new research directions, particularly in terms of using the positions and kinematics of massive binary products to constrain aspects of binary evolution. The tools that we have developed are broadly applicable, capable of exploring several areas in binary stellar evolution, galactic dynamics and potentially even galaxy evolution. As such, I hope that the most fascinating future direction to come from this thesis is one that I have not yet envisioned!

CHAPTER 9

CONCLUSION & SUMMARY

> *There, peeping among the cloud-wrack above a dark tor high up in the mountains, Sam saw a white star twinkle for a while. The beauty of it smote his heart, as he looked up out of the forsaken land, and hope returned to him. For like a shaft, clear and cold, the thought pierced him that in the end the Shadow was only a small and passing thing: there was light and high beauty for ever beyond its reach.*
>
> – J. R. R. Tolkein, *The Return of the King*

The purpose of this thesis has been to outline our efforts in exploring the impact of binary interactions on a variety of massive stellar populations, both through rapid population synthesis and detailed 1D stellar evolution models. In particular, we consider the sensitivity of these results to major uncertainties in binary physics, with a view towards placing constraints on binary physics parameters. The culmination of this work is not only a series of scientific results and predictions, but also two new open-source codes (`LEGWORK` and `cogsworth`) that enable future community-driven investigations into these matters.

We started by considering the population of gravitational-wave sources that will be detectable by the future spaced-based detector *LISA* (Chapters 2–3). Our new open-source code, `LEGWORK`, provides the community with a reliable resource for calculating the detectability of *LISA* sources and computing their evolution due to gravitational-wave emission. We demonstrated a series of potential use cases of the code, including investigating the role of eccentricity on detectability and comparing different sensitivity curves. We synthesised a population of binaries and used `LEGWORK` to explore how the rates and demographics of the galactic population of black hole and neutron star binaries are sensitive to different aspects of binary physics. We predict that *LISA* will detect between 30 and 370 BH and NS binaries, and the rates of each population are sensitive to uncertainties in common-envelope evolution, supernova natal kicks and Wolf-Rayet winds. We additionally showed that extending the *LISA* mission to 10 years would not only increase the total number of detections, but also the fraction of that population that can be distinguished from WDWD binaries.

From here, we turned out attention from gravitational-waves to gravity-mode pulsations and the asteroseismology of massive stars that have accreted material from a companion (Chapter 4). With our proof-of-principle analysis, we established for the first time that mass transfer can leave a significant imprint on the asteroseismic signals of accretor stars. We showed that the rejuvenation of the convective core as a result of mass transfer leaves an imprint on the chemical composition gradient in the star. This change in gradient influences all pulsations that are sensitive to that region of the star, shifting their pulsation periods. We showed that typical asteroseismic methods for estimating the mass and age of the star (which assume single star evolution) can be significantly erroneous for accretor stars. In particular, the age of the accretor we evolved would be underestimated by more than 70 million years, even for a correctly estimated mass and central hydrogen composition. The exact impact of these effects are undoubtedly sensitive to the timing of mass transfer and the amount of material that is accreted. Therefore, future work considering a full grid of binaries could model this effect and use asteroseismic measurements to better understand binary mass transfer.

In the subsequent chapters, we developed and applied new techniques for predicting the *extrinsic* positions and kinematics of massive binaries, in addition to their intrinsic properties such as rates, masses and orbital periods. We presented a new open-source code, `cogsworth`, which can be used to perform self-consistent population synthesis and galactic dynamics simulation (Chapter 5). We used `cogsworth` to demonstrate how the positions of massive stellar populations can inform our understanding of binary interactions. We highlighted how `cogsworth` simulations can be initiated from a hydrodynamical simulation, improving the initial spatial correlations of stars and the detail of the galactic potential used. We additionally examined how ejection velocities of stars are relatively insensitive to the strength of supernova natal kicks and compared the prescriptions for these ejections across three different population synthesis codes (Chapter 6).

Finally, we applied `cogsworth` to explore how binary interactions impact the timing and location of core-collapse supernovae (Chapter 7). Current models for supernova feedback in hydrodynamical simulations assume that all massive stars are formed as single stars. We showed that binary interactions result in both late-time and spatially-displaced supernovae. We found that more than a quarter of supernovae occur after the time of the final single star supernova, while 13% explode more than 100pc from their parent star cluster. We assessed the robustness of these predictions to a plethora of variations in binary physics, initial conditions and galaxy parameters and found these results are surprisingly insensitive to uncertainties in the parameters we considered. Given that these supernovae could have a significantly different impact on galactic feedback, we developed a new analytic model for core-collapse supernova feedback. This model includes physically-motivated metallicity-dependent transitions and reproduces the timing and velocity distributions of supernovae to within 1% and 4% respectively. Our model can be used in future hydrodynamical simulations to better account for binary evolution. This may be particularly relevant for high redshift galaxies in which the spatial extent of the galaxy is reduced, while the low-metallicity environment produces even later and more distributed supernovae.

Overall, it is our hope that the findings in this thesis, as well as the new open-source codes that we released, will drive forward the field of massive binary evolution. As we discussed in Chapter 8, there are several observational datasets that will be ideal for constraining binary evolution parameters from population synthesis simulations such as those we presented in this thesis. We look forward with excitement to the coming decades of massive stellar evolution research as our collective understanding of binary interactions improves!

BIBLIOGRAPHY

Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016, *Binary Black Hole Mergers in the First Advanced LIGO Observing Run*, Physical Review X, 6, 041015, doi: `10.1103/PhysRevX.6.041015`

—. 2019, *GWTC-1: A Gravitational-Wave Transient Catalog of Compact Binary Mergers Observed by LIGO and Virgo during the First and Second Observing Runs*, Physical Review X, 9, 031040, doi: `10.1103/PhysRevX.9.031040`

—. 2020a, *GW190425: Observation of a Compact Binary Coalescence with Total Mass* $\sim 3.4\ M_{\odot}$, ApJ, 892, L3, doi: `10.3847/2041-8213/ab75f5`

—. 2017a, *GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral*, Phys. Rev. Lett., 119, 161101, doi: `10.1103/PhysRevLett.119.161101`

—. 2017b, *On the Progenitor of Binary Neutron Star Merger GW170817*, ApJ, 850, L40, doi: `10.3847/2041-8213/aa93fc`

Abbott, R., Abbott, T. D., Abraham, S., et al. 2020b, *GW190814: Gravitational Waves from the Coalescence of a 23 Solar Mass Black Hole with a 2.6 Solar Mass Compact Object*, ApJ, 896, L44, doi: `10.3847/2041-8213/ab960f`

—. 2021a, *GWTC-2: Compact Binary Coalescences Observed by LIGO and Virgo during the First Half of the Third Observing Run*, Physical Review X, 11, 021053, doi: `10.1103/PhysRevX.11.021053`

—. 2021b, *Population Properties of Compact Objects from the Second LIGO-Virgo Gravitational-Wave Transient Catalog*, ApJ, 913, L7, doi: `10.3847/2041-8213/abe949`

—. 2021c, *Observation of Gravitational Waves from Two Neutron Star-Black Hole Coalescences*, ApJ, 915, L5, doi: `10.3847/2041-8213/ac082e`

Abbott, R., Abbott, T. D., Acernese, F., et al. 2023, *Population of Merging Compact Binaries Inferred Using Gravitational Waves through GWTC-3*, Physical Review X, 13, 011048, doi: `10.1103/PhysRevX.13.011048`

—. 2024, *GWTC-2.1: Deep extended catalog of compact binary coalescences observed by LIGO and Virgo during the first half of the third observing run*, Phys. Rev. D, 109, 022001, doi: `10.1103/PhysRevD.109.022001`

Abdul-Masih, M., Banyard, G., Bodensteiner, J., et al. 2020, *On the signature of a 70-solar-mass black hole in LB-1*, Nature, 580, E11, doi: `10.1038/s41586-020-2216-x`

Abt, H. A. 1983, *Normal and abnormal binary frequencies.*, ARA&A, 21, 343, doi: `10.1146/annurev.aa.21.090183.002015`

Aerts, C., Christensen-Dalsgaard, J., & Kurtz, D. W. 2010, Asteroseismology (Springer), doi: `10.1007/978-1-4020-5803-5`

Aerts, C., Mathis, S., & Rogers, T. M. 2019, *Angular Momentum Transport in Stellar Interiors*, ARA&A, 57, 35, doi: `10.1146/annurev-astro-091918-104359`

Aerts, C., & Tkachenko, A. 2024, *Asteroseismic modelling of fast rotators and its opportunities for astrophysics*, A&A, 692, R1, doi: `10.1051/0004-6361/202348575`

Aghakhanloo, M., Murphy, J. W., Smith, N., & Hložek, R. 2017, *Modelling luminous-blue-variable isolation*, MNRAS, 472, 591, doi: `10.1093/mnras/stx2050`

Agrawal, P., Hurley, J., Stevenson, S., et al. 2023, *Modelling stellar evolution in mass-transferring binaries and gravitational-wave progenitors with METISSE*, MNRAS, 525, 933, doi: `10.1093/mnras/stad2334`

Agrawal, P., Hurley, J., Stevenson, S., Szécsi, D., & Flynn, C. 2020, *The fates of massive stars: exploring uncertainties in stellar evolution with METISSE*, MNRAS, 497, 4549, doi: `10.1093/mnras/staa2264`

Almeida, L. A., Sana, H., Taylor, W., et al. 2017, *The Tarantula Massive Binary Monitoring. I. Observational campaign and OB-type spectroscopic binaries*, A&A, 598, A84, doi: `10.1051/0004-6361/201629844`

Amaro-Seoane, P., Audley, H., Babak, S., et al. 2017, *Laser Interferometer Space Antenna*, arXiv e-prints, arXiv:1702.00786, doi: `10.48550/arXiv.1702.00786`

Andersson, E. P., Agertz, O., & Renaud, F. 2020, *How runaway stars boost galactic outflows*, MNRAS, 494, 3328, doi: `10.1093/mnras/staa889`

Andersson, E. P., Agertz, O., Renaud, F., & Teyssier, R. 2023, *INFERNO: Galactic winds in dwarf galaxies with star-by-star simulations including runaway stars*, MNRAS, 521, 2196, doi: `10.1093/mnras/stad692`

Andrews, J. J., Breivik, K., Pankow, C., D'Orazio, D. J., & Safarzadeh, M. 2020, *LISA and the Existence of a Fast-merging Double Neutron Star Formation Channel*, ApJ, 892, L9, doi: `10.3847/2041-8213/ab5b9a`

Andrews, J. J., & Kalogera, V. 2022, *Constraining Black Hole Natal Kicks with Astrometric Microlensing*, ApJ, 930, 159, doi: `10.3847/1538-4357/ac66d6`

Andrews, J. J., Bavera, S. S., Briel, M., et al. 2024, *POSYDON Version 2: Population Synthesis with Detailed Binary-Evolution Simulations across a Cosmological Range of Metallicities*, arXiv e-prints, arXiv:2411.02376, doi: `10.48550/arXiv.2411.02376`

Angulo, C., Arnould, M., Rayet, M., et al. 1999, *A compilation of charged-particle induced thermonuclear reaction rates*, Nucl. Phys. A, 656, 3, doi: `10.1016/S0375-9474(99)00030-5`

Antoni, A., & Quataert, E. 2022, *Numerical simulations of the random angular momentum in convection: Implications for supergiant collapse to form black holes*, MNRAS, 511, 176, doi: `10.1093/mnras/stab3776`

—. 2023, *Numerical simulations of the random angular momentum in convection - II. Delayed explosions of red supergiants following 'failed' supernovae*, MNRAS, 525, 1229, doi: `10.1093/mnras/stad2328`

Antoniadis, J., Tauris, T. M., Ozel, F., et al. 2016, *The millisecond pulsar mass distribution: Evidence for bimodality and constraints on the maximum neutron star mass*, arXiv e-prints, arXiv:1605.01665. `https://arxiv.org/abs/1605.01665`

Antonini, F., & Rasio, F. A. 2016, *Merging Black Hole Binaries in Galactic Nuclei: Implications for Advanced-LIGO Detections*, ApJ, 831, 187, doi: `10.3847/0004-637X/831/2/187`

Antonini, F., Toonen, S., & Hamers, A. S. 2017, *Binary Black Hole Mergers from Field Triples: Properties, Rates, and the Impact of Stellar Evolution*, ApJ, 841, 77, doi: `10.3847/1538-4357/aa6f5e`

Applebaum, E., Brooks, A. M., Christensen, C. R., et al. 2021, *Ultrafaint Dwarfs in a Milky Way Context: Introducing the Mint Condition DC Justice League Simulations*, ApJ, 906, 96, doi: `10.3847/1538-4357/abcafa`

Arora, A., Sanderson, R. E., Panithanpaisal, N., et al. 2022, *On the Stability of Tidal Streams in Action Space*, ApJ, 939, 2, doi: `10.3847/1538-4357/ac93fb`

Artymowicz, P., & Lubow, S. H. 1994, *Dynamics of Binary-Disk Interaction. I. Resonances and Disk Gap Sizes*, ApJ, 421, 651, doi: `10.1086/173679`

Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, *Astropy: A community Python package for astronomy*, A&A, 558, A33, doi: `10.1051/0004-6361/201322068`

Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, *The Astropy Project: Building an Open-science Project and Status of the v2.0 Core Package*, AJ, 156, 123, doi: `10.3847/1538-3881/aabc4f`

Astropy Collaboration, Price-Whelan, A. M., Lim, P. L., et al. 2022, *The Astropy Project: Sustaining and Growing a Community-oriented Open-source Project and the Latest Major Release (v5.0) of the Core Package*, apj, 935, 167, doi: `10.3847/1538-4357/ac7c74`

Atri, P., Miller-Jones, J. C. A., Bahramian, A., et al. 2019, *Potential kick velocity distribution of black hole X-ray binaries and implications for natal kicks*, MNRAS, 489, 3116, doi: `10.1093/mnras/stz2335`

Baade, D. 1982, *An usually short stable period of absorption line asymmetries and V/Rvariations in the spectrum of the Be star 28 CMa.*, A&A, 105, 65

Babak, S., Hewitson, M., & Petiteau, A. 2021, *LISA Sensitivity and SNR Calculations*, arXiv e-prints, arXiv:2108.01167. `https://arxiv.org/abs/2108.01167`

Babak, S., Gair, J., Sesana, A., et al. 2017, *Science with the space-based interferometer LISA. V. Extreme mass-ratio inspirals*, Phys. Rev. D, 95, 103012, doi: `10.1103/PhysRevD.95.103012`

Banerjee, S. 2020, *LISA sources from young massive and open stellar clusters*, Phys. Rev. D, 102, 103002, doi: `10.1103/PhysRevD.102.103002`

Barack, L., & Cutler, C. 2004, *LISA capture sources: Approximate waveforms, signal-to-noise ratios, and parameter estimation accuracy*, Phys. Rev. D, 69, 082005, doi: `10.1103/PhysRevD.69.082005`

—. 2007, *Using LISA extreme-mass-ratio inspiral sources to test off-Kerr deviations in the geometry of massive black holes*, Phys. Rev. D, 75, 042003, doi: `10.1103/PhysRevD.75.042003`

Baraffe, I., & El Eid, M. F. 1991, *Evolution of massive stars with variable initial compositions*, A&A, 245, 548

Bartolo, N., Caprini, C., Domcke, V., et al. 2016, *Science with the space-based interferometer LISA. IV: probing inflation with gravitational waves*, J. Cosmology Astropart. Phys., 2016, 026, doi: `10.1088/1475-7516/2016/12/026`

Bastian, N., Covey, K. R., & Meyer, M. R. 2010, *A Universal Stellar Initial Mass Function? A Critical Look at Variations*, ARA&A, 48, 339, doi: `10.1146/annurev-astro-082708-101642`

Bastian, N., Cabrera-Ziri, I., Niederhofer, F., et al. 2017, *A high fraction of Be stars in young massive clusters: evidence for a large population of near-critically rotating stars*, MNRAS, 465, 4795, doi: `10.1093/mnras/stw3042`

Basu, S., Verner, G. A., Chaplin, W. J., & Elsworth, Y. 2012, *Effect of Uncertainties in Stellar Model Parameters on Estimated Masses and Radii of Single Stars*, ApJ, 746, 76, doi: `10.1088/0004-637X/746/1/76`

Bates, S. D., Lorimer, D. R., Rane, A., & Swiggum, J. 2014, *PSRPOPPy: an open-source package for pulsar population simulations*, MNRAS, 439, 2893, doi: `10.1093/mnras/stu157`

Beasor, E. R., Davies, B., Smith, N., et al. 2020, *A new mass-loss rate prescription for red supergiants*, MNRAS, 492, 5994, doi: `10.1093/mnras/staa255`

Begelman, M. C., Blandford, R. D., & Rees, M. J. 1980, *Massive black hole binaries in active galactic nuclei*, Nature, 287, 307, doi: `10.1038/287307a0`

Behroozi, P. S., Ramirez-Ruiz, E., & Fryer, C. L. 2014, *Interpreting Short Gamma-Ray Burst Progenitor Kicks and Time Delays using the Host Galaxy-Dark Matter Halo Connection*, ApJ, 792, 123, doi: `10.1088/0004-637X/792/2/123`

Belczynski, K., Benacquista, M., & Bulik, T. 2010, *Double Compact Objects as Low-frequency Gravitational Wave Sources*, ApJ, 725, 816, doi: `10.1088/0004-637X/725/1/816`

Belczyński, K., & Bulik, T. 1999, *The effect of supernova natal kicks on compact object merger rate*, A&A, 346, 91, doi: `10.48550/arXiv.astro-ph/9901193`

Belczynski, K., Holz, D. E., Bulik, T., & O'Shaughnessy, R. 2016, *The first gravitational-wave source from the isolated evolution of two stars in the 40-100 solar mass range*, Nature, 534, 512, doi: `10.1038/nature18322`

Belczynski, K., Kalogera, V., & Bulik, T. 2002, *A Comprehensive Study of Binary Compact Objects as Gravitational Wave Sources: Evolutionary Channels, Rates, and Physical Properties*, ApJ, 572, 407, doi: `10.1086/340304`

Belczynski, K., Kalogera, V., Rasio, F. A., et al. 2008, *Compact Object Modeling with the StarTrack Population Synthesis Code*, ApJS, 174, 223, doi: `10.1086/521026`

Belczynski, K., Taam, R. E., Kalogera, V., Rasio, F. A., & Bulik, T. 2007, *On the Rarity of Double Black Hole Binaries: Consequences for Gravitational Wave Detection*, ApJ, 662, 504, doi: `10.1086/513562`

Bellinger, E. P., de Mink, S. E., van Rossem, W. E., & Justham, S. 2024, *The Potential of Asteroseismology to Resolve the Blue Supergiant Problem*, ApJ, 967, L39, doi: `10.3847/2041-8213/ad4990`

Bellovary, J. M., Cleary, C. E., Munshi, F., et al. 2019, *Multimessenger signatures of massive black holes in dwarf galaxies*, MNRAS, 482, 2913, doi: `10.1093/mnras/sty2842`

Benson, R. S. 1970, PhD thesis, University of California, Berkeley

Berger, E. 2014, *Short-Duration Gamma-Ray Bursts*, ARA&A, 52, 43, doi: `10.1146/annurev-astro-081913-035926`

Bertelli, G., Bressan, A., Chiosi, C., Fagotto, F., & Nasi, E. 1994, *Theoretical isochrones from models with new radiative opacities.*, A&AS, 106, 275

Berti, E., Cardoso, V., & Will, C. M. 2006, *Gravitational-wave spectroscopy of massive black holes with the space interferometer LISA*, Phys. Rev. D, 73, 064030, doi: `10.1103/PhysRevD.73.064030`

Bertoldi, F., & McKee, C. F. 1992, *Pressure-confined Clumps in Magnetized Molecular Clouds*, ApJ, 395, 140, doi: `10.1086/171638`

Beyer, W. H. 1987, CRC Handbook of mathematical sciences

Blaauw, A. 1961, *On the origin of the O- and B-type stars with high velocities (the "run-away" stars), and some related problems*, Bull. Astron. Inst. Netherlands, 15, 265

Bland-Hawthorn, J., & Gerhard, O. 2016, *The Galaxy in Context: Structural, Kinematic, and Integrated Properties*, ARA&A, 54, 529, doi: `10.1146/annurev-astro-081915-023441`

Bloom, J. S., Kulkarni, S. R., & Djorgovski, S. G. 2002, *The Observed Offset Distribution of Gamma-Ray Bursts from Their Host Galaxies: A Robust Clue to the Nature of the Progenitors*, AJ, 123, 1111, doi: `10.1086/338893`

Blouin, S., Shaffer, N. R., Saumon, D., & Starrett, C. E. 2020, *New Conductive Opacities for White Dwarf Envelopes*, ApJ, 899, 46, doi: `10.3847/1538-4357/ab9e75`

Bodensteiner, J., Shenar, T., & Sana, H. 2020, *Investigating the lack of main-sequence companions to massive Be stars*, A&A, 641, A42, doi: `10.1051/0004-6361/202037640`

Boersma, J. 1961, *Mathematical theory of the two-body problem with one of the masses decreasing with time*, Bull. Astron. Inst. Netherlands, 15, 291

Bolton, C. T. 1972, *Identification of Cygnus X-1 with HDE 226868*, Nature, 235, 271, doi: `10.1038/235271b0`

Borucki, W. J., Koch, D., Basri, G., et al. 2010, *Kepler Planet-Detection Mission: Introduction and First Results*, Science, 327, 977, doi: `10.1126/science.1185402`

Bouabid, M. P., Dupret, M. A., Salmon, S., et al. 2013, *Effects of the Coriolis force on high-order g modes in γ Doradus stars*, MNRAS, 429, 2500, doi: `10.1093/mnras/sts517`

Boubert, D., & Evans, N. W. 2018, *On the kinematics of a runaway Be star population*, MNRAS, 477, 5261, doi: `10.1093/mnras/sty980`

Bovy, J., Leung, H. W., Hunt, J. A. S., et al. 2019, *Life in the fast lane: a direct view of the dynamics, formation, and evolution of the Milky Way's bar*, MNRAS, 490, 4740, doi: `10.1093/mnras/stz2891`

Bovy, J., Rix, H.-W., Schlafly, E. F., et al. 2016, *The Stellar Population Structure of the Galactic Disk*, ApJ, 823, 30, doi: `10.3847/0004-637X/823/1/30`

Brandt, N., & Podsiadlowski, P. 1995, *The effects of high-velocity supernova kicks on the orbital properties and sky distributions of neutron-star binaries*, MNRAS, 274, 461, doi: `10.1093/mnras/274.2.461`

Braun, H., & Langer, N. 1995, *Effects of accretion onto massive main sequence stars.*, A&A, 297, 483

Bray, J. C., & Eldridge, J. J. 2016, *Neutron star kicks and their relationship to supernovae ejecta mass*, MNRAS, 461, 3747, doi: `10.1093/mnras/stw1275`

Breivik, K., Chatterjee, S., & Larson, S. L. 2017, *Revealing Black Holes with Gaia*, ApJ, 850, L13, doi: `10.3847/2041-8213/aa97d5`

Breivik, K., Mingarelli, C. M. F., & Larson, S. L. 2020a, *Constraining Galactic Structure with the LISA White Dwarf Foreground*, ApJ, 901, 4, doi: `10.3847/1538-4357/abab99`

Breivik, K., Rodriguez, C. L., Larson, S. L., Kalogera, V., & Rasio, F. A. 2016, *Distinguishing between Formation Channels for Binary Black Holes with LISA*, ApJ, 830, L18, doi: `10.3847/2041-8205/830/1/L18`

Breivik, K., Coughlin, S., Zevin, M., et al. 2020b, *COSMIC Variance in Binary Population Synthesis*, ApJ, 898, 71, doi: `10.3847/1538-4357/ab9d85`

Bressan, A., Marigo, P., Girardi, L., et al. 2012, *PARSEC: stellar tracks and isochrones with the PAdova and TRieste Stellar Evolution Code*, MNRAS, 427, 127, doi: `10.1111/j.1365-2966.2012.21948.x`

Broekgaarden, F. S. 2021a, *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`

—. 2021b, *BHNS 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.5178777`

—. 2021c, *NSNS 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.5189849`

Broekgaarden, F. S., Berger, E., Neijssel, C. J., et al. 2021, *Impact of massive binary star and cosmic evolution on gravitational wave observations I: black hole-neutron star mergers*, MNRAS, 508, 5028, doi: `10.1093/mnras/stab2716`

Broekgaarden, F. S., Berger, E., Stevenson, S., et al. 2022, *Impact of massive binary star and cosmic evolution on gravitational wave observations - II. Double compact object rates and properties*, MNRAS, 516, 5737, doi: `10.1093/mnras/stac1677`

Broekgaarden, F. S., Justham, S., de Mink, S. E., et al. 2019, *STROOPWAFEL: simulating rare outcomes from astrophysical populations, with application to gravitational-wave sources*, MNRAS, 490, 5228, doi: `10.1093/mnras/stz2558`

Brown, G. E. 1995, *Neutron Star Accretion and Binary Pulsar Formation*, ApJ, 440, 270, doi: `10.1086/175268`

Brunish, W. M., & Truran, J. W. 1982, *The evolution of massive stars. II. The influence of initial composition and mass loss.*, ApJS, 49, 447, doi: `10.1086/190806`

Brunt, D. 1927, *The period of simple vertical oscillations in the atmosphere*, Quarterly Journal of the Royal Meteorological Society, 53, 30, doi: `10.1002/qj.49705322103`

Bunker, A. J., Stanway, E. R., Ellis, R. S., & McMahon, R. G. 2004, *The star formation rate of the Universe at $z\sim6$ from the Hubble Ultra-Deep Field*, MNRAS, 355, 374, doi: `10.1111/j.1365-2966.2004.08326.x`

Burrows, A., & Hayes, J. 1996, *Pulsar Recoil and Gravitational Radiation Due to Asymmetrical Stellar Collapse and Explosion*, Phys. Rev. Lett., 76, 352, doi: `10.1103/PhysRevLett.76.352`

Burssens, S., Bowman, D. M., Michielsen, M., et al. 2023, *A calibration point for stellar evolution from massive star asteroseismology*, Nature Astronomy, 7, 913, doi: `10.1038/s41550-023-01978-y`

Burt, C. A., Renzo, M., Grichener, A., & Shah, N. 2025, *On the Prevalence of Early Mass Transfer for Very Massive Binaries*, Research Notes of the AAS, 9, 75, doi: `10.3847/2515-5172/adc921`

Buysschaert, B., Aerts, C., Bowman, D. M., et al. 2018, *Forward seismic modeling of the pulsating magnetic B-type star HD 43317*, A&A, 616, A148, doi: `10.1051/0004-6361/201832642`

Byrne, C. M., Stanway, E. R., Eldridge, J. J., McSwiney, L., & Townsend, O. T. 2022, *The dependence of theoretical synthetic spectra on α-enhancement in young, binary stellar populations*, MNRAS, 512, 5329, doi: `10.1093/mnras/stac807`

Caldwell, R. R., Smith, T. L., & Walker, D. G. E. 2019, *Using a primordial gravitational wave background to illuminate new physics*, Phys. Rev. D, 100, 043513, doi: `10.1103/PhysRevD.100.043513`

Cantat-Gaudin, T., Fouesneau, M., Rix, H.-W., et al. 2023, *An empirical model of the Gaia DR3 selection function*, A&A, 669, A55, doi: `10.1051/0004-6361/202244784`

Cantiello, M., Yoon, S. C., Langer, N., & Livio, M. 2007, *Binary star progenitors of long gamma-ray bursts*, A&A, 465, L29, doi: `10.1051/0004-6361:20077115`

Caprini, C., Hindmarsh, M., Huber, S., et al. 2016, *Science with the space-based interferometer eLISA. II: gravitational waves from cosmological phase transitions*, J. Cosmology Astropart. Phys., 2016, 001, doi: `10.1088/1475-7516/2016/04/001`

Carretero-Castrillo, M., Ribó, M., & Paredes, J. M. 2023, *Galactic runaway O and Be stars found using Gaia DR3*, A&A, 679, A109, doi: `10.1051/0004-6361/202346613`

Cassisi, S., Potekhin, A. Y., Pietrinferni, A., Catelan, M., & Salaris, M. 2007, *Updated Electron-Conduction Opacities: The Impact on Low-Mass Stellar Models*, ApJ, 661, 1094, doi: `10.1086/516819`

Castor, J. I., Abbott, D. C., & Klein, R. I. 1975, *Radiation-driven winds in Of stars.*, ApJ, 195, 157, doi: `10.1086/153315`

Ceverino, D., & Klypin, A. 2009, *The Role of Stellar Feedback in the Formation of Galaxies*, ApJ, 695, 292, doi: `10.1088/0004-637X/695/1/292`

Ceverino, D., Klypin, A., Klimek, E. S., et al. 2014, *Radiative feedback and the low efficiency of galaxy formation in low-mass haloes at high redshift*, MNRAS, 442, 1545, doi: `10.1093/mnras/stu956`

Chan, T. K., Kereš, D., Wetzel, A., et al. 2018, *The origin of ultra diffuse galaxies: stellar feedback and quenching*, MNRAS, 478, 906, doi: `10.1093/mnras/sty1153`

Chandrasekhar, S. 1931, *The Maximum Mass of Ideal White Dwarfs*, ApJ, 74, 81, doi: `10.1086/143324`

Chattopadhyay, D., Stevenson, S., Hurley, J. R., Bailes, M., & Broekgaarden, F. 2021, *Modelling neutron star-black hole binaries: future pulsar surveys and gravitational wave detectors*, MNRAS, 504, 3682, doi: `10.1093/mnras/stab973`

Chattopadhyay, D., Stevenson, S., Hurley, J. R., Rossi, L. J., & Flynn, C. 2020, *Modelling double neutron stars: radio and gravitational waves*, MNRAS, 494, 1587, doi: `10.1093/mnras/staa756`

Chawla, C., Chatterjee, S., Breivik, K., et al. 2022, *Gaia May Detect Hundreds of Well-characterized Stellar Black Holes*, ApJ, 931, 107, doi: `10.3847/1538-4357/ac60a5`

Chen, W.-C., Liu, D.-D., & Wang, B. 2020, *Detectability of Ultra-compact X-Ray Binaries as LISA Sources*, ApJ, 900, L8, doi: `10.3847/2041-8213/abae66`

Chen, X., Zhang, X., Li, Y., et al. 2021, *OO Dra: An Algol-type Binary Formed through an Extremely Helium-poor Mass Accretion Revealed by Asteroseismology*, ApJ, 920, 76, doi: `10.3847/1538-4357/ac1baa`

Chevance, M., Kruijssen, J. M. D., Vazquez-Semadeni, E., et al. 2020, *The Molecular Cloud Lifecycle*, Space Sci. Rev., 216, 50, doi: `10.1007/s11214-020-00674-x`

Choi, J., Dotter, A., Conroy, C., et al. 2016, *Mesa Isochrones and Stellar Tracks (MIST). I. Solar-scaled Models*, ApJ, 823, 102, doi: `10.3847/0004-637X/823/2/102`

Christensen, C., Munshi, F., & Brooks, A. 2023, in American Astronomical Society Meeting Abstracts, Vol. 55, American Astronomical Society Meeting Abstracts, 407.04

Christensen, C. R., Brooks, A. M., Munshi, F., et al. 2024, *Environment Matters: Predicted Differences in the Stellar Mass–Halo Mass Relation and History of Star Formation for Dwarf Galaxies*, ApJ, 961, 236, doi: `10.3847/1538-4357/ad0c5a`

Christensen-Dalsgaard, J. 2008, *ADIPLS—the Aarhus adiabatic oscillation package*, Ap&SS, 316, 113, doi: `10.1007/s10509-007-9689-z`

Christian, P., & Loeb, A. 2017, *LISA detection of binary black holes in the Milky Way galaxy*, MNRAS, 469, 930, doi: `10.1093/mnras/stx910`

Chruslinska, M., Belczynski, K., Bulik, T., & Gladysz, W. 2017, *Constraints on the Formation of Double Neutron Stars from the Observed Eccentricities and Current Limits on Merger Rates*, Acta Astron., 67, 37

Chugunov, A. I., Dewitt, H. E., & Yakovlev, D. G. 2007, *Coulomb tunneling for fusion reactions in dense matter: Path integral MonteCarlo versus mean field*, Phys. Rev. D, 76, 025028, doi: `10.1103/PhysRevD.76.025028`

Claeys, J. S. W., de Mink, S. E., Pols, O. R., Eldridge, J. J., & Baes, M. 2011, *Binary progenitor models of type IIb supernovae*, A&A, 528, A131, doi: `10.1051/0004-6361/201015410`

Claeys, J. S. W., Pols, O. R., Izzard, R. G., Vink, J., & Verbunt, F. W. M. 2014, *Theoretical uncertainties of the Type Ia supernova rate*, A&A, 563, A83, doi: `10.1051/0004-6361/201322714`

Claret, A., & Torres, G. 2017, *The Dependence of Convective Core Overshooting on Stellar Mass: A Semi-empirical Determination Using the Diffusive Approach with Two Different Element Mixtures*, ApJ, 849, 18, doi: `10.3847/1538-4357/aa8770`

Coleman, M. S. B., & Burrows, A. 2022, *Kicks and induced spins of neutron stars at birth*, MNRAS, 517, 3938, doi: `10.1093/mnras/stac2573`

Collette, A. 2013, Python and HDF5 (O'Reilly)

Colpi, M., Holley-Bockelmann, K., Bogdanovic, T., et al. 2019, *Astro2020 science white paper: The gravitational wave view of massive black holes*, arXiv e-prints, arXiv:1903.06867. `https://arxiv.org/abs/1903.06867`

Cornish, N. J., & Larson, S. L. 2003, *LISA data analysis: Source identification and subtraction*, Phys. Rev. D, 67, 103001, doi: `10.1103/PhysRevD.67.103001`

Corral-Santana, J. M., Casares, J., Muñoz-Darias, T., et al. 2016, *BlackCAT: A catalogue of stellar-mass black holes in X-ray transients*, A&A, 587, A61, doi: 10.1051/0004-6361/201527130

Cox, A. N., Morgan, S. M., Rogers, F. J., & Iglesias, C. A. 1992, *An Opacity Mechanism for the Pulsations of OB Stars*, ApJ, 393, 272, doi: 10.1086/171504

Crowe, R. A., & Matalas, R. 1982, *Semiconvection in Low-Mass Main Sequence Stars*, A&A, 108, 55

Crowther, P. A. 2007, *Physical Properties of Wolf-Rayet Stars*, ARA&A, 45, 177, doi: 10.1146/annurev.astro.45.051806.110615

Cyburt, R. H., Amthor, A. M., Ferguson, R., et al. 2010, *The JINA REACLIB Database: Its Recent Updates and Impact on Type-I X-ray Bursts*, ApJS, 189, 240, doi: 10.1088/0067-0049/189/1/240

Da Rio, N., Tan, J. C., Covey, K. R., et al. 2017, *IN-SYNC. V. Stellar Kinematics and Dynamics in the Orion A Molecular Cloud*, ApJ, 845, 105, doi: 10.3847/1538-4357/aa7a5b

Dabringhausen, J., Marks, M., & Kroupa, P. 2022, *BIPOS1 - a computer programme for the dynamical processing of the initial binary star population*, MNRAS, 510, 413, doi: 10.1093/mnras/stab3288

Darragh-Ford, E., Hunt, J. A. S., Price-Whelan, A. M., & Johnston, K. V. 2023, *ESCARGOT: Mapping Vertical Phase Spiral Characteristics Throughout the Real and Simulated Milky Way*, ApJ, 955, 74, doi: 10.3847/1538-4357/acf1fc

Davies, M. B., Benz, W., & Hills, J. G. 1992, *Stellar Encounters Involving Neutron Stars in Globular Cluster Cores*, ApJ, 401, 246, doi: 10.1086/172056

De Donder, E., & Vanbeveren, D. 2003, *The galactic evolution of the supernova rates*, New A, 8, 817, doi: 10.1016/S1384-1076(03)00069-1

De Donder, E., Vanbeveren, D., & van Bever, J. 1997, *The number of O-type runaways, the number of O and Wolf-Rayet stars with a compact companion and the formation rate of double pulsars predicted by massive close binary evolution.*, A&A, 318, 812

de Kool, M. 1990, *Common Envelope Evolution and Double Cores of Planetary Nebulae*, ApJ, 358, 189, doi: 10.1086/168974

De Marco, O., Passy, J.-C., Moe, M., et al. 2011, *On the α formalism for the common envelope interaction*, MNRAS, 411, 2277, doi: 10.1111/j.1365-2966.2010.17891.x

de Mink, S. E., & Belczynski, K. 2015, *Merger Rates of Double Neutron Stars and Stellar Origin Black Holes: The Impact of Initial Conditions on Binary Evolution Predictions*, ApJ, 814, 58, doi: 10.1088/0004-637X/814/1/58

de Mink, S. E., Cantiello, M., Langer, N., et al. 2009, *Rotational mixing in massive binaries. Detached short-period systems*, A&A, 497, 243, doi: 10.1051/0004-6361/200811439

de Mink, S. E., Langer, N., Izzard, R. G., Sana, H., & de Koter, A. 2013, *The Rotation Rates of Massive Stars: The Role of Binary Interaction through Tides, Mass Transfer, and Mergers*, ApJ, 764, 166, doi: 10.1088/0004-637X/764/2/166

de Mink, S. E., Pols, O. R., & Hilditch, R. W. 2007, *Efficiency of mass transfer in massive close binaries. Tests from double-lined eclipsing binaries in the SMC*, A&A, 467, 1181, doi: 10.1051/0004-6361:20067007

de Mink, S. E., Pols, O. R., & Yoon, S. C. 2008, in American Institute of Physics Conference Series, Vol. 990, First Stars III, ed. B. W. O'Shea & A. Heger (AIP), 230–232, doi: 10.1063/1.2905549

de Mink, S. E., Sana, H., Langer, N., Izzard, R. G., & Schneider, F. R. N. 2014, *The Incidence of Stellar Mergers and Mass Gainers among Massive Stars*, ApJ, 782, 7, doi: 10.1088/0004-637X/782/1/7

Dekel, A., & Silk, J. 1986, *The Origin of Dwarf Galaxies, Cold Dark Matter, and Biased Galaxy Formation*, ApJ, 303, 39, doi: 10.1086/164050

Dervişoğlu, A., Tout, C. A., & Ibanoğlu, C. 2010, *Spin angular momentum evolution of the long-period Algols*, MNRAS, 406, 1071, doi: 10.1111/j.1365-2966.2010.16732.x

Dewdney, P. E., Hall, P. J., Schilizzi, R. T., & Lazio, T. J. L. W. 2009, *The Square Kilometre Array*, IEEE Proceedings, 97, 1482, doi: 10.1109/JPROC.2009.2021005

Dewey, R. J., & Cordes, J. M. 1987, *Monte Carlo Simulations of Radio Pulsars and Their Progenitors*, ApJ, 321, 780, doi: 10.1086/165671

Dewi, J. D. M., Pols, O. R., Savonije, G. J., & van den Heuvel, E. P. J. 2002, *The evolution of naked helium stars with a neutron star companion in close binary systems*, MNRAS, 331, 1027, doi: 10.1046/j.1365-8711.2002.05257.x

Di Carlo, U. N., Mapelli, M., Giacobbo, N., et al. 2020, *Binary black holes in young star clusters: the impact of metallicity*, MNRAS, 498, 495, doi: 10.1093/mnras/staa2286

Diemer, B., More, S., & Kravtsov, A. V. 2013, *The Pseudo-evolution of Halo Mass*, ApJ, 766, 25, doi: 10.1088/0004-637X/766/1/25

Disberg, P., Gaspari, N., & Levan, A. J. 2024, *Deceleration of kicked objects due to the Galactic potential*, A&A, 687, A272, doi: 10.1051/0004-6361/202449996

Dominik, M., Belczynski, K., Fryer, C., et al. 2013, *Double Compact Objects. II. Cosmological Merger Rates*, ApJ, 779, 72, doi: 10.1088/0004-637X/779/1/72

Dominik, M., Berti, E., O'Shaughnessy, R., et al. 2015, *Double Compact Objects III: Gravitational-wave Detection Rates*, ApJ, 806, 263, doi: 10.1088/0004-637X/806/2/263

Dormand, J., & Prince, P. 1980, *A family of embedded Runge-Kutta formulae*, Journal of Computational and Applied Mathematics, 6, 19, doi: https://doi.org/10.1016/0771-050X(80)90013-3

Dosopoulou, F., & Kalogera, V. 2016a, *Orbital Evolution of Mass-transferring Eccentric Binary Systems. I. Phase-dependent Evolution*, ApJ, 825, 70, doi: 10.3847/0004-637X/825/1/70

—. 2016b, *Orbital Evolution of Mass-transferring Eccentric Binary Systems. II. Secular Evolution*, ApJ, 825, 71, doi: 10.3847/0004-637X/825/1/71

Dotter, A. 2016, *MESA Isochrones and Stellar Tracks (MIST) 0: Methods for the Construction of Stellar Isochrones*, ApJS, 222, 8, doi: 10.3847/0067-0049/222/1/8

Drout, M. R., Götberg, Y., Ludwig, B. A., et al. 2023, *An observed population of intermediate-mass helium stars that have been stripped in binaries*, Science, 382, 1287, doi: 10.1126/science.ade4970

du Buisson, L., Marchant, P., Podsiadlowski, P., et al. 2020, *Cosmic rates of black hole mergers and pair-instability supernovae from chemically homogeneous binary evolution*, MNRAS, 499, 5941, doi: 10.1093/mnras/staa3225

Duchêne, G., & Kraus, A. 2013, *Stellar Multiplicity*, ARA&A, 51, 269, doi: 10.1146/annurev-astro-081710-102602

Dziembowski, W. A. 1971, *Nonradial Oscillations of Evolved Stars. I. Quasiadiabatic Approximation*, Acta Astron., 21, 289

Dziembowski, W. A., Moskalik, P., & Pamyatnykh, A. A. 1993, *The opacity mechanism in B-type stars - II. Excitation of high-order g-modes in main-sequence stars.*, MNRAS, 265, 588, doi: 10.1093/mnras/265.3.588

Eddington, A. S. 1926, The Internal Constitution of the Stars

Edelman, B., Doctor, Z., Godfrey, J., & Farr, B. 2022, *Ain't No Mountain High Enough: Semiparametric Modeling of LIGO-Virgo's Binary Black Hole Mass Distribution*, ApJ, 924, 101, doi: 10.3847/1538-4357/ac3667

Eggenberger, P., Meynet, G., Maeder, A., et al. 2008, *The Geneva stellar evolution code*, Ap&SS, 316, 43, doi: 10.1007/s10509-007-9511-y

Eggleton, P. P. 1971, *The evolution of low mass stars*, MNRAS, 151, 351, doi: 10.1093/mnras/151.3.351

—. 1983, *Approximations to the radii of Roche lobes*, ApJ, 268, 368, doi: 10.1086/160960

Eggleton, P. P., Fitchett, M. J., & Tout, C. A. 1989, *The Distribution of Visual Binaries with Two Bright Components*, ApJ, 347, 998, doi: 10.1086/168190

Eichler, D., Livio, M., Piran, T., & Schramm, D. N. 1989, *Nucleosynthesis, neutrino bursts and γ-rays from coalescing neutron stars*, Nature, 340, 126, doi: 10.1038/340126a0

Eilers, A.-C., Hogg, D. W., Rix, H.-W., & Ness, M. K. 2019, *The Circular Velocity Curve of the Milky Way from 5 to 25 kpc*, ApJ, 871, 120, doi: 10.3847/1538-4357/aaf648

Ekström, S., Georgy, C., Eggenberger, P., et al. 2012, *Grids of stellar models with rotation. I. Models from 0.8 to 120 $M_\odot$ at solar metallicity ($Z = 0.014$)*, A&A, 537, A146, doi: 10.1051/0004-6361/201117751

El-Badry, K., & Quataert, E. 2020, *Not so fast: LB-1 is unlikely to contain a 70 $M_\odot$ black hole*, MNRAS, 493, L22, doi: 10.1093/mnrasl/slaa004

El-Badry, K., Quataert, E., Wetzel, A., et al. 2018, *Gas kinematics, morphology and angular momentum in the FIRE simulations*, MNRAS, 473, 1930, doi: 10.1093/mnras/stx2482

El-Badry, K., Rix, H.-W., Quataert, E., et al. 2023, *A Sun-like star orbiting a black hole*, MNRAS, 518, 1057, doi: 10.1093/mnras/stac3140

El-Badry, K., Rix, H.-W., Latham, D. W., et al. 2024, *A population of neutron star candidates in wide orbits from Gaia astrometry*, The Open Journal of Astrophysics, 7, 58, doi: 10.33232/001c.121261

Eldridge, J. J., Fraser, M., Smartt, S. J., Maund, J. R., & Crockett, R. M. 2013, *The death of massive stars - II. Observational constraints on the progenitors of Type Ibc supernovae*, MNRAS, 436, 774, doi: 10.1093/mnras/stt1612

Eldridge, J. J., Izzard, R. G., & Tout, C. A. 2008, *The effect of massive binaries on stellar populations and supernova progenitors*, MNRAS, 384, 1109, doi: 10.1111/j.1365-2966.2007.12738.x

Eldridge, J. J., Langer, N., & Tout, C. A. 2011, *Runaway stars as progenitors of supernovae and gamma-ray bursts*, MNRAS, 414, 3501, doi: 10.1111/j.1365-2966.2011.18650.x

Eldridge, J. J., Stanway, E. R., Breivik, K., et al. 2020, *Weighing in on black hole binaries with BPASS: LB-1 does not contain a 70 Msun black hole*, MNRAS, 495, 2786, doi: 10.1093/mnras/staa1324

Eldridge, J. J., Stanway, E. R., Xiao, L., et al. 2017, *Binary Population and Spectral Synthesis Version 2.1: Construction, Observational Verification, and New Results*, PASA, 34, e058, doi: 10.1017/pasa.2017.51

Eldridge, J. J., & Tout, C. A. 2004, *A more detailed look at the opacities for enriched carbon and oxygen mixtures*, MNRAS, 348, 201, doi: 10.1111/j.1365-2966.2004.07344.x

Ertl, T., Janka, H. T., Woosley, S. E., Sukhbold, T., & Ugliano, M. 2016, *A Two-parameter Criterion for Classifying the Explodability of Massive Stars by the Neutrino-driven Mechanism*, ApJ, 818, 124, doi: 10.3847/0004-637X/818/2/124

Ertl, T., Woosley, S. E., Sukhbold, T., & Janka, H. T. 2020, *The Explosion of Helium Stars Evolved with Mass Loss*, ApJ, 890, 51, doi: 10.3847/1538-4357/ab6458

Evans, F. A., Renzo, M., & Rossi, E. M. 2020, *Core-collapse supernovae in binaries as the origin of galactic hyper-runaway stars*, MNRAS, 497, 5344, doi: 10.1093/mnras/staa2334

Fall, S. M., & Efstathiou, G. 1980, *Formation and rotation of disc galaxies with haloes.*, MNRAS, 193, 189, doi: 10.1093/mnras/193.2.189

Fang, X., Thompson, T. A., & Hirata, C. M. 2019, *The Population of Eccentric Binary Black Holes: Implications for mHz Gravitational-wave Experiments*, ApJ, 875, 75, doi: 10.3847/1538-4357/ab0e6a

Farah, A., Fishbach, M., Essick, R., Holz, D. E., & Galaudage, S. 2022, *Bridging the Gap: Categorizing Gravitational-wave Events at the Transition between Neutron Stars and Black Holes*, ApJ, 931, 108, doi: 10.3847/1538-4357/ac5f03

Farmer, R., Renzo, M., de Mink, S. E., Marchant, P., & Justham, S. 2019, *Mind the Gap: The Location of the Lower Edge of the Pair-instability Supernova Black Hole Mass Gap*, ApJ, 887, 53, doi: 10.3847/1538-4357/ab518b

Farr, W. M., Sravan, N., Cantrell, A., et al. 2011, *The Mass Distribution of Stellar-mass Black Holes*, ApJ, 741, 103, doi: 10.1088/0004-637X/741/2/103

Ferguson, J. W., Alexander, D. R., Allard, F., et al. 2005, *Low-Temperature Opacities*, ApJ, 623, 585, doi: 10.1086/428642

Finn, L. S., & Thorne, K. S. 2000, *Gravitational waves from a compact star in a circular, inspiral orbit, in the equatorial plane of a massive, spinning black hole, as observed by LISA*, Phys. Rev. D, 62, 124021, doi: 10.1103/PhysRevD.62.124021

Flanagan, É. É., & Hughes, S. A. 1998, *Measuring gravitational waves from binary black hole coalescences. I. Signal to noise for inspiral, merger, and ringdown*, Phys. Rev. D, 57, 4535, doi: 10.1103/PhysRevD.57.4535

Flannery, B. P., & van den Heuvel, E. P. J. 1975, *On the origin of the binary pulsar PSR 1913+16.*, A&A, 39, 61

Fong, W.-f., Nugent, A. E., Dong, Y., et al. 2022, *Short GRB Host Galaxies. I. Photometric and Spectroscopic Catalogs, Host Associations, and Galactocentric Offsets*, ApJ, 940, 56, doi: 10.3847/1538-4357/ac91d0

Fragos, T., Andrews, J. J., Ramirez-Ruiz, E., et al. 2019, *The Complete Evolution of a Neutron-star Binary through a Common Envelope Phase Using 1D Hydrodynamic Simulations*, ApJ, 883, L45, doi: 10.3847/2041-8213/ab40d1

Fragos, T., Lehmer, B., Tremmel, M., et al. 2013, *X-Ray Binary Evolution Across Cosmic Time*, ApJ, 764, 41, doi: `10.1088/0004-637X/764/1/41`

Fragos, T., Andrews, J. J., Bavera, S. S., et al. 2023, *POSYDON: A General-purpose Population Synthesis Code with Detailed Binary-evolution Simulations*, ApJS, 264, 45, doi: `10.3847/1538-4365/ac90c1`

Frank, J., King, A. R., & Raine, D. J. 1985, Accretion power in astrophysics

Frankel, N., Rix, H.-W., Ting, Y.-S., Ness, M., & Hogg, D. W. 2018, *Measuring Radial Orbit Migration in the Galactic Disk*, ApJ, 865, 96, doi: `10.3847/1538-4357/aadba5`

Frankel, N., Sanders, J., Rix, H.-W., Ting, Y.-S., & Ness, M. 2019, *The Inside-out Growth of the Galactic Disk*, ApJ, 884, 99, doi: `10.3847/1538-4357/ab4254`

Fryer, C., Burrows, A., & Benz, W. 1998, *Population Syntheses for Neutron Star Systems with Intrinsic Kicks*, ApJ, 496, 333, doi: `10.1086/305348`

Fryer, C. L., Belczynski, K., Ramirez-Ruiz, E., et al. 2015, *The Fate of the Compact Remnant in Neutron Star Mergers*, ApJ, 812, 24, doi: `10.1088/0004-637X/812/1/24`

Fryer, C. L., Belczynski, K., Wiktorowicz, G., et al. 2012, *Compact Remnant Mass Function: Dependence on the Explosion Mechanism and Metallicity*, ApJ, 749, 91, doi: `10.1088/0004-637X/749/1/91`

Fujii, M. S., & Portegies Zwart, S. 2011, *The Origin of OB Runaway Stars*, Science, 334, 1380, doi: `10.1126/science.1211927`

Fuller, G. M., Fowler, W. A., & Newman, M. J. 1985, *Stellar weak interaction rates for intermediate-mass nuclei. IV - Interpolation procedures for rapidly varying lepton capture rates using effective log (ft)-values*, ApJ, 293, 1, doi: `10.1086/163208`

Gerosa, D., Ma, S., Wong, K. W. K., et al. 2019, *Multiband gravitational-wave event rates and stellar physics*, Phys. Rev. D, 99, 103004, doi: `10.1103/PhysRevD.99.103004`

Ginat, Y. B., Glanz, H., Perets, H. B., Grishin, E., & Desjacques, V. 2020, *Gravitational waves from in-spirals of compact objects in binary common-envelope evolution*, MNRAS, 493, 4861, doi: `10.1093/mnras/staa465`

Girichidis, P., Walch, S., Naab, T., et al. 2016, *The SILCC (SImulating the LifeCycle of molecular Clouds) project - II. Dynamical evolution of the supernova-driven ISM and the launching of outflows*, MNRAS, 456, 3432, doi: `10.1093/mnras/stv2742`

Goldberg, D., & Mazeh, T. 1994, *The mass-ratio distribution of the spectroscopic binaries in the Pleiades.*, A&A, 282, 801

Goldstein, J., & Townsend, R. H. D. 2020, *The Contour Method: a New Approach to Finding Modes of Nonadiabatic Stellar Pulsations*, ApJ, 899, 116, doi: `10.3847/1538-4357/aba748`

Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, *HEALPix: A Framework for High-Resolution Discretization and Fast Analysis of Data Distributed on the Sphere*, ApJ, 622, 759, doi: `10.1086/427976`

Gosnell, N. M., Leiner, E. M., Mathieu, R. D., et al. 2019, *Constraining Mass-transfer Histories of Blue Straggler Stars with COS Spectroscopy of White Dwarf Companions*, ApJ, 885, 45, doi: `10.3847/1538-4357/ab4273`

Götberg, Y., de Mink, S. E., & Groh, J. H. 2017, *Ionizing spectra of stars that lose their envelope through interaction with a binary companion: role of metallicity*, A&A, 608, A11, doi: `10.1051/0004-6361/201730472`

Götberg, Y., de Mink, S. E., Groh, J. H., Leitherer, C., & Norman, C. 2019, *The impact of stars stripped in binaries on the integrated spectra of stellar populations*, A&A, 629, A134, doi: `10.1051/0004-6361/201834525`

Götberg, Y., Korol, V., Lamberts, A., et al. 2020, *Stars Stripped in Binaries: The Living Gravitational-wave Sources*, ApJ, 904, 56, doi: `10.3847/1538-4357/abbda5`

Götberg, Y., Drout, M. R., Ji, A. P., et al. 2023, *Stellar Properties of Observed Stars Stripped in Binaries in the Magellanic Clouds*, ApJ, 959, 125, doi: `10.3847/1538-4357/ace5a3`

Green, G. 2018, *dustmaps: A Python interface for maps of interstellar dust*, The Journal of Open Source Software, 3, 695, doi: `10.21105/joss.00695`

Green, G. M., Schlafly, E., Zucker, C., Speagle, J. S., & Finkbeiner, D. 2019, *A 3D Dust Map Based on Gaia, Pan-STARRS 1, and 2MASS*, ApJ, 887, 93, doi: `10.3847/1538-4357/ab5362`

Griffith, E., Weinberg, D. H., Johnson, J. A., et al. 2021, *The Similarity of Abundance Ratio Trends and Nucleosynthetic Patterns in the Milky Way Disk and Bulge*, ApJ, 909, 77, doi: `10.3847/1538-4357/abd6be`

Grudić, M. Y., Hafen, Z., Rodriguez, C. L., et al. 2023, *Great balls of FIRE - I. The formation of star clusters across cosmic time in a Milky Way-mass galaxy*, MNRAS, 519, 1366, doi: `10.1093/mnras/stac3573`

Gunawardhana, M. L. P., Hopkins, A. M., Sharp, R. G., et al. 2011, *Galaxy and Mass Assembly (GAMA): the star formation rate dependence of the stellar initial mass function*, MNRAS, 415, 1647, doi: `10.1111/j.1365-2966.2011.18800.x`

Gunn, J. E., & Ostriker, J. P. 1970, *On the Nature of Pulsars. III. Analysis of Observations*, ApJ, 160, 979, doi: `10.1086/150487`

Guo, Y., Wang, L., Liu, C., et al. 2024, *A Catalog of Early-type Runaway Stars from LAMOST DR8*, ApJS, 272, 45, doi: `10.3847/1538-4365/ad46f8`

Guo, Z., Gies, D. R., & Matson, R. A. 2017a, *Gravity Modes Reveal the Internal Rotation of a Post-mass-transfer Gamma Doradus/Delta Scuti Hybrid Pulsator in Kepler Eclipsing Binary KIC 9592855*, ApJ, 851, 39, doi: `10.3847/1538-4357/aa978c`

Guo, Z., Gies, D. R., Matson, R. A., et al. 2017b, *KIC 8262223: A Post-mass Transfer Eclipsing Binary Consisting of a Delta Scuti Pulsator and a Helium White Dwarf Precursor*, ApJ, 837, 114, doi: `10.3847/1538-4357/aa61a4`

Guo, Z., & Li, G. 2019, *A Mass-accreting Gamma Doradus Pulsator with a Synchronized Core in Kepler Eclipsing Binary KIC 7385478*, ApJ, 882, L5, doi: `10.3847/2041-8213/ab3a53`

Guzik, J. A., Kaye, A. B., Bradley, P. A., Cox, A. N., & Neuforge, C. 2000, *Driving the Gravity-Mode Pulsations in γ Doradus Variables*, ApJ, 542, L57, doi: `10.1086/312908`

Haiman, Z., & Loeb, A. 1997, *Signatures of Stellar Reionization of the Universe*, ApJ, 483, 21, doi: `10.1086/304238`

Hamada, T., & Salpeter, E. E. 1961, *Models for Zero-Temperature Stars.*, ApJ, 134, 683, doi: `10.1086/147195`

Hamann, W. R., Gräfener, G., Liermann, A., et al. 2019, *The Galactic WN stars revisited. Impact of Gaia distances on fundamental stellar parameters*, A&A, 625, A57, doi: `10.1051/0004-6361/201834850`

Hamann, W. R., & Koesterke, L. 1998, *Spectrum formation in clumped stellar winds: consequences for the analyses of Wolf-Rayet spectra*, A&A, 335, 1003

Hamers, A. S., Rantala, A., Neunteufel, P., Preece, H., & Vynatheya, P. 2021, *Multiple Stellar Evolution: a population synthesis algorithm to model the stellar, binary, and dynamical evolution of multiple-star systems*, MNRAS, 502, 4479, doi: `10.1093/mnras/stab287`

Hansen, B. M. S., & Phinney, E. S. 1997, *The pulsar kick velocity distribution*, MNRAS, 291, 569, doi: `10.1093/mnras/291.3.569`

Hatta, Y. 2023, *Semi-analytical Expression of G-mode Period Spacing: The Case of Brunt-Väisälä Frequency with Not a Jump but a Ramp*, ApJ, 950, 165, doi: `10.3847/1538-4357/acd4b9`

Heber, U. 2016, *Hot Subluminous Stars*, PASP, 128, 082001, doi: `10.1088/1538-3873/128/966/082001`

Heisler, J., & Tremaine, S. 1986, *The influence of the Galactic tidal field on the Oort comet cloud*, Icarus, 65, 13, doi: `10.1016/0019-1035(86)90060-6`

Hellings, P. 1983, *Phenomenological Study of Massive Accretion Stars*, Ap&SS, 96, 37, doi: `10.1007/BF00661941`

Hendriks, D., & Izzard, R. 2023, *binary_c-python: A Python-based stellar population synthesis tool and interface to binary_c*, The Journal of Open Source Software, 8, 4642, doi: `10.21105/joss.04642`

Hennebelle, P., & Grudić, M. Y. 2024, *The Physical Origin of the Stellar Initial Mass Function*, ARA&A, 62, 63, doi: `10.1146/annurev-astro-052622-031748`

Henyey, L. G., Forbes, J. E., & Gould, N. L. 1964, *A New Method of Automatic Computation of Stellar Evolution.*, ApJ, 139, 306, doi: `10.1086/147754`

Henyey, L. G., Wilets, L., Böhm, K. H., Lelevier, R., & Levee, R. D. 1959, *A Method for Automatic Computation of Stellar Evolution.*, ApJ, 129, 628, doi: `10.1086/146661`

Hernquist, L., & Ostriker, J. P. 1992, *A Self-consistent Field Method for Galactic Dynamics*, ApJ, 386, 375, doi: `10.1086/171025`

Herwig, F. 2000, *The evolution of AGB stars with convective overshoot*, A&A, 360, 952, doi: `10.48550/arXiv.astro-ph/0007139`

Hills, J. G. 1980, *The effect of mass loss on the dynamical evolution of a stellar system - Analytic approximations*, ApJ, 235, 986, doi: `10.1086/157703`

—. 1983, *The effects of sudden mass loss and a random kick velocity produced in a supernova explosion on the dynamics of a binary star of arbitrary orbital eccentricity. Applications to X-ray binaries and to the binarypulsars.*, ApJ, 267, 322, doi: `10.1086/160871`

Hirai, R., & Podsiadlowski, P. 2022, *Neutron stars colliding with binary companions: formation of hypervelocity stars, pulsar planets, bumpy superluminous supernovae and Thorne-Żytkow objects*, MNRAS, 517, 4544, doi: `10.1093/mnras/stac3007`

Hirai, R., Podsiadlowski, P., & Yamada, S. 2018, *Comprehensive Study of Ejecta-companion Interaction for Core-collapse Supernovae in Massive Binaries*, ApJ, 864, 119, doi: `10.3847/1538-4357/aad6a0`

Hjellming, M. S., & Webbink, R. F. 1987, *Thresholds for Rapid Mass Transfer in Binary System. I. Polytropic Models*, ApJ, 318, 794, doi: `10.1086/165412`

Hobbs, G., Lorimer, D. R., Lyne, A. G., & Kramer, M. 2005, *A statistical study of 233 pulsar proper motions*, MNRAS, 360, 974, doi: `10.1111/j.1365-2966.2005.09087.x`

Hofmeister, E., Kippenhahn, R., & Weigert, A. 1964, *Sternentwicklung I. Ein Programm zur Lösung der zeitabhängigen Aufbaugleichungen. Mit 3 Textabbildungen*, ZAp, 59, 215

Hoogerwerf, R., de Bruijne, J. H. J., & de Zeeuw, P. T. 2001, *On the origin of the O and B-type stars with high velocities. II. Runaway stars and pulsars ejected from the nearby young stellar groups*, A&A, 365, 49, doi: `10.1051/0004-6361:20000014`

Hopkins, P. F., Kereš, D., Oñorbe, J., et al. 2014, *Galaxies on FIRE (Feedback In Realistic Environments): stellar feedback explains cosmologically inefficient star formation*, MNRAS, 445, 581, doi: `10.1093/mnras/stu1738`

Hopkins, P. F., Quataert, E., & Murray, N. 2012, *Stellar feedback in galaxies and the origin of galaxy-scale winds*, MNRAS, 421, 3522, doi: `10.1111/j.1365-2966.2012.20593.x`

Hopkins, P. F., Wetzel, A., Kereš, D., et al. 2018a, *How to model supernovae in simulations of star and galaxy formation*, MNRAS, 477, 1578, doi: `10.1093/mnras/sty674`

—. 2018b, *FIRE-2 simulations: physics versus numerics in galaxy formation*, MNRAS, 480, 800, doi: `10.1093/mnras/sty1690`

Hopkins, P. F., Wetzel, A., Wheeler, C., et al. 2023a, *FIRE-3: updated stellar evolution models, yields, and microphysics and fitting functions for applications in galaxy simulations*, MNRAS, 519, 3154, doi: `10.1093/mnras/stac3489`

Hopkins, P. F., Gurvich, A. B., Shen, X., et al. 2023b, *What causes the formation of discs and end of bursty star formation?*, MNRAS, 525, 2241, doi: `10.1093/mnras/stad1902`

Hotokezaka, K., Nissanke, S., Hallinan, G., et al. 2016, *Radio Counterparts of Compact Binary Mergers Detectable in Gravitational Waves: A Simulation for an Optimized Survey*, ApJ, 831, 190, doi: `10.3847/0004-637X/831/2/190`

Hu, C.-Y. 2019, *Supernova-driven winds in simulated dwarf galaxies*, MNRAS, 483, 3363, doi: `10.1093/mnras/sty3252`

Hu, C.-Y., Naab, T., Glover, S. C. O., Walch, S., & Clark, P. C. 2017, *Variable interstellar radiation fields in simulated dwarf galaxies: supernovae versus photoelectric heating*, MNRAS, 471, 2151, doi: `10.1093/mnras/stx1773`

Hu, C.-Y., Naab, T., Walch, S., Glover, S. C. O., & Clark, P. C. 2016, *Star formation and molecular hydrogen in dwarf galaxies: a non-equilibrium view*, MNRAS, 458, 3528, doi: `10.1093/mnras/stw544`

Hu, C.-Y., Zhukovska, S., Somerville, R. S., & Naab, T. 2019, *Thermal and non-thermal dust sputtering in hydrodynamical simulations of the multiphase interstellar medium*, MNRAS, 487, 3252, doi: `10.1093/mnras/stz1481`

Hu, C.-Y., Smith, M. C., Teyssier, R., et al. 2023, *Code Comparison in Galaxy-scale Simulations with Resolved Supernova Feedback: Lagrangian versus Eulerian Methods*, ApJ, 950, 132, doi: `10.3847/1538-4357/accf9e`

Huang, S.-J., Hu, Y.-M., Korol, V., et al. 2020, *Science with the TianQin Observatory: Preliminary results on Galactic double white dwarf binaries*, Phys. Rev. D, 102, 063021, doi: 10.1103/PhysRevD.102.063021

Huang, S.-S. 1963, *Modes of Mass Ejection by Binary Stars and the Effect on Their Orbital Periods.*, ApJ, 138, 471, doi: 10.1086/147659

Hulse, R. A., & Taylor, J. H. 1975, *Discovery of a pulsar in a binary system.*, ApJ, 195, L51, doi: 10.1086/181708

Hurley, J. R., Pols, O. R., & Tout, C. A. 2000, *Comprehensive analytic formulae for stellar evolution as a function of mass and metallicity*, MNRAS, 315, 543, doi: 10.1046/j.1365-8711.2000.03426.x

Hurley, J. R., Tout, C. A., & Pols, O. R. 2002, *Evolution of binary stars and the effect of tides on binary populations*, MNRAS, 329, 897, doi: 10.1046/j.1365-8711.2002.05038.x

Iaconi, R., & De Marco, O. 2019, *Speaking with one voice: simulations and observations discuss the common envelope α parameter*, MNRAS, 490, 2550, doi: 10.1093/mnras/stz2756

Iben, Jr., I., & Ehrman, J. R. 1962, *The Internal Structure of Middle Main-Sequence Stars.*, ApJ, 135, 770, doi: 10.1086/147320

Iglesias, C. A., & Rogers, F. J. 1993, *Radiative opacities for carbon- and oxygen-rich mixtures*, ApJ, 412, 752, doi: 10.1086/172958

—. 1996, *Updated Opal Opacities*, ApJ, 464, 943, doi: 10.1086/177381

Igoshev, A. P. 2020, *The observed velocity distribution of young pulsars - II. Analysis of complete PSRπ*, MNRAS, 494, 3663, doi: 10.1093/mnras/staa958

Iorio, G., Mapelli, M., Costa, G., et al. 2023, *Compact object mergers: exploring uncertainties from stellar and binary evolution with SEVN*, MNRAS, 524, 426, doi: 10.1093/mnras/stad1630

Irwin, A. W. 2004, *The FreeEOS Code for Calculating the Equation of State for Stellar Interiors.* http://freeeos.sourceforge.net/

Itoh, N., Hayashi, H., Nishikawa, A., & Kohyama, Y. 1996, *Neutrino Energy Loss in Stellar Interiors. VII. Pair, Photo-, Plasma, Bremsstrahlung, and Recombination Neutrino Processes*, ApJS, 102, 411, doi: 10.1086/192264

Ivanova, N., Heinke, C. O., Rasio, F. A., Belczynski, K., & Fregeau, J. M. 2008, *Formation and evolution of compact binaries in globular clusters - II. Binaries with neutron stars*, MNRAS, 386, 553, doi: 10.1111/j.1365-2966.2008.13064.x

Ivanova, N., Justham, S., & Ricker, P. 2020, Common Envelope Evolution, doi: 10.1088/2514-3433/abb6f0

Ivanova, N., & Taam, R. E. 2004, *Thermal Timescale Mass Transfer and the Evolution of White Dwarf Binaries*, ApJ, 601, 1058, doi: 10.1086/380561

Ivanova, N., Justham, S., Chen, X., et al. 2013, *Common envelope evolution: where we stand and how we can move forward*, A&A Rev., 21, 59, doi: 10.1007/s00159-013-0059-2

Izzard, R. G., Dray, L. M., Karakas, A. I., Lugaro, M., & Tout, C. A. 2006, *Population nucleosynthesis in single and binary stars. I. Model*, A&A, 460, 565, doi: 10.1051/0004-6361:20066129

Izzard, R. G., Glebbeek, E., Stancliffe, R. J., & Pols, O. R. 2009, *Population synthesis of binary carbon-enhanced metal-poor stars*, A&A, 508, 1359, doi: 10.1051/0004-6361/200912827

Izzard, R. G., & Jermyn, A. S. 2023, *Circumbinary discs for stellar population models*, MNRAS, 521, 35, doi: 10.1093/mnras/stac2899

Izzard, R. G., Preece, H., Jofre, P., et al. 2018, *Binary stars in the Galactic thick disc*, MNRAS, 473, 2984, doi: 10.1093/mnras/stx2355

Izzard, R. G., Tout, C. A., Karakas, A. I., & Pols, O. R. 2004, *A new synthetic model for asymptotic giant branch stars*, MNRAS, 350, 407, doi: 10.1111/j.1365-2966.2004.07446.x

Janka, H.-T. 2012, *Explosion Mechanisms of Core-Collapse Supernovae*, Annual Review of Nuclear and Particle Science, 62, 407, doi: 10.1146/annurev-nucl-102711-094901

—. 2013, *Natal kicks of stellar mass black holes by asymmetric mass ejection in fallback supernovae*, MNRAS, 434, 1355, doi: 10.1093/mnras/stt1106

—. 2017, *Neutron Star Kicks by the Gravitational Tug-boat Mechanism in Asymmetric Supernova Explosions: Progenitor and Explosion Dependence*, ApJ, 837, 84, doi: 10.3847/1538-4357/aa618e

Janka, H. T., & Mueller, E. 1994, *Neutron star recoils from anisotropic supernovae.*, A&A, 290, 496

Jayasinghe, T., Stanek, K. Z., Thompson, T. A., et al. 2021, *A unicorn in monoceros: the 3 $M_{\odot}$ dark companion to the bright, nearby red giant V723 Mon is a non-interacting, mass-gap black hole candidate*, MNRAS, 504, 2577

Jeans, J. H. 1924, *Cosmogonic problems associated with a secular decrease of mass*, MNRAS, 85, 2, doi: 10.1093/mnras/85.1.2

—. 1925, *The effect of varying mass on a binary system*, MNRAS, 85, 912, doi: 10.1093/mnras/85.9.912

Jermyn, A. S., Schwab, J., Bauer, E., Timmes, F. X., & Potekhin, A. Y. 2021, *Skye: A Differentiable Equation of State*, ApJ, 913, 72, doi: 10.3847/1538-4357/abf48e

Jermyn, A. S., Bauer, E. B., Schwab, J., et al. 2023, *Modules for Experiments in Stellar Astrophysics (MESA): Time-dependent Convection, Energy Conservation, Automatic Differentiation, and Infrastructure*, ApJS, 265, 15, doi: 10.3847/1538-4365/acae8d

Jeřábková, T., Zonoozi, A. H., Kroupa, P., et al. 2018, *Impact of metallicity and star formation rate on the time-dependent, galaxy-wide stellar initial mass function*, A&A, 620, A39, doi: 10.1051/0004-6361/201833055

Jiang, Y.-F., & Tremaine, S. 2010, *The evolution of wide binary stars*, MNRAS, 401, 977, doi: 10.1111/j.1365-2966.2009.15744.x

Johnston, C. 2021, *One size does not fit all: Evidence for a range of mixing efficiencies in stellar evolution calculations*, A&A, 655, A29, doi: 10.1051/0004-6361/202141080

Johnston, C., Pavlovski, K., & Tkachenko, A. 2019a, *Modelling of the B-type binaries CW Cephei and U Ophiuchi. A critical view on dynamical masses, core boundary mixing, and core mass*, A&A, 628, A25, doi: 10.1051/0004-6361/201935235

Johnston, C., Tkachenko, A., Aerts, C., et al. 2019b, *Binary asteroseismic modelling: isochrone-cloud methodology and application to Kepler gravity mode pulsators*, MNRAS, 482, 1231, doi: 10.1093/mnras/sty2671

Kaiser, E. A., Hirschi, R., Arnett, W. D., et al. 2020, *Relative importance of convective uncertainties in massive stars*, MNRAS, 496, 1967, doi: 10.1093/mnras/staa1595

Kalogera, V. 1996, *Orbital Characteristics of Binary Systems after Asymmetric Supernova Explosions*, ApJ, 471, 352, doi: 10.1086/177974

Kalogera, V., & Baym, G. 1996, *The Maximum Mass of a Neutron Star*, ApJ, 470, L61, doi: 10.1086/310296

Kalogera, V., Belczynski, K., Kim, C., O'Shaughnessy, R., & Willems, B. 2007, *Formation of double compact objects*, Phys. Rep., 442, 75, doi: 10.1016/j.physrep.2007.02.008

Katsuda, S., Morii, M., Janka, H.-T., et al. 2018, *Intermediate-mass Elements in Young Supernova Remnants Reveal Neutron Star Kicks by Asymmetric Explosions*, ApJ, 856, 18, doi: 10.3847/1538-4357/aab092

Katz, J. I. 1975, *Two kinds of stellar collapse*, Nature, 253, 698, doi: 10.1038/253698a0

Katz, N., Weinberg, D. H., & Hernquist, L. 1996, *Cosmological Simulations with TreeSPH*, ApJS, 105, 19, doi: 10.1086/192305

Keane, E., Bhattacharyya, B., Kramer, M., et al. 2015, in Advancing Astrophysics with the Square Kilometre Array (AASKA14), 40. https://arxiv.org/abs/1501.00056

Keith, B., Munshi, F., Brooks, A. M., et al. 2025, *A MARVEL-ous study of how well galaxy shapes reflect Dark Matter halo shapes in Cold Dark Matter Simulations*, arXiv e-prints, arXiv:2501.16317, doi: 10.48550/arXiv.2501.16317

Kiel, P. D., & Hurley, J. R. 2009, *Populating the Galaxy with pulsars - II. Galactic dynamics*, MNRAS, 395, 2326, doi: 10.1111/j.1365-2966.2009.14711.x

Kim, C.-G., & Ostriker, E. C. 2017, *Three-phase Interstellar Medium in Galaxies Resolving Evolution with Star Formation and Supernova Feedback (TIGRESS): Algorithms, Fiducial Model, and Convergence*, ApJ, 846, 133, doi: 10.3847/1538-4357/aa8599

Kim, C.-G., Ostriker, E. C., Somerville, R. S., et al. 2020, *First Results from SMAUG: Characterization of Multiphase Galactic Outflows from a Suite of Local Star-forming Galactic Disk Simulations*, ApJ, 900, 61, doi: 10.3847/1538-4357/aba962

Kimm, T., Cen, R., Devriendt, J., Dubois, Y., & Slyz, A. 2015, *Towards simulating star formation in turbulent high-z galaxies with mechanical supernova feedback*, MNRAS, 451, 2900, doi: 10.1093/mnras/stv1211

Kippenhahn, R. 1969, *Mass Exchange in a Massive Close Binary System*, A&A, 3, 83

Kippenhahn, R., Ruschenplatt, G., & Thomas, H. C. 1980, *The time scale of thermohaline mixing in stars*, A&A, 91, 175

Kippenhahn, R., & Weigert, A. 1967, *Entwicklung in engen Doppelsternsystemen I. Massenaustausch vor und nach Beendigung des zentralen Wasserstoff-Brennens*, ZAp, 65, 251

Klein, A., Barausse, E., Sesana, A., et al. 2016, *Science with the space-based interferometer eLISA: Supermassive black hole binaries*, Phys. Rev. D, 93, 024003, doi: 10.1103/PhysRevD.93.024003

Klencki, J., Istrate, A., Nelemans, G., & Pols, O. 2022, *Partial-envelope stripping and nuclear-timescale mass transfer from evolved supergiants at low metallicity*, A&A, 662, A56, doi: 10.1051/0004-6361/202142701

Klencki, J., Moe, M., Gladysz, W., et al. 2018, *Impact of inter-correlated initial binary parameters on double black hole and neutron star mergers*, A&A, 619, A77

Klencki, J., Nelemans, G., Istrate, A. G., & Chruslinska, M. 2021, *It has to be cool: Supergiant progenitors of binary black hole mergers from common-envelope evolution*, A&A, 645, A54, doi: 10.1051/0004-6361/202038707

Kobulnicky, H. A., & Fryer, C. L. 2007, *A New Look at the Binary Characteristics of Massive Stars*, ApJ, 670, 747, doi: 10.1086/522073

Kolb, U., & Ritter, H. 1990, *A comparative study of the evolution of a close binary using a standard and an improved technique for computing mass transfer.*, A&A, 236, 385

Korol, V., Koop, O., & Rossi, E. M. 2018, *Detectability of Double White Dwarfs in the Local Group with LISA*, ApJ, 866, L20, doi: 10.3847/2041-8213/aae587

Korol, V., Rossi, E. M., & Barausse, E. 2019, *A multimessenger study of the Milky Way's stellar disc and bulge with LISA, Gaia, and LSST*, MNRAS, 483, 5518, doi: 10.1093/mnras/sty3440

Korol, V., Rossi, E. M., Groot, P. J., et al. 2017, *Prospects for detection of detached double white dwarf binaries with Gaia, LSST and LISA*, MNRAS, 470, 1894, doi: 10.1093/mnras/stx1285

Korol, V., & Safarzadeh, M. 2021, *How can LISA probe a population of GW190425-like binary neutron stars in the Milky Way?*, MNRAS, 502, 5576, doi: 10.1093/mnras/stab310

Korol, V., Toonen, S., Klein, A., et al. 2020, *Populations of double white dwarfs in Milky Way satellites and their detectability with LISA*, A&A, 638, A153, doi: 10.1051/0004-6361/202037764

Kreidberg, L., Bailyn, C. D., Farr, W. M., & Kalogera, V. 2012, *Mass Measurements of Black Holes in X-Ray Transients: Is There a Mass Gap?*, ApJ, 757, 36, doi: 10.1088/0004-637X/757/1/36

Kremer, K., Breivik, K., Larson, S. L., & Kalogera, V. 2017, *Accreting Double White Dwarf Binaries: Implications for LISA*, ApJ, 846, 95, doi: 10.3847/1538-4357/aa8557

Kremer, K., Chatterjee, S., Breivik, K., et al. 2018, *LISA Sources in Milky Way Globular Clusters*, Phys. Rev. Lett., 120, 191103, doi: 10.1103/PhysRevLett.120.191103

Kremer, K., Lombardi, J. C., Lu, W., Piro, A. L., & Rasio, F. A. 2022, *Hydrodynamics of Collisions and Close Encounters between Stellar Black Holes and Main-sequence Stars*, ApJ, 933, 203, doi: 10.3847/1538-4357/ac714f

Kroupa, P. 1995, *Inverse dynamical population synthesis and star formation*, MNRAS, 277, 1491, doi: 10.1093/mnras/277.4.1491

—. 2001, *On the variation of the initial mass function*, MNRAS, 322, 231, doi: 10.1046/j.1365-8711.2001.04022.x

Kroupa, P., Jeřábková, T., Dinnbier, F., Beccari, G., & Yan, Z. 2018, *Evidence for feedback and stellar-dynamically regulated bursty star cluster formation: the case of the Orion Nebula Cluster*, A&A, 612, A74, doi: 10.1051/0004-6361/201732151

Krtička, J., & Kubát, J. 2017, *Comoving frame models of hot star winds. II. Reduction of O star wind mass-loss rates in global models*, A&A, 606, A31, doi: 10.1051/0004-6361/201730723

—. 2018, *Global hot-star wind models for stars from Magellanic Clouds*, A&A, 612, A20, doi: 10.1051/0004-6361/201731969

Krtička, J., Kubát, J., & Krtičková, I. 2024, *New mass-loss rates of Magellanic Cloud B supergiants from global wind models*, A&A, 681, A29, doi: 10.1051/0004-6361/202347916

Kruckow, M. U., Tauris, T. M., Langer, N., Kramer, M., & Izzard, R. G. 2018, *Progenitors of gravitational wave mergers: binary evolution with the stellar grid-based code COMBINE*, MNRAS, 481, 1908, doi: 10.1093/mnras/sty2190

Krumholz, M. R., McKee, C. F., & Bland-Hawthorn, J. 2019, *Star Clusters Across Cosmic Time*, ARA&A, 57, 227, doi: 10.1146/annurev-astro-091918-104430

Kuhn, M. A., Hillenbrand, L. A., Sills, A., Feigelson, E. D., & Getman, K. V. 2019, *Kinematics in Young Star Clusters and Associations with Gaia DR2*, ApJ, 870, 32, doi: 10.3847/1538-4357/aaef8c

Kuiper, G. P. 1941, *On the Interpretation of β Lyrae and Other Close Binaries.*, ApJ, 93, 133, doi: 10.1086/144252

Kummer, F. 2020, *Master thesis*, UvA Scripties

Kupfer, T., Korol, V., Shah, S., et al. 2018, *LISA verification binaries with updated distances from Gaia Data Release 2*, MNRAS, 480, 302, doi: 10.1093/mnras/sty1545

Kurtz, D. W. 2022, *Asteroseismology Across the Hertzsprung-Russell Diagram*, ARA&A, 60, 31, doi: 10.1146/annurev-astro-052920-094232

Kyutoku, K., Kiuchi, K., Sekiguchi, Y., Shibata, M., & Taniguchi, K. 2018, *Neutrino transport in black hole-neutron star binaries: Neutrino emission and dynamical mass ejection*, Phys. Rev. D, 97, 023009, doi: 10.1103/PhysRevD.97.023009

Kyutoku, K., Nishino, Y., & Seto, N. 2019, *How to detect the shortest period binary pulsars in the era of LISA*, MNRAS, 483, 2615, doi: 10.1093/mnras/sty3322

Labadie-Bartz, J., Carciofi, A. C., Henrique de Amorim, T., et al. 2022, *Classifying Be Star Variability With TESS. I. The Southern Ecliptic*, AJ, 163, 226, doi: 10.3847/1538-3881/ac5abd

Lada, C. J., & Lada, E. A. 2003, *Embedded Clusters in Molecular Clouds*, ARA&A, 41, 57, doi: 10.1146/annurev.astro.41.011802.094844

Lamb, J. B., Oey, M. S., Segura-Cox, D. M., et al. 2016, *The Runaways and Isolated O-Type Star Spectroscopic Survey of the SMC (RIOTS4)*, ApJ, 817, 113, doi: 10.3847/0004-637X/817/2/113

Lamberts, A., Blunt, S., Littenberg, T. B., et al. 2019, *Predicting the LISA white dwarf binary population in the Milky Way with cosmological simulations*, MNRAS, 490, 5888, doi: 10.1093/mnras/stz2834

Lamberts, A., Garrison-Kimmel, S., Hopkins, P. F., et al. 2018, *Predicting the binary black hole population of the Milky Way with cosmological simulations*, MNRAS, 480, 2704, doi: 10.1093/mnras/sty2035

Landri, C., Ricker, P. M., Renzo, M., Rau, S., & Vigna-Gómez, A. 2025, *The Effect of Donor Star Rejuvenation on Common Envelope Evolution*, ApJ, 979, 57, doi: 10.3847/1538-4357/ad9d3c

Langanke, K., & Martínez-Pinedo, G. 2000, *Shell-model calculations of stellar weak interaction rates: II. Weak rates for nuclei in the mass range /A=45-65 in supernovae environments*, Nuclear Physics A, 673, 481, doi: 10.1016/S0375-9474(00)00131-7

Langer, N. 1998, *Coupled mass and angular momentum loss of massive main sequence stars*, A&A, 329, 551

—. 2012, *Presupernova Evolution of Massive Single and Binary Stars*, ARA&A, 50, 107

Langer, N., Fricke, K. J., & Sugimoto, D. 1983, *Semiconvective diffusion and energy transport*, A&A, 126, 207

Laplace, E., Götberg, Y., de Mink, S. E., Justham, S., & Farmer, R. 2020, *The expansion of stripped-envelope stars: Consequences for supernovae and gravitational-wave progenitors*, A&A, 637, A6

Laplace, E., Justham, S., Renzo, M., et al. 2021, *Different to the core: The pre-supernova structures of massive single and binary-stripped stars*, A&A, 656, A58, doi: 10.1051/0004-6361/202140506

Laplace, E., Schneider, F. R. N., & Podsiadlowski, P. 2025, *It's written in the massive stars: The role of stellar physics in the formation of black holes*, A&A, 695, A71, doi: 10.1051/0004-6361/202451077

Lau, H. H. B., Izzard, R. G., & Schneider, F. R. N. 2014, *Numerical tests of rotational mixing in massive stars with the new population synthesis code BONNFIRES*, A&A, 570, A125, doi: 10.1051/0004-6361/201424161

Lau, M. Y. M., Hirai, R., González-Bolívar, M., et al. 2022, *Common envelopes in massive stars: towards the role of radiation pressure and recombination energy in ejecting red supergiant envelopes*, MNRAS, 512, 5462, doi: 10.1093/mnras/stac049

Lau, M. Y. M., Hirai, R., Mandel, I., & Tout, C. A. 2024, *Expansion of Accreting Main-sequence Stars during Rapid Mass Transfer*, ApJ, 966, L7, doi: 10.3847/2041-8213/ad3d50

Lau, M. Y. M., Mandel, I., Vigna-Gómez, A., et al. 2020, *Detecting double neutron stars with LISA*, MNRAS, 492, 3061, doi: 10.1093/mnras/staa002

Lauterborn, D. 1970, *Evolution with mass exchange of case C for a binary system of total mass 7 M sun.*, A&A, 7, 150

Law-Smith, J. A. P., Everson, R. W., Ramirez-Ruiz, E., et al. 2020, *Successful Common Envelope Ejection and Binary Neutron Star Formation in 3D Hydrodynamics*, arXiv e-prints, arXiv:2011.06630, doi: 10.48550/arXiv.2011.06630

Ledoux, P. 1947, *Stellar Models with Convection and with Discontinuity of the Mean Molecular Weight*, ApJ, 105, 305, doi: 10.1086/144905

Leitherer, C., Ekström, S., Meynet, G., et al. 2014, *The Effects of Stellar Rotation. II. A Comprehensive Set of Starburst99 Models*, ApJS, 212, 14, doi: 10.1088/0067-0049/212/1/14

Leitherer, C., Robert, C., & Drissen, L. 1992, *Deposition of Mass, Momentum, and Energy by Massive Stars into the Interstellar Medium*, ApJ, 401, 596, doi: 10.1086/172089

Leitherer, C., Schaerer, D., Goldader, J. D., et al. 1999, *Starburst99: Synthesis Models for Galaxies with Active Star Formation*, ApJS, 123, 3, doi: 10.1086/313233

Leroy, A. K., Schinnerer, E., Hughes, A., et al. 2021, *PHANGS-ALMA: Arcsecond CO(2-1) Imaging of Nearby Star-forming Galaxies*, ApJS, 257, 43, doi: 10.3847/1538-4365/ac17f3

Levesque, E. M., & Lamers, H. J. 2017, Understanding Stellar Evolution, 2514-3433 (IOP Publishing), doi: `10.1088/978-0-7503-1278-3`

Li, L.-X., & Paczyński, B. 1998, *Transient Events from Neutron Star Mergers*, ApJ, 507, L59, doi: `10.1086/311680`

Li, Y.-J., Wang, Y.-Z., Han, M.-Z., et al. 2021, *A Flexible Gaussian Process Reconstruction Method and the Mass Function of the Coalescing Binary Black Hole Systems*, ApJ, 917, 33, doi: `10.3847/1538-4357/ac0971`

Licquia, T. C., & Newman, J. A. 2015, *Improved Estimates of the Milky Way's Stellar Mass and Star Formation Rate from Hierarchical Bayesian Meta-Analysis*, ApJ, 806, 96, doi: `10.1088/0004-637X/806/1/96`

Lin, D. N. C., & Papaloizou, J. 1979, *Tidal torques on accretion discs in binary systems with extreme mass ratios.*, MNRAS, 186, 799, doi: `10.1093/mnras/186.4.799`

Liu, J. 2009, *Gravitational wave radiation from close double white dwarfs in the Galaxy*, MNRAS, 400, 1850, doi: `10.1111/j.1365-2966.2009.15574.x`

Liu, J., Soria, R., Zheng, Z., et al. 2020, *Reply to: On the signature of a 70-solar-mass black hole in LB-1*, Nature, 580, E16, doi: `10.1038/s41586-020-2217-9`

Liu, J., Zhang, H., Howard, A. W., et al. 2019, *A wide star-black-hole binary system from radial-velocity measurements*, Nature, 575, 618, doi: `10.1038/s41586-019-1766-2`

Liu, J., & Zhang, Y. 2014, *Gravitational-wave radiation from double compact objects with eLISA in the Galaxy*, PASP, 126, 211, doi: `10.1086/675721`

Liu, Z.-W., Röpke, F. K., & Han, Z. 2023, *Type Ia Supernova Explosions in Binary Systems: A Review*, Research in Astronomy and Astrophysics, 23, 082001, doi: `10.1088/1674-4527/acd89e`

Liu, Z.-W., Tauris, T. M., Röpke, F. K., et al. 2015, *The interaction of core-collapse supernova ejecta with a companion star*, A&A, 584, A11, doi: `10.1051/0004-6361/201526757`

Lovegrove, E., & Woosley, S. E. 2013, *Very Low Energy Supernovae from Neutrino Mass Loss*, ApJ, 769, 109, doi: `10.1088/0004-637X/769/2/109`

Lowing, B., Jenkins, A., Eke, V., & Frenk, C. 2011, *A halo expansion technique for approximating simulated dark matter haloes*, MNRAS, 416, 2697, doi: `10.1111/j.1365-2966.2011.19222.x`

Lucy, L. B., & Solomon, P. M. 1970, *Mass Loss by Hot Stars*, ApJ, 159, 879, doi: `10.1086/150365`

Lyne, A. G., & Lorimer, D. R. 1994, *High birth velocities of radio pulsars*, Nature, 369, 127, doi: `10.1038/369127a0`

Maeder, A., & Meynet, G. 2000, *The Evolution of Rotating Stars*, ARA&A, 38, 143, doi: `10.1146/annurev.astro.38.1.143`

Majewski, S. R., Schiavon, R. P., Frinchaboy, P. M., et al. 2017, *The Apache Point Observatory Galactic Evolution Experiment (APOGEE)*, AJ, 154, 94, doi: `10.3847/1538-3881/aa784d`

Manchester, R. N., Hobbs, G. B., Teoh, A., & Hobbs, M. 2005, *The Australia Telescope National Facility Pulsar Catalogue*, AJ, 129, 1993, doi: `10.1086/428488`

Mandel, I. 2016, *Estimates of black hole natal kick velocities from observations of low-mass X-ray binaries*, MNRAS, 456, 578, doi: `10.1093/mnras/stv2733`

—. 2021, *An Accurate Analytical Fit to the Gravitational-wave Inspiral Duration for Eccentric Binaries*, Research Notes of the American Astronomical Society, 5, 223, doi: `10.3847/2515-5172/ac2d35`

Mandel, I., & de Mink, S. E. 2016, *Merging binary black holes formed through chemically homogeneous evolution in short-period stellar binaries*, MNRAS, 458, 2634, doi: `10.1093/mnras/stw379`

Mandel, I., & Müller, B. 2020, *Simple recipes for compact remnant masses and natal kicks*, MNRAS, 499, 3214, doi: `10.1093/mnras/staa3043`

Mandel, I., Sesana, A., & Vecchio, A. 2018, *The astrophysical science case for a decihertz gravitational-wave detector*, Classical and Quantum Gravity, 35, 054004, doi: `10.1088/1361-6382/aaa7e0`

Mapelli, M. 2021, Formation Channels of Single and Binary Stellar-Mass Black Holes (Springer), 4, doi: `10.1007/978-981-15-4702-7_16-1`

Marchant, P., & Bodensteiner, J. 2024, *The Evolution of Massive Binary Stars*, ARA&A, 62, 21, doi: `10.1146/annurev-astro-052722-105936`

Marchant, P., Langer, N., Podsiadlowski, P., et al. 2017, *Ultra-luminous X-ray sources and neutron-star-black-hole mergers from very massive close binaries at low metallicity*, A&A, 604, A55, doi: `10.1051/0004-6361/201630188`

Marchant, P., Langer, N., Podsiadlowski, P., Tauris, T. M., & Moriya, T. J. 2016, *A new route towards merging massive black holes*, A&A, 588, A50, doi: `10.1051/0004-6361/201628133`

Marchant, P., Pappas, K. M. W., Gallegos-Garcia, M., et al. 2021, *The role of mass transfer and common envelope evolution in the formation of merging binary black holes*, A&A, 650, A107, doi: `10.1051/0004-6361/202039992`

Marchant, P., Renzo, M., Farmer, R., et al. 2019, *Pulsational Pair-instability Supernovae in Very Close Binaries*, ApJ, 882, 36, doi: `10.3847/1538-4357/ab3426`

Margalit, B., & Metzger, B. D. 2017, *Constraining the Maximum Mass of Neutron Stars from Multi-messenger Observations of GW170817*, ApJ, 850, L19, doi: `10.3847/2041-8213/aa991c`

Marks, M., Kroupa, P., Dabringhausen, J., & Pawlowski, M. S. 2012, *Evidence for top-heavy stellar initial mass functions with increasing density and decreasing metallicity*, MNRAS, 422, 2246, doi: `10.1111/j.1365-2966.2012.20767.x`

Marsh, T. R., Nelemans, G., & Steeghs, D. 2004, *Mass transfer between double white dwarfs*, MNRAS, 350, 113, doi: `10.1111/j.1365-2966.2004.07564.x`

Mason, B. D., Hartkopf, W. I., Gies, D. R., Henry, T. J., & Helsel, J. W. 2009, *The High Angular Resolution Multiplicity of Massive Stars*, AJ, 137, 3358, doi: `10.1088/0004-6256/137/2/3358`

Massevitch, A., & Yungelson, L. 1975, *On the evolution of close binaries with mass and momentum loss from the system*, Mem. Soc. Astron. Italiana, 46, 217

Massey, P., Neugent, K. F., Levesque, E. M., Drout, M. R., & Courteau, S. 2021, *The Red Supergiant Content of M31 and M33*, AJ, 161, 79, doi: `10.3847/1538-3881/abd01f`

Mayker Chen, N., Leroy, A. K., Lopez, L. A., et al. 2023, *Comparing the Locations of Supernovae to CO (2-1) Emission in Their Host Galaxies*, ApJ, 944, 110, doi: `10.3847/1538-4357/acab00`

Mazeh, T., Goldberg, D., Duquennoy, A., & Mayor, M. 1992, *On the Mass-Ratio Distribution of Spectroscopic Binaries with Solar-Type Primaries*, ApJ, 401, 265, doi: 10.1086/172058

McClelland, L. A. S., & Eldridge, J. J. 2016, *Helium stars: towards an understanding of Wolf-Rayet evolution*, MNRAS, 459, 1505, doi: 10.1093/mnras/stw618

McKernan, B., Ford, K. E. S., & O'Shaughnessy, R. 2020, *Black hole, neutron star, and white dwarf merger rates in AGN discs*, MNRAS, 498, 4088, doi: 10.1093/mnras/staa2681

McMillan, P. J. 2011, *Mass models of the Milky Way*, MNRAS, 414, 2446, doi: 10.1111/j.1365-2966.2011.18564.x

Menon, H., Wesolowski, L., Zheng, G., et al. 2015, *Adaptive techniques for clustered N-body cosmological simulations*, Computational Astrophysics and Cosmology, 2, 1, doi: 10.1186/s40668-015-0007-9

Metzger, B. D. 2017, *Kilonovae*, Living Reviews in Relativity, 20, 3, doi: 10.1007/s41114-017-0006-z

Michielsen, M., Aerts, C., & Bowman, D. M. 2021, *Probing the temperature gradient in the core boundary layer of stars with gravito-inertial modes. The case of KIC 7760680*, A&A, 650, A175, doi: 10.1051/0004-6361/202039926

Miglio, A., Montalbán, J., Eggenberger, P., & Noels, A. 2008, *Gravity modes and mixed modes as probes of stellar cores in main-sequence stars: From solar-like to β Cep stars*, Astronomische Nachrichten, 329, 529, doi: 10.1002/asna.200710991

Miller, M. C., & Lauburg, V. M. 2009, *Mergers of Stellar-Mass Black Holes in Nuclear Star Clusters*, ApJ, 692, 917, doi: 10.1088/0004-637X/692/1/917

Miller-Jones, J. C. A., Bahramian, A., Orosz, J. A., et al. 2021, *Cygnus X-1 contains a 21–solar mass black hole—Implications for massive star winds*, Science, 371, 1046, doi: 10.1126/science.abb3363

Misra, D., Kovlakas, K., Fragos, T., et al. 2023, *X-ray luminosity function of high-mass X-ray binaries: Studying the signatures of different physical processes using detailed binary evolution calculations*, A&A, 672, A99, doi: 10.1051/0004-6361/202244929

Miszuda, A., Kołaczek-Szymański, P. A., Szewczuk, W., & Daszyńska-Daszkiewicz, J. 2022, *The eclipsing binary systems with δ Scuti component - II. AB Cas*, MNRAS, 514, 622, doi: 10.1093/mnras/stac1197

Miszuda, A., Szewczuk, W., & Daszyńska-Daszkiewicz, J. 2021, *The eclipsing binary systems with δ Scuti component - I. KIC 10661783*, MNRAS, 505, 3206, doi: 10.1093/mnras/stab1597

Mitalas, R. 1972, *On Semiconvection*, ApJ, 177, 693, doi: 10.1086/151748

Modak, S., & Hamilton, C. 2023, *Eccentricity dynamics of wide binaries - I. The effect of Galactic tides*, MNRAS, 524, 3102, doi: 10.1093/mnras/stad2073

Moe, M., & Di Stefano, R. 2017, *Mind Your Ps and Qs: The Interrelation between Period (P) and Mass-ratio (Q) Distributions of Binary Stars*, ApJS, 230, 15, doi: 10.3847/1538-4365/aa6fb6

Mombarg, J. S. G. 2023, *Calibrating angular momentum transport in intermediate-mass stars from gravity-mode asteroseismology*, A&A, 677, A63, doi: 10.1051/0004-6361/202345956

Mombarg, J. S. G., Van Reeth, T., Pedersen, M. G., et al. 2019, *Asteroseismic masses, ages, and core properties of γ Doradus stars using gravito-inertial dipole modes and spectroscopy*, MNRAS, 485, 3248, doi: 10.1093/mnras/stz501

Moore, C. J., Chua, A. J. K., & Gair, J. R. 2017, *Gravitational waves from extreme mass ratio inspirals around bumpy black holes*, Classical and Quantum Gravity, 34, 195009, doi: 10.1088/1361-6382/aa85fa

Moore, C. J., Cole, R. H., & Berry, C. P. L. 2015, *Gravitational-wave sensitivity curves*, Classical and Quantum Gravity, 32, 015014, doi: 10.1088/0264-9381/32/1/015014

Morris, M. 1993, *Massive Star Formation near the Galactic Center and the Fate of the Stellar Remnants*, ApJ, 408, 496, doi: 10.1086/172607

Morton, D. C. 1960, *Evolutionary Mass Exchange in Close Binary Systems.*, ApJ, 132, 146, doi: 10.1086/146908

—. 1967, *Mass Loss from Three OB Supergiants in Orion*, ApJ, 150, 535, doi: 10.1086/149356

Morton, T. D. 2015, *isochrones: Stellar model grid package*, Astrophysics Source Code Library, record ascl:1503.010

Moyano, F. D., Eggenberger, P., & Salmon, S. J. A. J. 2024, *Angular momentum transport near convective-core boundaries of Gamma Doradus stars*, A&A, 681, L16, doi: 10.1051/0004-6361/202348704

Munari, U., Henden, A., Kiyota, S., et al. 2002, *The mysterious eruption of V838 Mon*, A&A, 389, L51, doi: 10.1051/0004-6361:20020715

Naab, T., & Ostriker, J. P. 2017, *Theoretical Challenges in Galaxy Formation*, ARA&A, 55, 59, doi: 10.1146/annurev-astro-081913-040019

Nandez, J. L. A., & Ivanova, N. 2016, *Common envelope events with low-mass giants: understanding the energy budget*, MNRAS, 460, 3992, doi: 10.1093/mnras/stw1266

Narayan, R., Piran, T., & Shemi, A. 1991, *Neutron Star and Black Hole Binaries in the Galaxy*, ApJ, 379, L17, doi: 10.1086/186143

Navarro, J. F., Frenk, C. S., & White, S. D. M. 1997, *A Universal Density Profile from Hierarchical Clustering*, ApJ, 490, 493, doi: 10.1086/304888

Neijssel, C. J., Vigna-Gómez, A., Stevenson, S., et al. 2019, *The effect of the metallicity-specific star formation history on double compact object mergers*, MNRAS, 490, 3740, doi: 10.1093/mnras/stz2840

Nelemans, G., Yungelson, L. R., & Portegies Zwart, S. F. 2001, *The gravitational wave signal from the Galactic disk population of binaries containing two compact objects*, A&A, 375, 890, doi: 10.1051/0004-6361:20010683

Neo, S., Miyaji, S., Nomoto, K., & Sugimoto, D. 1977, *Effect of Rapid Mass Accretion onto the Main-Sequence Stars*, PASJ, 29, 249

Neugent, K. F., & Massey, P. 2011, *The Wolf-Rayet Content of M33*, ApJ, 733, 123, doi: 10.1088/0004-637X/733/2/123

—. 2023, *Newly Discovered Wolf-Rayet Stars in M31*, AJ, 166, 68, doi: 10.3847/1538-3881/ace25f

Neuhäuser, R., Gießler, F., & Hambaryan, V. V. 2020, *A nearby recent supernova that ejected the runaway star ζ Oph, the pulsar PSR B1706-16, and ^{60}Fe found on Earth*, MNRAS, 498, 899, doi: 10.1093/mnras/stz2629

Nguyen, C. T., Costa, G., Girardi, L., et al. 2022, *PARSEC V2.0: Stellar tracks and isochrones of low- and intermediate-mass stars with rotation*, A&A, 665, A126, doi: 10.1051/0004-6361/202244166

Nishizawa, A., Berti, E., Klein, A., & Sesana, A. 2016, *eLISA eccentricity measurements as tracers of binary black hole formation*, Phys. Rev. D, 94, 064020, doi: `10.1103/PhysRevD.94.064020`

Nissanke, S., Vallisneri, M., Nelemans, G., & Prince, T. A. 2012, *Gravitational-wave Emission from Compact Galactic Binaries*, ApJ, 758, 131, doi: `10.1088/0004-637X/758/2/131`

Nomoto, K. 1984, *Evolution of 8-10 solar mass stars toward electron capture supernovae. I - Formation of electron-degenerate O + NE + MG cores.*, ApJ, 277, 791, doi: `10.1086/161749`

—. 1987, *Evolution of 8–10* M_{sun} *Stars toward Electron Capture Supernovae. II. Collapse of an O + NE + MG Core*, ApJ, 322, 206, doi: `10.1086/165716`

Nomoto, K., Kobayashi, C., & Tominaga, N. 2013, *Nucleosynthesis in Stars and the Chemical Enrichment of Galaxies*, ARA&A, 51, 457, doi: `10.1146/annurev-astro-082812-140956`

Nugent, A. E., Fong, W.-f., Castrejon, C., et al. 2024, *A Population of Short-duration Gamma-Ray Bursts with Dwarf Host Galaxies*, ApJ, 962, 5, doi: `10.3847/1538-4357/ad17c0`

O'Connor, E., & Ott, C. D. 2011, *Black Hole Formation in Failing Core-Collapse Supernovae*, ApJ, 730, 70, doi: `10.1088/0004-637X/730/2/70`

O'Connor, E., Bollig, R., Burrows, A., et al. 2018, *Global comparison of core-collapse supernova simulations in spherical symmetry*, Journal of Physics G Nuclear Physics, 45, 104001, doi: `10.1088/1361-6471/aadeae`

Oda, T., Hino, M., Muto, K., Takahara, M., & Sato, K. 1994, *Rate Tables for the Weak Processes of sd-Shell Nuclei in Stellar Matter*, Atomic Data and Nuclear Data Tables, 56, 231, doi: `10.1006/adnd.1994.1007`

Offner, S. S. R., Moe, M., Kratter, K. M., et al. 2023, in Astronomical Society of the Pacific Conference Series, Vol. 534, Protostars and Planets VII, ed. S. Inutsuka, Y. Aikawa, T. Muto, K. Tomida, & M. Tamura, 275, doi: `10.48550/arXiv.2203.10066`

Ogata, M., Hirai, R., & Hijikawa, K. 2021, *Observability of inflated companion stars after supernovae in massive binaries*, MNRAS, 505, 2485, doi: `10.1093/mnras/stab1439`

Oh, S., & Kroupa, P. 2016, *Dynamical ejections of massive stars from young star clusters under diverse initial conditions*, A&A, 590, A107, doi: `10.1051/0004-6361/201628233`

Ohlmann, S. T., Röpke, F. K., Pakmor, R., & Springel, V. 2016, *Hydrodynamic Moving-mesh Simulations of the Common Envelope Phase in Binary Stellar Systems*, ApJ, 816, L9, doi: `10.3847/2041-8205/816/1/L9`

Oliva, G. A., & Kuiper, R. 2020, *Modeling disk fragmentation and multiplicity in massive star formation*, A&A, 644, A41, doi: `10.1051/0004-6361/202038103`

Öpik, E. 1924, *Statistical Studies of Double Stars: On the Distribution of Relative Luminosities and Distances of Double Stars in the Harvard Revised Photometry North of Declination -31°*, Publications of the Tartu Astrofizica Observatory, 25, 1

Orr, M. E., Fielding, D. B., Hayward, C. C., & Burkhart, B. 2022, *Bursting Bubbles: Feedback from Clustered Supernovae and the Trade-off Between Turbulence and Outflows*, ApJ, 932, 88, doi: `10.3847/1538-4357/ac6c26`

Orr, M. E., Hayward, C. C., Hopkins, P. F., et al. 2018, *What FIREs up star formation: the emergence of the Kennicutt-Schmidt law from feedback*, MNRAS, 478, 3653, doi: `10.1093/mnras/sty1241`

Ouazzani, R. M., Lignières, F., Dupret, M. A., et al. 2020, *First evidence of inertial modes in* γ *Doradus stars: The core rotation revealed*, A&A, 640, A49, doi: `10.1051/0004-6361/201936653`

Özel, F., Psaltis, D., Narayan, R., & McClintock, J. E. 2010, *The Black Hole Mass Distribution in the Galaxy*, ApJ, 725, 1918, doi: `10.1088/0004-637X/725/2/1918`

Packet, W. 1981, *On the spin-up of the mass accreting component in a close binary system*, A&A, 102, 17

Paczyński, B. 1966, *Evolution of Close Binaries. I.*, Acta Astron., 16, 231

Paczynski, B. 1976, in IAU Symposium, Vol. 73, Structure and Evolution of Close Binary Systems, ed. P. Eggleton, S. Mitton, & J. Whelan, 75

Paczynski, B. 1991, *A Polytropic Model of an Accretion Disk, a Boundary Layer, and a Star*, ApJ, 370, 597, doi: `10.1086/169846`

Paczyński, B., & Sienkiewicz, R. 1972, *Evolution of Close Binaries VIII. Mass Exchange on the Dynamical Time Scale*, Acta Astron., 22, 73

Pamyatnykh, A. A. 1999, *Pulsational Instability Domains in the Upper Main Sequence*, Acta Astron., 49, 119

Parravano, A. 1996, *Star Formation Bursts Driven by Delayed Self-Regulation*, ApJ, 462, 594, doi: `10.1086/177175`

Pascale, R., Binney, J., Nipoti, C., & Posti, L. 2019, *Action-based models for dwarf spheroidal galaxies and globular clusters*, MNRAS, 488, 2423, doi: `10.1093/mnras/stz1617`

Patton, R. A., & Sukhbold, T. 2020, *Towards a realistic explosion landscape for binary population synthesis*, MNRAS, 499, 2803, doi: `10.1093/mnras/staa3029`

Paxton, B. 2023, *Modules for Experiments in Stellar Astrophysics (MESA)*, r23.05.1, Zenodo, doi: `10.5281/zenodo.7983526`

Paxton, B., Bildsten, L., Dotter, A., et al. 2011, *Modules for Experiments in Stellar Astrophysics (MESA)*, ApJS, 192, 3, doi: `10.1088/0067-0049/192/1/3`

Paxton, B., Cantiello, M., Arras, P., et al. 2013, *Modules for Experiments in Stellar Astrophysics (MESA): Planets, Oscillations, Rotation, and Massive Stars*, ApJS, 208, 4, doi: `10.1088/0067-0049/208/1/4`

Paxton, B., Marchant, P., Schwab, J., et al. 2015, *Modules for Experiments in Stellar Astrophysics (MESA): Binaries, Pulsations, and Explosions*, ApJS, 220, 15, doi: `10.1088/0067-0049/220/1/15`

Paxton, B., Schwab, J., Bauer, E. B., et al. 2018, *Modules for Experiments in Stellar Astrophysics (MESA): Convective Boundaries, Element Diffusion, and Massive Star Explosions*, ApJS, 234, 34, doi: `10.3847/1538-4365/aaa5a8`

Paxton, B., Smolec, R., Schwab, J., et al. 2019, *Modules for Experiments in Stellar Astrophysics (MESA): Pulsating Variable Stars, Rotation, Convective Boundaries, and Energy Conservation*, ApJS, 243, 10, doi: `10.3847/1538-4365/ab2241`

Pedersen, M. G. 2022, *On the Diversity of Mixing and Helium Core Masses of B-type Dwarfs from Gravity-mode Asteroseismology*, ApJ, 930, 94, doi: `10.3847/1538-4357/ac5b05`

Pedersen, M. G., Aerts, C., Pápics, P. I., & Rogers, T. M. 2018, *The shape of convective core overshooting from gravity-mode period spacings*, A&A, 614, A128, doi: `10.1051/0004-6361/201732317`

Peters, P. C. 1964, *Gravitational Radiation and the Motion of Two Point Masses*, Physical Review, 136, 1224, doi: `10.1103/PhysRev.136.B1224`

Peters, P. C., & Mathews, J. 1963, *Gravitational Radiation from Point Masses in a Keplerian Orbit*, Physical Review, 131, 435, doi: `10.1103/PhysRev.131.435`

Petrovic, J., Langer, N., & van der Hucht, K. A. 2005, *Constraining the mass transfer in massive binaries through progenitor evolution models of Wolf-Rayet+O binaries*, A&A, 435, 1013, doi: `10.1051/0004-6361:20042368`

Pettini, M., Steidel, C. C., Adelberger, K. L., Dickinson, M., & Giavalisco, M. 2000, *The Ultraviolet Spectrum of MS 1512-CB58: An Insight into Lyman-Break Galaxies*, ApJ, 528, 96, doi: `10.1086/308176`

Pfahl, E., Podsiadlowski, P., & Rappaport, S. 2005, *Relativistic Binary Pulsars with Black Hole Companions*, ApJ, 628, 343, doi: `10.1086/430515`

Pfahl, E., Rappaport, S., & Podsiadlowski, P. 2002, *A Comprehensive Study of Neutron Star Retention in Globular Clusters*, The Astrophysical Journal, 573, 283, doi: `10.1086/340494`

Pfahl, E., Rappaport, S., Podsiadlowski, P., & Spruit, H. 2002, *A New Class of High-Mass X-Ray Binaries: Implications for Core Collapse and Neutron Star Recoil*, ApJ, 574, 364, doi: `10.1086/340794`

Piro, A. L. 2013, *Taking the "Un" out of "Unnovae"*, ApJ, 768, L14, doi: `10.1088/2041-8205/768/1/L14`

Plummer, H. C. 1911, *On the problem of distribution in globular star clusters*, MNRAS, 71, 460, doi: `10.1093/mnras/71.5.460`

Podsiadlowski, P., Joss, P. C., & Hsu, J. J. L. 1992, *Presupernova Evolution in Massive Interacting Binaries*, ApJ, 391, 246, doi: `10.1086/171341`

Podsiadlowski, P., Langer, N., Poelarends, A. J. T., et al. 2004, *The Effects of Binary Evolution on the Dynamics of Core Collapse and Neutron Star Kicks*, ApJ, 612, 1044, doi: `10.1086/421713`

Pol, N., McLaughlin, M., & Lorimer, D. 2021, *Insights into the Milky Way pulsar–black hole population using radio and gravitational wave observations*, arXiv e-prints, arXiv:2109.04512. `https://arxiv.org/abs/2109.04512`

Pols, O. R. 1994, *Case A evolution of massive close binaries: formation of contact systems and possible reversal of the supernova order*, A&A, 290, 119

Pols, O. R., Schröder, K.-P., Hurley, J. R., Tout, C., & Eggleton, P. 1998, *Stellar evolution models for* $Z = 0.0001$ *to 0.03*, MNRAS, 298, 525, doi: `10.1046/j.1365-8711.1998.01658.x`

Pols, O. R., Tout, C. A., Eggleton, P. P., & Han, Z. 1995, *Approximate input physics for stellar modelling*, MNRAS, 274, 964, doi: `10.1093/mnras/274.3.964`

Pontzen, A., Roškar, R., Stinson, G. S., et al. 2013, *pynbody: Astrophysics Simulation Analysis for Python*

Popham, R., & Narayan, R. 1991, *Does Accretion Cease When a Star Approaches Breakup?*, ApJ, 370, 604, doi: `10.1086/169847`

Portegies Zwart, S. F., & McMillan, S. L. W. 2000, *Black Hole Mergers in the Universe*, ApJ, 528, L17, doi: `10.1086/312422`

Portegies Zwart, S. F., McMillan, S. L. W., & Gieles, M. 2010, *Young Massive Star Clusters*, ARA&A, 48, 431, doi: `10.1146/annurev-astro-081309-130834`

Portegies Zwart, S. F., & Verbunt, F. 1996, *Population synthesis of high-mass binaries.*, A&A, 309, 179

Potekhin, A. Y., & Chabrier, G. 2010, *Thermodynamic Functions of Dense Plasmas: Analytic Approximations for Astrophysical Applications*, Contributions to Plasma Physics, 50, 82, doi: `10.1002/ctpp.201010017`

Poutanen, J. 2017, *Rosseland and Flux Mean Opacities for Compton Scattering*, ApJ, 835, 119, doi: `10.3847/1538-4357/835/2/119`

Poveda, A., Ruiz, J., & Allen, C. 1967, *Run-away Stars as the Result of the Gravitational Collapse of Proto-stellar Clusters*, Boletin de los Observatorios Tonantzintla y Tacubaya, 4, 86

Preece, H. P., Vigna-Gómez, A., Rajamuthukumar, A. S., Vynatheya, P., & Klencki, J. 2024, *The Evolution of Massive Stellar Multiplicity in the Field I. Numerical simulations, long-term evolution and final outcomes*, arXiv e-prints, arXiv:2412.14022, doi: `10.48550/arXiv.2412.14022`

Press, W. H., & Thorne, K. S. 1972, *Gravitational-Wave Astronomy*, ARA&A, 10, 335, doi: `10.1146/annurev.aa.10.090172.002003`

Price-Whelan, A., Souchereau, H., Wagg, T., et al. 2024, *adrn/gala: v1.9.1*, v1.9.1, Zenodo, doi: `10.5281/zenodo.13377376`

Price-Whelan, A. M. 2017, *Gala: A Python package for galactic dynamics*, The Journal of Open Source Software, 2, 388, doi: `10.21105/joss.00388`

Quillen, A. C., & Bland-Hawthorn, J. 2008, *When is star formation episodic? A delay differential equation 'negative feedback' model*, MNRAS, 386, 2227, doi: `10.1111/j.1365-2966.2008.13193.x`

Rastello, S., Mapelli, M., Di Carlo, U. N., et al. 2020, *Dynamics of black hole-neutron star binaries in young star clusters*, MNRAS, 497, 1563, doi: `10.1093/mnras/staa2018`

—. 2021, *Dynamics of binary black holes in low-mass young star clusters*, MNRAS, 507, 3612, doi: `10.1093/mnras/stab2355`

Renzo, M., & Götberg, Y. 2021, *Evolution of Accretor Stars in Massive Binaries: Broader Implications from Modeling* ζ *Ophiuchi*, ApJ, 923, 277, doi: `10.3847/1538-4357/ac29c5`

Renzo, M., Zapartas, E., Justham, S., et al. 2023, *Rejuvenated Accretors Have Less Bound Envelopes: Impact of Roche Lobe Overflow on Subsequent Common Envelope Events*, ApJ, 942, L32, doi: `10.3847/2041-8213/aca4d3`

Renzo, M., Zapartas, E., de Mink, S. E., et al. 2019, *Massive runaway and walkaway stars. A study of the kinematical imprints of the physical processes governing the evolution and explosion of their binary progenitors*, A&A, 624, A66, doi: `10.1051/0004-6361/201833297`

Repetto, S., Davies, M. B., & Sigurdsson, S. 2012, *Investigating stellar-mass black hole kicks*, MNRAS, 425, 2799, doi: `10.1111/j.1365-2966.2012.21549.x`

Repetto, S., Igoshev, A. P., & Nelemans, G. 2017, *The Galactic distribution of X-ray binaries and its implications for compact object formation and natal kicks*, MNRAS, 467, 298, doi: 10.1093/mnras/stx027

Repetto, S., & Nelemans, G. 2015, *Constraining the formation of black holes in short-period black hole low-mass X-ray binaries*, MNRAS, 453, 3341, doi: 10.1093/mnras/stv1753

Ricker, G. R., Winn, J. N., Vanderspek, R., et al. 2015, *Transiting Exoplanet Survey Satellite (TESS)*, Journal of Astronomical Telescopes, Instruments, and Systems, 1, 014003, doi: 10.1117/1.JATIS.1.1.014003

Riley, J., Agrawal, P., Barrett, J. W., et al. 2022, *Rapid Stellar and Binary Population Synthesis with COMPAS*, ApJS, 258, 34, doi: 10.3847/1538-4365/ac416c

Robson, T., Cornish, N. J., & Liu, C. 2019, *The construction and use of LISA sensitivity curves*, Classical and Quantum Gravity, 36, 105011, doi: 10.1088/1361-6382/ab1101

Roche, E. 1849, *La figure d'une masse fluide soumise à l'attraction d'un point éloigné, première partie*, Mémoires de la section des sciences, Académie des sciences de Montpellier, 1, 243

—. 1850, *La figure d'une masse fluide soumise à l'attraction d'un point éloigné, deuxième partie*, Mémoires de la section des sciences, Académie des sciences de Montpellier, 1, 333

—. 1851, *La figure d'une masse fluide soumise à l'attraction d'un point éloigné, troisième partie*, Mémoires de la section des sciences, Académie des sciences de Montpellier, 2, 21

Rodriguez, C. L., Amaro-Seoane, P., Chatterjee, S., & Rasio, F. A. 2018, *Post-Newtonian Dynamics in Dense Star Clusters: Highly Eccentric, Highly Spinning, and Repeated Binary Black Hole Mergers*, Phys. Rev. Lett., 120, 151101, doi: 10.1103/PhysRevLett.120.151101

Rodriguez, C. L., Hafen, Z., Grudić, M. Y., et al. 2023, *Great balls of FIRE II: The evolution and destruction of star clusters across cosmic time in a Milky Way-mass galaxy*, MNRAS, 521, 124, doi: 10.1093/mnras/stad578

Rodriguez, C. L., Morscher, M., Pattabiraman, B., et al. 2015, *Binary Black Hole Mergers from Globular Clusters: Implications for Advanced LIGO*, Phys. Rev. Lett., 115, 051101, doi: 10.1103/PhysRevLett.115.051101

Roebber, E., Buscicchio, R., Vecchio, A., et al. 2020, *Milky Way Satellites Shining Bright in Gravitational Waves*, ApJ, 894, L15, doi: 10.3847/2041-8213/ab8ac9

Rogers, F. J., & Nayfonov, A. 2002, *Updated and Expanded OPAL Equation-of-State Tables: Implications for Helioseismology*, ApJ, 576, 1064, doi: 10.1086/341894

Röpke, F. K., & De Marco, O. 2023, *Simulations of common-envelope evolution in binary stellar systems: physical models and numerical techniques*, Living Reviews in Computational Astrophysics, 9, 2, doi: 10.1007/s41115-023-00017-x

Rossi, L. J. 2015, *NIGO: A Numerical Integrator of Galactic Orbits*, Astronomy and Computing, 12, 11, doi: 10.1016/j.ascom.2015.03.008

Ruan, D., Brooks, A. M., Cruz, A., et al. 2025, *Predictions for Detecting a Turndown in the Baryonic Tully Fisher Relation*, arXiv e-prints, arXiv:2503.16607, doi: 10.48550/arXiv.2503.16607

Rui, N. Z., & Fuller, J. 2021, *Asteroseismic fingerprints of stellar mergers*, MNRAS, 508, 1618, doi: 10.1093/mnras/stab2528

Ruiter, A. J., Belczynski, K., Benacquista, M., Larson, S. L., & Williams, G. 2010, *The LISA Gravitational Wave Foreground: A Study of Double White Dwarfs*, ApJ, 717, 1006, doi: 10.1088/0004-637X/717/2/1006

Ruiter, A. J., & Seitenzahl, I. R. 2025, *Type Ia supernova progenitors: a contemporary view of a long-standing puzzle*, A&A Rev., 33, 1, doi: 10.1007/s00159-024-00158-9

Salmon, S. J. A. J., Moyano, F. D., Eggenberger, P., Haemmerlé, L., & Buldgen, G. 2022, *Backtracing the internal rotation history of the β Cep star HD 129929*, A&A, 664, L1, doi: 10.1051/0004-6361/202243961

Salpeter, E. E. 1955, *The Luminosity Function and Stellar Evolution.*, ApJ, 121, 161, doi: 10.1086/145971

Sana, H., de Mink, S. E., de Koter, A., et al. 2012, *Binary Interaction Dominates the Evolution of Massive Stars*, Science, 337, 444, doi: 10.1126/science.1223344

Sana, H., Le Bouquin, J. B., Lacour, S., et al. 2014, *Southern Massive Stars at High Angular Resolution: Observational Campaign and Companion Detection*, ApJS, 215, 15, doi: 10.1088/0067-0049/215/1/15

Sana, H., Ramírez-Agudelo, O. H., Hénault-Brunet, V., et al. 2022, *The VLT-FLAMES Tarantula Survey. Observational evidence for two distinct populations of massive runaway stars in 30 Doradus*, A&A, 668, L5, doi: 10.1051/0004-6361/202244677

Sander, A. A. C., & Vink, J. S. 2020, *On the nature of massive helium star winds and Wolf-Rayet-type mass-loss*, MNRAS, 499, 873, doi: 10.1093/mnras/staa2712

Sanders, J. L., & Binney, J. 2015, *Extended distribution functions for our Galaxy*, MNRAS, 449, 3479, doi: 10.1093/mnras/stv578

Sanderson, R. E., Wetzel, A., Loebman, S., et al. 2020, *Synthetic Gaia Surveys from the FIRE Cosmological Simulations of Milky Way-mass Galaxies*, ApJS, 246, 6, doi: 10.3847/1538-4365/ab5b9d

Sarbadhicary, S. K., Wagner, J., Koch, E. W., et al. 2023, *Where do stars explode in the ISM? – The distribution of dense gas around massive stars and supernova remnants in M33*, arXiv e-prints, arXiv:2310.17694, doi: 10.48550/arXiv.2310.17694

Saumon, D., Chabrier, G., & van Horn, H. M. 1995, *An Equation of State for Low-Mass Stars and Giant Planets*, ApJS, 99, 713, doi: 10.1086/192204

Schaller, G., Schaerer, D., Meynet, G., & Maeder, A. 1992, *New Grids of Stellar Models from 0.8-SOLAR-MASS to 120-SOLAR-MASSES at Z=0.020 and Z=0.001*, A&AS, 96, 269

Scheck, L., Kifonidis, K., Janka, H. T., & Müller, E. 2006, *Multidimensional supernova simulations with approximative neutrino transport. I. Neutron star kicks and the anisotropy of neutrino-driven explosions in two spatial dimensions*, A&A, 457, 963, doi: 10.1051/0004-6361:20064855

Schneider, F. R. N., Izzard, R. G., Langer, N., & de Mink, S. E. 2015, *Evolution of Mass Functions of Coeval Stars through Wind Mass Loss and Binary Interactions*, ApJ, 805, 20, doi: 10.1088/0004-637X/805/1/20

Schneider, F. R. N., Podsiadlowski, P., Langer, N., Castro, N., & Fossati, L. 2016, *Rejuvenation of stellar mergers and the origin of magnetic fields in massive stars*, MNRAS, 457, 2355, doi: 10.1093/mnras/stw148

Schneider, F. R. N., Ramírez-Agudelo, O. H., Tramper, F., et al. 2018a, *The VLT-FLAMES Tarantula Survey. XXIX. Massive star formation in the local 30 Doradus starburst*, A&A, 618, A73, doi: 10.1051/0004-6361/201833433

Schneider, F. R. N., Sana, H., Evans, C. J., et al. 2018b, *An excess of massive stars in the local 30 Doradus starburst*, Science, 359, 69, doi: 10.1126/science.aan0106

Schoettler, C., Parker, R. J., & de Bruijne, J. 2022, *Constraining the initial conditions of NGC 2264 using ejected stars found in Gaia DR2*, MNRAS, 510, 3178, doi: 10.1093/mnras/stab3529

Schreier, R., Hillel, S., Shiber, S., & Soker, N. 2021, *Simulating highly eccentric common envelope jet supernova impostors*, MNRAS, 508, 2386, doi: 10.1093/mnras/stab2687

Schuster, E. F. 1985, *Incorporating support constraints into nonparametric estimators of densities*, Communications in Statistics-theory and Methods, 14, 1123

Schwarzschild, M. 1958, Structure and evolution of the stars.

Sekaran, S., Tkachenko, A., Johnston, C., & Aerts, C. 2021, *A comparison of the dynamical and model-derived parameters of the pulsating eclipsing binary KIC 9850387*, A&A, 648, A91, doi: 10.1051/0004-6361/202040154

Sellwood, J. A., & Binney, J. J. 2002, *Radial mixing in galactic discs*, MNRAS, 336, 785, doi: 10.1046/j.1365-8711.2002.05806.x

Sesana, A. 2016, *Prospects for Multiband Gravitational-Wave Astronomy after GW150914*, Phys. Rev. Lett., 116, 231102, doi: 10.1103/PhysRevLett.116.231102

Sesana, A., Lamberts, A., & Petiteau, A. 2020, *Finding binary black holes in the Milky Way with LISA*, MNRAS, 494, L75, doi: 10.1093/mnrasl/slaa039

Shao, Y., & Li, X.-D. 2021, *Population Synthesis of Black Hole Binaries with Compact Star Companions*, ApJ, 920, 81, doi: 10.3847/1538-4357/ac173e

Shapley, A. E., Steidel, C. C., Pettini, M., & Adelberger, K. L. 2003, *Rest-Frame Ultraviolet Spectra of z~3 Lyman Break Galaxies*, ApJ, 588, 65, doi: 10.1086/373922

Sharma, S., Hayden, M. R., & Bland-Hawthorn, J. 2021, *Chemical enrichment and radial migration in the Galactic disc - the origin of the [αFe] double sequence*, MNRAS, 507, 5882, doi: 10.1093/mnras/stab2015

Shenar, T., Bodensteiner, J., Abdul-Masih, M., et al. 2020, *The "hidden" companion in LB-1 unveiled by spectral disentangling*, A&A, 639, L6, doi: 10.1051/0004-6361/202038275

Shenar, T., Sablowski, D. P., Hainich, R., et al. 2019, *The Wolf-Rayet binaries of the nitrogen sequence in the Large Magellanic Cloud. Spectroscopy, orbital analysis, formation, and evolution*, A&A, 627, A151, doi: 10.1051/0004-6361/201935684

Shi, X.-d., Qian, S.-b., & Li, L.-J. 2022, *New Pulsating Stars Detected in EA-type Eclipsing-binary Systems Based on TESS Data*, ApJS, 259, 50, doi: 10.3847/1538-4365/ac59b9

Shklovskii, I. S. 1970, *Possible Causes of the Secular Increase in Pulsar Periods.*, Soviet Ast., 13, 562

Siegel, D. M., Agarwal, A., Barnes, J., et al. 2022, *"Super-kilonovae" from Massive Collapsars as Signatures of Black Hole Birth in the Pair-instability Mass Gap*, ApJ, 941, 100, doi: 10.3847/1538-4357/ac8d04

Sigurdsson, S., & Hernquist, L. 1993, *Primordial black holes in globular clusters*, Nature, 364, 423, doi: 10.1038/364423a0

Silsbee, K., & Tremaine, S. 2017, *Lidov-Kozai Cycles with Gravitational Radiation: Merging Black Holes in Isolated Triple Systems*, ApJ, 836, 39, doi: 10.3847/1538-4357/aa5729

Silva Aguirre, V., Ballot, J., Serenelli, A. M., & Weiss, A. 2011, *Constraining mixing processes in stellar cores using asteroseismology. Impact of semiconvection in low-mass stars*, A&A, 529, A63, doi: 10.1051/0004-6361/201015847

Smak, J. 1962, *Close Binaries. II. A Preliminary Discussion of the Subgiants in Close-binary Systems*, Acta Astron., 12, 28

Smarr, L. L., & Blandford, R. 1976, *The binary pulsar: physical processes, possible companions, and evolutionary histories.*, ApJ, 207, 574, doi: 10.1086/154524

Smith, M. C., Bryan, G. L., Somerville, R. S., et al. 2021, *Efficient early stellar feedback can suppress galactic outflows by reducing supernova clustering*, MNRAS, 506, 3882, doi: 10.1093/mnras/stab1896

Smits, R., Kramer, M., Stappers, B., et al. 2009, *Pulsar searches and timing with the square kilometre array*, A&A, 493, 1161, doi: 10.1051/0004-6361:200810383

Snaith, O. N., Haywood, M., Di Matteo, P., et al. 2014, *The Dominant Epoch of Star Formation in the Milky Way Formed the Thick Disk*, ApJ, 781, L31, doi: 10.1088/2041-8205/781/2/L31

Soberman, G. E., Phinney, E. S., & van den Heuvel, E. P. J. 1997, *Stability criteria for mass transfer in binary stellar evolution.*, A&A, 327, 620. https://arxiv.org/abs/astro-ph/9703016

Sobolev, V. V. 1960, Moving Envelopes of Stars, doi: 10.4159/harvard.9780674864658

Soker, N., & Tylenda, R. 2003, *Main-Sequence Stellar Eruption Model for V838 Monocerotis*, ApJ, 582, L105, doi: 10.1086/367759

Somerville, R. S., & Davé, R. 2015, *Physical Models of Galaxy Formation in a Cosmological Framework*, ARA&A, 53, 51, doi: 10.1146/annurev-astro-082812-140951

Song, Y., Stevenson, S., & Chattopadhyay, D. 2024, *Binary population synthesis of the Galactic canonical pulsar population*, arXiv e-prints, arXiv:2406.11428, doi: 10.48550/arXiv.2406.11428

Spera, M., Mapelli, M., & Bressan, A. 2022, *SEVN: Stellar EVolution for N-body*, Astrophysics Source Code Library, record ascl:2206.019

Srinivasan, G. 1989, *Pulsars: their origin and evolution*, A&A Rev., 1, 209, doi: 10.1007/BF00873079

Staritsin, E. I. 2019, *The spin-up of a star gaining mass in a close binary*, Ap&SS, 364, 110, doi: 10.1007/s10509-019-3600-6

Stegmann, J., & Antonini, F. 2024, *TSE: A triple stellar evolution code*, The Journal of Open Source Software, 9, 7102, doi: 10.21105/joss.07102

Stegmann, J., Vigna-Gómez, A., Rantala, A., et al. 2024, *Close Encounters of Wide Binaries Induced by the Galactic Tide: Implications for Stellar Mergers and Gravitational-wave Sources*, ApJ, 972, L19, doi: 10.3847/2041-8213/ad70bb

Steinwandel, U. P., Bryan, G. L., Somerville, R. S., Hayward, C. C., & Burkhart, B. 2023, *On the impact of runaway stars on dwarf galaxies with resolved interstellar medium*, MNRAS, 526, 1408, doi: 10.1093/mnras/stad2744

Stephan, A. P., Naoz, S., Ghez, A. M., et al. 2016, *Merging binaries in the Galactic Center: the eccentric Kozai-Lidov mechanism with stellar evolution*, MNRAS, 460, 3494, doi: `10.1093/mnras/stw1220`

Stevenson, S., Sampson, M., Powell, J., et al. 2019, *The Impact of Pair-instability Mass Loss on the Binary Black Hole Mass Distribution*, ApJ, 882, 121, doi: `10.3847/1538-4357/ab3981`

Stevenson, S., Vigna-Gómez, A., Mandel, I., et al. 2017, *Formation of the first three gravitational-wave observations through isolated binary evolution*, Nature Communications, 8, 14906, doi: `10.1038/ncomms14906`

Storck, A., & Church, R. P. 2023, *The effects of Galactic model uncertainties on LISA observations of double neutron stars*, MNRAS, 521, 2368, doi: `10.1093/mnras/stad663`

Struck-Marcell, C., & Scalo, J. M. 1987, *Continuum Models for Gas in Disturbed Galaxies. III. Bifurcations and Chaos in a Deterministic Model for Bursts of Star Formation*, ApJS, 64, 39, doi: `10.1086/191191`

Struve, O. 1941, *The Spectrum of β Lyrae.*, ApJ, 93, 104, doi: `10.1086/144249`

Sugimoto, D. 1970, *On the Numerical Stability of Computations of Stellar Evolution*, ApJ, 159, 619, doi: `10.1086/150336`

Sukhbold, T., Ertl, T., Woosley, S. E., Brown, J. M., & Janka, H. T. 2016, *Core-collapse Supernovae from 9 to 120 Solar Masses Based on Neutrino-powered Explosions*, ApJ, 821, 38, doi: `10.3847/0004-637X/821/1/38`

Sukhbold, T., & Woosley, S. E. 2014, *The Compactness of Presupernova Stellar Cores*, ApJ, 783, 10, doi: `10.1088/0004-637X/783/1/10`

Sun, M., Townsend, R. H. D., & Guo, Z. 2023, *gyre_tides: Modeling Binary Tides within the GYRE Stellar Oscillation Code*, ApJ, 945, 43, doi: `10.3847/1538-4357/acb33a`

Sutantyo, W. 1978, *Asymmetric Supernova Explosions and the Origin of Binary Pulsars*, Ap&SS, 54, 479, doi: `10.1007/BF00639450`

Suwa, Y., Yoshida, T., Shibata, M., Umeda, H., & Takahashi, K. 2015, *Neutrino-driven explosions of ultra-stripped Type Ic supernovae generating binary neutron stars*, MNRAS, 454, 3073, doi: `10.1093/mnras/stv2195`

Sweeney, D., Tuthill, P., Sharma, S., & Hirai, R. 2022, *The Galactic underworld: the spatial distribution of compact remnants*, MNRAS, 516, 4971, doi: `10.1093/mnras/stac2092`

Taam, R. E., & Sandquist, E. L. 2000, *Common Envelope Evolution of Massive Binary Stars*, ARA&A, 38, 113, doi: `10.1146/annurev.astro.38.1.113`

Takahashi, R., & Seto, N. 2002, *Parameter Estimation for Galactic Binaries by the Laser Interferometer Space Antenna*, ApJ, 575, 1030, doi: `10.1086/341483`

Tamanini, N., & Danielski, C. 2019, *The gravitational-wave detection of exoplanets orbiting white dwarf binaries using LISA*, Nature Astronomy, 3, 858, doi: `10.1038/s41550-019-0807-y`

Tassoul, M. 1980, *Asymptotic approximations for stellar nonradial pulsations.*, ApJS, 43, 469, doi: `10.1086/190678`

Tauris, T. M., Fender, R. P., van den Heuvel, E. P. J., Johnston, H. M., & Wu, K. 1999, *Circinus X-1: survivor of a highly asymmetric supernova*, MNRAS, 310, 1165, doi: `10.1046/j.1365-8711.1999.03068.x`

Tauris, T. M., Langer, N., Moriya, T. J., et al. 2013, *Ultra-stripped Type Ic Supernovae from Close Binary Evolution*, ApJ, 778, L23, doi: `10.1088/2041-8205/778/2/L23`

Tauris, T. M., Langer, N., & Podsiadlowski, P. 2015, *Ultra-stripped supernovae: progenitors and fate*, MNRAS, 451, 2123, doi: `10.1093/mnras/stv990`

Tauris, T. M., & Takens, R. J. 1998, *Runaway velocities of stellar components originating from disrupted binaries via asymmetric supernova explosions*, A&A, 330, 1047

Tauris, T. M., & van den Heuvel, E. P. J. 2014, *Formation of the Galactic Millisecond Pulsar Triple System PSR J0337+1715—A Neutron Star with Two Orbiting White Dwarfs*, ApJ, 781, L13, doi: `10.1088/2041-8205/781/1/L13`

Tauris, T. M., Kramer, M., Freire, P. C. C., et al. 2017, *Formation of Double Neutron Star Systems*, ApJ, 846, 170, doi: `10.3847/1538-4357/aa7e89`

Thiele, S., Breivik, K., Sanderson, R. E., & Luger, R. 2023, *Applying the Metallicity-dependent Binary Fraction to Double White Dwarf Formation: Implications for LISA*, ApJ, 945, 162, doi: `10.3847/1538-4357/aca7be`

Thorne, K. S., & Zytkow, A. N. 1975, *Red giants and supergiants with degenerate neutron cores.*, ApJ, 199, L19, doi: `10.1086/181839`

Thrane, E., Osłowski, S., & Lasky, P. D. 2020, *Ultrarelativistic astrophysics using multimessenger observations of double neutron stars with LISA and the SKA*, MNRAS, 493, 5408, doi: `10.1093/mnras/staa593`

Timmes, F. X., & Swesty, F. D. 2000, *The Accuracy, Consistency, and Speed of an Electron-Positron Equation of State Based on Table Interpolation of the Helmholtz Free Energy*, ApJS, 126, 501, doi: `10.1086/313304`

Timmes, F. X., Woosley, S. E., & Weaver, T. A. 1996, *The Neutron Star and Black Hole Initial Mass Function*, ApJ, 457, 834, doi: `10.1086/176778`

Tiwari, V. 2022, *Exploring Features in the Binary Black Hole Population*, ApJ, 928, 155, doi: `10.3847/1538-4357/ac589a`

Toonen, S., Claeys, J. S. W., Mennekens, N., & Ruiter, A. J. 2014, *PopCORN: Hunting down the differences between binary population synthesis codes*, A&A, 562, A14, doi: `10.1051/0004-6361/201321576`

Toonen, S., Hamers, A., & Portegies Zwart, S. 2016, *The evolution of hierarchical triple star-systems*, Computational Astrophysics and Cosmology, 3, 6, doi: `10.1186/s40668-016-0019-0`

Toonen, S., & Nelemans, G. 2013, *The effect of common-envelope evolution on the visible population of post-common-envelope binaries*, A&A, 557, A87, doi: `10.1051/0004-6361/201321753`

Toonen, S., Nelemans, G., & Portegies Zwart, S. 2012, *Supernova Type Ia progenitors from merging double white dwarfs. Using a new population synthesis model*, A&A, 546, A70, doi: `10.1051/0004-6361/201218966`

Toonen, S., Portegies Zwart, S., Hamers, A. S., & Bandopadhyay, D. 2020, *The evolution of stellar triples. The most common evolutionary pathways*, A&A, 640, A16, doi: `10.1051/0004-6361/201936835`

Torres, G., Andersen, J., & Giménez, A. 2010, *Accurate masses and radii of normal stars: modern results and applications*, A&A Rev., 18, 67, doi: `10.1007/s00159-009-0025-1`

Tout, C. A., Aarseth, S. J., Pols, O. R., & Eggleton, P. P. 1997, *Rapid binary star evolution for N-body simulations and population synthesis*, MNRAS, 291, 732, doi: `10.1093/mnras/291.4.732`

Tout, C. A., Pols, O. R., Eggleton, P. P., & Han, Z. 1996, *Zero-age main-seqence radii and luminosities as analytic functions of mass and metallicity*, MNRAS, 281, 257, doi: 10.1093/mnras/281.1.257

Townsend, R. H. D. 2003, *A semi-analytical formula for the light variations due to low-frequency g modes in rotating stars*, MNRAS, 343, 125, doi: 10.1046/j.1365-8711.2003.06640.x

Townsend, R. H. D., Goldstein, J., & Zweibel, E. G. 2018, *Angular momentum transport by heat-driven g-modes in slowly pulsating B stars*, MNRAS, 475, 879, doi: 10.1093/mnras/stx3142

Townsend, R. H. D., & Teitler, S. A. 2013, *GYRE: an open-source stellar oscillation code based on a new Magnus Multiple Shooting scheme*, MNRAS, 435, 3406, doi: 10.1093/mnras/stt1533

Tutukov, A., & Yungelson, L. 1973, *Evolution of massive close binaries*, Nauchnye Informatsii, 27, 70

Tutukov, A. V., & Yungelson, L. R. 1993, *The merger rate of neutron star and black hole binaries.*, MNRAS, 260, 675, doi: 10.1093/mnras/260.3.675

Ugliano, M., Janka, H.-T., Marek, A., & Arcones, A. 2012, *Progenitor-explosion Connection and Remnant Birth Masses for Neutrino-driven Supernovae of Iron-core Progenitors*, ApJ, 757, 69, doi: 10.1088/0004-637X/757/1/69

Ugolini, C., Limongi, M., Schneider, R., et al. 2025, *The initial mass-remnant mass relation for core collapse supernovae*, arXiv e-prints, arXiv:2501.18689, doi: 10.48550/arXiv.2501.18689

Väisälä, V. 1925, *Über die Wirkung der Windschwankungen auf die Pilotbeobachtungen*, Societus Scientiarum Fennica Commentationes Physico-Mathematicae, 2, 19

van den Heuvel, E. P. J. 1969, *The Expected Fraction of Evolved Close Binaries among Main-Sequence Stars of Spectral Type Earlier than A5*, AJ, 74, 1095, doi: 10.1086/110909

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

—. 2011, Compact stars and the evolution of binary systems (World Scientific Publishing Co. Pte. Ltd), 55–73, doi: 10.1142/9789814374774_0006

van den Heuvel, E. P. J., Portegies Zwart, S. F., & de Mink, S. E. 2017, *Forming short-period Wolf-Rayet X-ray binaries and double black holes through stable mass transfer*, MNRAS, 471, 4256, doi: 10.1093/mnras/stx1430

van der Linden, T. J. 1987, *The evolution of intermediate mass Case B close binaries.*, A&A, 178, 170

van Son, L. A. C., de Mink, S. E., Callister, T., et al. 2022a, *The Redshift Evolution of the Binary Black Hole Merger Rate: A Weighty Matter*, ApJ, 931, 17, doi: 10.3847/1538-4357/ac64a3

van Son, L. A. C., De Mink, S. E., Broekgaarden, F. S., et al. 2020, *Polluting the Pair-instability Mass Gap for Binary Black Holes through Super-Eddington Accretion in Isolated Binaries*, ApJ, 897, 100

van Son, L. A. C., de Mink, S. E., Renzo, M., et al. 2022b, *No Peaks without Valleys: The Stable Mass Transfer Channel for Gravitational-wave Sources in Light of the Neutron Star-Black Hole Mass Gap*, ApJ, 940, 184, doi: 10.3847/1538-4357/ac9b0a

Vanbeveren, D., De Donder, E., Van Bever, J., Van Rensbergen, W., & De Loore, C. 1998, *The WR and O-type star population predicted by massive star evolutionary synthesis*, New A, 3, 443, doi: 10.1016/S1384-1076(98)00020-7

Vasiliev, E. 2019, *AGAMA: action-based galaxy modelling architecture*, MNRAS, 482, 1525, doi: 10.1093/mnras/sty2672

Veske, D., Bartos, I., M'arka, Z., & M'arka, S. 2021, *Characterizing the Observation Bias in Gravitational-wave Detections and Finding Structured Population Properties*, ApJ, 922, 258, doi: 10.3847/1538-4357/ac27ac

Vigna-Gómez, A., Neijssel, C. J., Stevenson, S., et al. 2018, *On the formation history of Galactic double neutron stars*, MNRAS, 481, 4009, doi: 10.1093/mnras/sty2463

Vigna-Gómez, A., Willcox, R., Tamborra, I., et al. 2024, *Constraints on Neutrino Natal Kicks from Black-Hole Binary VFTS 243*, Phys. Rev. Lett., 132, 191403, doi: 10.1103/PhysRevLett.132.191403

Vink, J. S. 2017, *Winds from stripped low-mass helium stars and Wolf-Rayet stars*, A&A, 607, L8, doi: 10.1051/0004-6361/201731902

Vink, J. S., & de Koter, A. 2005, *On the metallicity dependence of Wolf-Rayet winds*, A&A, 442, 587, doi: 10.1051/0004-6361:20052862

Vink, J. S., de Koter, A., & Lamers, H. J. G. L. M. 2000, *New theoretical mass-loss rates of O and B stars*, A&A, 362, 295. https://arxiv.org/abs/astro-ph/0008183

—. 2001, *Mass-loss predictions for O and B stars as a function of metallicity*, A&A, 369, 574, doi: 10.1051/0004-6361:20010127

Vynatheya, P., & Hamers, A. S. 2022, *How Important Is Secular Evolution for Black Hole and Neutron Star Mergers in 2+2 and 3+1 Quadruple-star Systems?*, ApJ, 926, 195, doi: 10.3847/1538-4357/ac4892

Waelkens, C. 1991, *Slowly pulsating B stars.*, A&A, 246, 453

Waelkens, C., & Rufener, F. 1985, *Photometric variability of mid-B stars.*, A&A, 152, 6

Wagg, T., Breivik, K., & de Mink, S. 2022a, *LEGWORK: A python package for computing the evolution and detectability of stellar-origin gravitational-wave sources with space-based detectors*, The Journal of Open Source Software, 7, 3998, doi: 10.21105/joss.03998

Wagg, T., Breivik, K., & de Mink, S. E. 2022b, *LEGWORK: A Python Package for Computing the Evolution and Detectability of Stellar-origin Gravitational-wave Sources with Space-based Detectors*, ApJS, 260, 52, doi: 10.3847/1538-4365/ac5c52

Wagg, T., Broekgaarden, F., de Mink, S., et al. 2021, *Dataset from: Gravitational wave sources in our Galactic backyard - Predictions for BHBH, BHNS and NSNS binaries in LISA*, 0.0.1, Zenodo, doi: 10.5281/zenodo.5725607

Wagg, T., Broekgaarden, F. S., de Mink, S. E., et al. 2022c, *Gravitational Wave Sources in Our Galactic Backyard: Predictions for BHBH, BHNS, and NSNS Binaries Detectable with LISA*, ApJ, 937, 118, doi: 10.3847/1538-4357/ac8675

Wagg, T., Hendriks, D. D., Renzo, M., & Breivik, K. 2025, *Stellar ejection velocities from the binary supernova scenario: A comparison across population synthesis codes*, arXiv e-prints, arXiv:2504.16161. https://arxiv.org/abs/2504.16161

Wagg, T., Johnston, C., Bellinger, E. P., et al. 2024, *The asteroseismic imprints of mass transfer. A case study of a binary mass-gainer in the SPB instability strip*, A&A, 687, A222, doi: 10.1051/0004-6361/202449912

Walch, S., Girichidis, P., Naab, T., et al. 2015, *The SILCC (SImulating the LifeCycle of molecular Clouds) project - I. Chemical evolution of the supernova-driven ISM*, MNRAS, 454, 238, doi: `10.1093/mnras/stv1975`

Wang, B. 2018, *Mass-accreting white dwarfs and type Ia supernovae*, Research in Astronomy and Astrophysics, 18, 049, doi: `10.1088/1674-4527/18/5/49`

Weaver, T. A., Zimmerman, G. B., & Woosley, S. E. 1978, *Presupernova evolution of massive stars.*, ApJ, 225, 1021, doi: `10.1086/156569`

Webbink, R. F. 1984, *Double white dwarfs as progenitors of R Coronae Borealis stars and type I supernovae.*, ApJ, 277, 355, doi: `10.1086/161701`

Webster, B. L., & Murdin, P. 1972, *Cygnus X-1-a Spectroscopic Binary with a Heavy Companion ?*, Nature, 235, 37, doi: `10.1038/235037a0`

Wegg, C., Gerhard, O., & Portail, M. 2015, *The structure of the Milky Way's bar outside the bulge*, MNRAS, 450, 4050, doi: `10.1093/mnras/stv745`

Weinberg, M. D., Shapiro, S. L., & Wasserman, I. 1987, *The Dynamical Fate of Wide Binaries in the Solar Neighborhood*, ApJ, 312, 367, doi: `10.1086/164883`

Weisz, D. R., Fouesneau, M., Hogg, D. W., et al. 2013, *The Panchromatic Hubble Andromeda Treasury. IV. A Probabilistic Approach to Inferring the High-mass Stellar Initial Mass Function and Other Power-law Functions*, ApJ, 762, 123, doi: `10.1088/0004-637X/762/2/123`

Wetzel, A., Hayward, C. C., Sanderson, R. E., et al. 2023, *Public Data Release of the FIRE-2 Cosmological Zoom-in Simulations of Galaxy Formation*, ApJS, 265, 44, doi: `10.3847/1538-4365/acb99a`

Wetzel, A. R., Hopkins, P. F., Kim, J.-h., et al. 2016, *Reconciling Dwarf Galaxies with ΛCDM Cosmology: Simulating a Realistic Population of Satellites around a Milky Way-mass Galaxy*, ApJ, 827, L23, doi: `10.3847/2041-8205/827/2/L23`

Wex, N., Kalogera, V., & Kramer, M. 2000, *Constraints on Supernova Kicks from the Double Neutron Star System PSR B1913+16*, ApJ, 528, 401, doi: `10.1086/308148`

White, R. L., Long, K. S., Becker, R. H., et al. 2019, *A New, Deep JVLA Radio Survey of M33*, ApJS, 241, 37, doi: `10.3847/1538-4365/ab0e89`

Whyte, C. A., & Eggleton, P. P. 1985, *A simple model for binary star evolution*, MNRAS, 214, 357, doi: `10.1093/mnras/214.3.357`

Wiggins, B. K., Fryer, C. L., Smidt, J. M., et al. 2018, *The Location and Environments of Neutron Star Mergers in an Evolving Universe*, ApJ, 865, 27, doi: `10.3847/1538-4357/aad2d4`

Wijers, R. A. M. J., van Paradijs, J., & van den Heuvel, E. P. J. 1992, *Binary pulsars as probes of neutron star birth.*, A&A, 261, 145

Willems, B., Kalogera, V., Vecchio, A., et al. 2007, *Eccentric Double White Dwarfs as LISA Sources in Globular Clusters*, ApJ, 665, L59, doi: `10.1086/521049`

Wong, T. L. S., White, C. J., & Bildsten, L. 2024, *Shocking and Mass Loss of Compact Donor Stars in Type Ia Supernovae*, ApJ, 973, 65, doi: `10.3847/1538-4357/ad6a11`

Wong, T.-W., Valsecchi, F., Fragos, T., & Kalogera, V. 2012, *Understanding Compact Object Formation and Natal Kicks. III. The Case of Cygnus X-1*, ApJ, 747, 111, doi: `10.1088/0004-637X/747/2/111`

Wongwathanarat, A., Janka, H. T., & Müller, E. 2013, *Three-dimensional neutrino-driven supernovae: Neutron star kicks, spins, and asymmetric ejection of nucleosynthesis products*, A&A, 552, A126, doi: `10.1051/0004-6361/201220636`

Woosley, S. E., Blinnikov, S., & Heger, A. 2007, *Pulsational pair instability as an explanation for the most luminous supernovae*, Nature, 450, 390, doi: `10.1038/nature06333`

Xiao, L., & Eldridge, J. J. 2015, *Core-collapse supernova rate synthesis within 11 Mpc*, MNRAS, 452, 2597, doi: `10.1093/mnras/stv1425`

Xin, C., Renzo, M., & Metzger, B. D. 2022, *Dissecting the microphysics behind the metallicity-dependence of massive stars radii*, MNRAS, 516, 5816, doi: `10.1093/mnras/stac2551`

Xu, X.-J., & Li, X.-D. 2010a, *On the Binding Energy Parameter λ of Common Envelope Evolution*, ApJ, 716, 114, doi: `10.1088/0004-637X/716/1/114`

—. 2010b, *ERRATUM: "On the Binding Energy Parameter λ of Common Envelope Evolution"*, ApJ, 722, 1985, doi: `10.1088/0004-637X/722/2/1985`

Yang, M., Bonanos, A. Z., Jiang, B., et al. 2023, *Evolved massive stars at low-metallicity. V. Mass-loss rate of red supergiant stars in the Small Magellanic Cloud*, A&A, 676, A84, doi: `10.1051/0004-6361/202244770`

Ye, C., & Fishbach, M. 2022, *Inferring the Neutron Star Maximum Mass and Lower Mass Gap in Neutron Star-Black Hole Systems with Spin*, ApJ, 937, 73, doi: `10.3847/1538-4357/ac7f99`

Yi, S.-X., Nelemans, G., Brinkerink, C., et al. 2022, *The Gravitational Wave Universe Toolbox. A software package to simulate observations of the gravitational wave universe with different detectors<xref ref-type="fn" rid="FN2">,<xref ref-type="fn" rid="FN3">**, A&A, 663, A155, doi: `10.1051/0004-6361/202141634`

Yoon, S.-C., Dessart, L., & Clocchiatti, A. 2017, *Type Ib and IIb Supernova Progenitors in Interacting Binary Systems*, ApJ, 840, 10, doi: `10.3847/1538-4357/aa6afe`

Yoon, S. C., Woosley, S. E., & Langer, N. 2010, *Type Ib/c Supernovae in Binary Systems. I. Evolution and Properties of the Progenitor Stars*, ApJ, 725, 940, doi: `10.1088/0004-637X/725/1/940`

Yu, S., & Jeffery, C. S. 2010, *The gravitational wave signal from diverse populations of double white dwarf binaries in the Galaxy*, A&A, 521, A85, doi: `10.1051/0004-6361/201014827`

Yungelson, L. 1973, *Evolution of the secondary component of a close binary system*, Nauchnye Informatsii, 27, 93

Zapartas, E., de Mink, S. E., Izzard, R. G., et al. 2017, *Delay-time distribution of core-collapse supernovae with late events resulting from binary interaction*, A&A, 601, A29, doi: `10.1051/0004-6361/201629685`

Zehe, T., Mugrauer, M., Neuhäuser, R., et al. 2018, *The radial and rotational velocity of ζ Ophiuchi*, Astronomische Nachrichten, 339, 46, doi: `10.1002/asna.201713383`

Ziosi, B. M., Mapelli, M., Branchesi, M., & Tormen, G. 2014, *Dynamics of stellar black holes in young star clusters with different metallicities - II. Black hole-black hole binaries*, MNRAS, 441, 3703, doi: `10.1093/mnras/stu824`

Zolotov, A., Dekel, A., Mandelker, N., et al. 2015, *Compaction and quenching of high-z galaxies in cosmological simulations: blue and red nuggets*, MNRAS, 450, 2327, doi: `10.1093/mnras/stv740`

Zonca, A., Singer, L., Lenz, D., et al. 2019, *healpy: equal area pixelization and spherical harmonics transforms for data on the sphere in Python*, Journal of Open Source Software, 4, 1298, doi: `10.21105/joss.01298`