What limits the luminosity of red supergiants?

Eric Svensson

Division of Astrophysics
Department of Physics
Lund University

2023-EXA201

Degree project of 15 higher education credits
June 2023

Supervisor: Ross Church

Division of Astrophysics
Department of Physics
Box 43
SE-221 00 Lund
Sweden

Abstract

The observed upper limit of the luminosity of red supergiants, often called the Humphreys-Davidson limit, has recently been shown to be independent of metallicity and to not be explainable by single star evolution – findings that both contradict established ideas of the origins of the limit. It is now known that the vast majority of stars with sufficient mass to possibly evolve into red supergiants above the limit interact with companion stars at some point during their lives. In this thesis, the possibility of binary star interactions causing the Humphreys-Davidson limit is investigated. This is done numerically with the use of synthesized populations at M31, LMC, and SMC metallicities, after confirming that the simulations work as intended by producing a number of single star results. It is found that while binary star interactions can not wholly reproduce the limit, such interactions are sufficient to reproduce the limit for any red supergiant that remains in a binary and does not actively transfer mass to its companion in the red supergiant phase. For such systems upper luminosity limits of $\log_{10}(L/L_{\odot}) = 5.43^{+0.06}_{-0.09}$ for M31, $\log_{10}(L/L_{\odot}) = 5.48^{+0.07}_{-0.09}$ for the LMC, and $\log_{10}(L/L_{\odot}) = 5.53^{+0.07}_{-0.11}$ for the SMC are found, which are consistent with recent estimates of the Humphreys-Davidson limit. However, stars that are the result of a binary merging are found to be able to surpass the limit. This may be due to the physics of mergers not being modelled appropriately. If such stars genuinely do exist, interactions with a tertiary star may be what stops them from surpassing the limit.

Populärvetenskaplig beskrivning

Där uppe finns de. Stjärnorna som om natten pryder himlavalvet är fängslande i deras mängd och mångfald. Om du befinner dig långt bort från städernas ljus så kan du se nattens sanna skönhet, och om du tittar riktigt noga kan du börja uppskatta att alla stjärnor inte är likformiga vita prickar. Vissa är ljusstarkare, vissa är ljussvagare. Vissa är rödare än andra och vissa ser blåare ut. För astronomer är det dessa skillnader i ljusstyrka och färg vad som separerar olika typer av stjärnor från varandra.

En särskild typ av stjärna är de röda superjättarna. Röda superjättar har bränt allt väte i sina kärnor till tyngre ämnen och sedan expanderat till jättelika versioner av sig själva. Som namnet antyder har de en rödaktig färg och kan nå gigantiska storlekar. Deras ljusstyrka beror på kärnans massa, som teoretiskt sett skulle kunna vara hur tung som helst. Om kärnan kan vara hur tung som helst, då borde röda superjättar kunna bli hur ljusa som helst. Men icke. Astronomer blev mycket förvånade på slutet av 70-talet när det visade sig att det faktiskt finns en gräns på hur ljusa röda superjättar kan bli. Det pågår något hos de röda superjättarna som astronomer ännu inte fullt ut förstår.

Länge trodde man att riktigt tunga stjärnor var för instabila för att kunna behålla sin stora radie och att de istället utvecklades till blåare, mindre stjärnor, och därmed skapade en övre gräns för röda superjättar. Modern forskning har motbevisat denna teori och frågan är återigen öppen. I mitt kandidatarbete använder jag mig av den nyligen upptäckta faktan att de allra flesta, kanske alla, röda superjättar befinner sig i binära stjärnsystem – stjärnsystem med två stjärnor. Jag undersöker om interaktioner med partnerstjärnan kan ligga bakom gränsen i ljusstyrka för de röda superjättarna och med det bidra till att äntligen sätta punkt på denna bestående astronomiska gåta.

Contents

Chapter 1

Introduction

Red giants are – as the name implies – red and very large stars. They make up a post-main sequence stage of stellar evolution, i.e., a stage of stars that have exhausted their core hydrogen. Red giants have dead cores: no nuclear fusion occurs there. But surrounding the core is a thin layer, called a shell, of hydrogen burning. Outside the shell there is then a large convective hydrogen envelope (Eldridge & Tout 2019 p. 210–213). The envelope makes up most of the giant's extensive radius and since the effective surface temperature of the envelope is relatively low the star appears red. If the progenitor star and thus the red giant is massive enough, it is known as a red supergiant (RSG).

The luminosity of a star burning hydrogen in a shell is proportional to the mass of its core (Kippenhahn et al. 2012 p. 426) which is inherited from the progenitor star. Since there is no such thing as a hard upper limit on how massive a star can be, there is a priori no upper mass limit of the cores of stars either. If the progenitor core mass can be arbitrarily large the resulting RSG should thus be able to be arbitrarily luminous. But this is not what has been observed. Humphreys & Davidson (1979) found that in the solar neighbourhood and in the LMC there is in actuality an upper limit of RSG luminosity. Subsequent studies have confirmed the limit – now denoted the Humphreys-Davidson limit or H-D limit – as empirical fact in multiple environments. Table 1.1 contains H-D limits in the galaxies M31, the LMC, and the SMC.

The H-D limit was originally thought by Humphreys & Davidson (1979) to be due to instabilities in extremely massive stars, causing the hydrogen envelope to be lost. Casting off the hydrogen envelope would expose the shell surrounding the core, which hosts temperatures orders of magnitudes higher than on the surface of an RSG. Such high temperatures implies a blackbody radiation peak in the bluer wavelengths, not the red. Giants that lose their hydrogen envelopes would thus evolve to have a significantly smaller radius and a bluer spectrum, possibly evolving into Wolf-Rayet stars (WRs). WRs are luminous stars with high effective temperature and relatively low hydrogen abundance on the surface (Crowther 2007). Humphreys & Davidson (1979) proposed that above some initial mass threshold the loss of the envelope occurs before the star ever reaches the RSG phase,

Table 1.1: Humphreys-Davidson limits in different environments with different chemical compositions. $Z_{\rm obs}$ is an average value of observed metallicites in the galaxies, while $Z_{\rm BSE}$ is the metallicity used in the simulations for the galaxies. See Section 2.1 for why these values differ. $L_{\rm max}$ refers to the upper luminosity limits.

Galaxy	$Z_{\rm obs}$	$Z_{\rm BSE}$	$\log_{10}(L_{\rm max}/L_\odot)$	Reference (for luminosities)
M31	$\sim Z_\odot$	0.02	5.53 ± 0.03	McDonald et al. (2022)
LMC	$\sim 0.5Z_\odot$	0.01	5.5	Davies et al. (2018)
SMC	$\sim 0.25Z_\odot$	0.004	5.5	Davies et al. (2018)

causing an effective disappearance in the observational record of all would-be RSGs above this threshold. Since RSG luminosity is dependent on core mass this would also create an effective upper limit of RSG luminosity – the H-D limit.

The mechanics of this mass loss was suggested by Humphreys & Davidson (1979) to be mainly due to stellar winds and periods of pulsational instability. They base their arguments on the properties of two well-studied massive stars: Eta Carina (η Car) and P Cygni (P Cyg). η Car and P Cyg are both counted among a class of stars called luminous blue variables (LBVs) characterized by their great mass and instability. It is now known that η Car's mass loss occurred in two main eruptions (Humphreys et al. 1999). The greatest of these occurred in the 1840s and the second around 1890. Eruptions like those of η Car – LBV eruptions – could explain the H-D limit if they are sufficiently dominant before the red giant phase for stars with masses $M > 50M_\odot$. In that case, such stars would likely evolve into Wolf-Rayet (WR) stars directly following the LBV phase and never become RSGs (Eldridge & Tout 2019 p. 255-256).

Gull et al. (2022) recently discussed η Car with respect to modern observations and understanding. It is now known that η Car is part of a binary star system with a smaller, though still massive, partner star η Car-B. The effect of this upon the evolution of the primary is however unclear. Additionally, the flux of far ultraviolet light from η Car has increased by an order of magnitude since the turn of the millennium though the stellar winds of neither star in the binary seem to have changed.

Beasor et al. (2020) derived a new mass-loss rate prescription for RSGs using newer observational data. Clusters with a high number of RSGs, mainly the Milky Way cluster RSGC1 and the LMC cluster NGC 2004, were investigated in the study since the age of the stars in a cluster can generally be known. The authors argue that all stars within a cluster currently in a RSG phase must have very similar initial masses since the RSG phase is very short and they all have roughly the same age. Using that the age and initial masses of the RSGs could be known the authors could then derive a mass-loss rate prescription based on initial mass. They found that this new mass-loss rate is much too low to alone explain the H-D limit. The results indicated no reason to expect that would-be

RSGs above the limit should lose their envelopes through internal processes alone, unless LBV-like eruptions before the RSG phase are prevalent enough to discard the envelope on their own.

This is supported by recent findings by McDonald et al. (2022). They studied RSGs in M31 using mid-infrared photometry and determined an upper luminosity limit and a luminosity function for these RSGs. The paper also includes upper luminosity limits for RSGs in the LMC and SMC, found by Davies et al. (2018). The authors concluded that both the luminosity function of RSGs and the upper luminosity limit are independent of metallicity. Furthermore, they argued that this suggests that stellar winds are not the cause of the H-D limit – since the strength of such winds should be proportional to metallicity. They also find that the limit is lower than predicted by Humphreys & Davidson (1979): at $\log_{10}(L/L_{\odot}) \approx 5.5$, regardless of metallicity, rather than $\log_{10}(L/L_{\odot}) \approx 5.8$.

If internal processes such as stellar winds are insufficient to be the underlying mechanism behind the H-D limit, some other explanation is needed. Binary star interactions may be such an alternative explanation. Sana et al. (2012) found that 71% of O-type stars – stars with masses $\gtrsim 16M_{\odot}$ (Moe & Di Stefano 2017) – exchange mass with a companion star during their lifetime. Let this be the definition of a star being in a close binary system. Humphreys & Davidson (1979) estimated that RSGs at the H-D limit have masses of up to $60M_{\odot}$, and unpublished work by Maxwell Moe suggests that the fraction of stars in close binary systems approaches unity as stellar mass increases. It can therefore be concluded that it is highly probable that any RSG in the mass range of the H-D limit will be interacting with a companion. Mass transfer onto the companion star could strip a highly massive star of its hydrogen envelope and theoretically cause it to evolve bluewards in much the same way as the previously proposed processes might. If close binary interactions can reproduce the H-D limit and a very high fraction of massive RSGs are in close binaries, then binary interactions may explain the H-D limit. This is the focus of this thesis.

Massive binary star systems are in and of themselves interesting to the scientific context at large, and the physics of binary star systems are becoming increasingly relevant. Since it is now known that most massive stars interact with a companion at some point (Sana et al. 2012), binary star interactions can be considered a prevalent phase in the evolution of massive stars in general. These massive stars are important supernova progenitors – the effect of binary star interactions have long been studied in exactly this context, see e.g. Podsiadlowski et al. (1992). But they are also the progenitor systems to an important source of gravitational waves: black hole binaries, as discovered by Abbott et al. (2016). The effect of binarity upon the evolutionary paths of massive stars is studied in this thesis primarily from the perspective of the luminosities of RSGs, but the results concern the physics of all massive binaries.

McDonald et al. (2022) performed some investigations into binary interactions as part of their discussion in light of the results of Sana et al. (2012) and similar studies. For

this they used the Binary Population and Spectral Analysis code (`BPASS`), which is a set of models of stellar evolution that is capable of modelling the effects of binary evolution. Using `BPASS` McDonald et al. (2022) found that the luminosity function of the binary evolutionary path behaves very similarly to that of the single star evolutionary path. This is noted as being surprising by the authors, but they otherwise leave investigations into this as a question for future work.

1.1 Close binary systems

1.1.1 Roche lobes

The gravitational potential Φ of a binary system is the sum of the individual potentials resulting from the mass of the stars and a rotational potential. Between the stars lies the inner Lagrangian point, L_1, where $\nabla\Phi = 0$. The two equipotential surfaces centered on each star that intersect L_1 are known as the Roche limits for the stars, and the volumes they bound are known as the Roche lobes (Hilditch 2001 p. 157-158). An expanding star will expand such that its surface is continually an equipotential. When the surface reaches the Roche limit – referred to as the star having filled its Roche lobe – any further expansion will cause the material at L_1 to move into the potential well of the partner star. This is known as Roche lobe overflow and the result is mass from one star accreting onto the other (Eldridge & Tout 2019 p. 297). A radius R_L can be defined as the maximum radius the star can reach before it fills its Roche lobe (Eldridge & Tout 2019 p. 293), if stars are approximated to always be spherical. If the radius of the star is larger than R_L, then the system is undergoing mass transfer.

1.1.2 Mass transfer

Mass transfer due to Roche lobe overflow occurs on dynamical timescales, much shorter than the nuclear timescales of stellar evolution on the main sequence (MS) (Eldridge & Tout 2019 p. 301). Such mass transfer can have significant effects on the evolution of both stars. The nature of the mass transfer is highly dependent on the size of the Roche lobes, which in turn is dependent on the size of the binary stars' orbits. Three cases are typically considered: Cases A, B, and C. Case A is characterized by a short initial orbital period – and thus a small orbital semi-major axis – that allows the primary star to fill its Roche lobe while still on the MS. Such mass transfers are believed to be stable and conserve mass in the system. Case B covers longer initial orbital periods wherein Roche lobe overflow occurs when the primary is transitioning from the MS to the red giant branch. Finally, Case C describes initial orbital periods long enough that stars will enter the red giant stage without having filled their Roche lobes (Podsiadlowski et al. 1992 and Hilditch 2001 p. 25).

Mass transfer after a star has stopped core hydrogen burning, Cases B and C, in systems where the primary and secondary originally have similar masses, is characterized by

the primary losing most of its envelope mass. However, some envelope mass is retained and the primary remains visually similar to what it would have looked like if it kept its envelope. Most of the lost envelope mass is accreted onto the companion star which as a result grows very luminous. The evolutionary effect the accretion has on the companion depends on in which stage the companion is at the time of accretion. If the companion is on the MS it will start to follow the evolutionary path of a higher-mass MS star than it was originally. If the companion is past the MS it may instead evolve into a blue supergiant (BSG), if enough mass is accreted (Podsiadlowski et al. 1992).

In the case that the envelope of the primary expands to enclose both stars and fills the Roche lobes of both stars, then the system enters a state referred to as common envelope (CE) evolution. The core of the primary star and the entirety of the secondary star orbit each other within the large and low density envelope of the primary. Two evolutionary paths are then available for the system: complete loss of the envelope or merging. Loss of the envelope will lead to a close binary system with a white dwarf – the remnant of the core of the primary – and the seconday star. If the orbit of this binary is shrunk further through magnetic breaking or radiation of gravitational waves, the Roche lobe of the secondary may shrink to the point that it overfills and the white dwarf primary starts accreting material from the secondary. The system then becomes a so-called cataclysmic variable. Merging, on the other hand, turns the binary system into a single star system. In such systems the core of the primary is enveloped by the secondary which in turn is still within the primary's envelope (Eldridge & Tout 2019 p. 315-316).

Chapter 2

Methods

2.1 BSE

All stellar evolution simulations in this thesis are carried out using the very fast binary evolution code BSE, written and maintained by Prof. Jarrod R. Hurley at the Swinburne University of Technology. BSE utilizes analytical formulae fitted to approximate stellar evolution in order to evolve a star from initial conditions through time. Formulae governing single star evolution and the corresponding evolutionary tracks were presented in Hurley et al. (2000), while a follow-up paper, Hurley et al. (2002), presented the binary star evolution algorithms.

The code will largely be used as-is with two exceptions. Firstly, the mass-loss rate ($\dot{M}$) prescription for giant stars. BSE includes an LBV-like $\dot{M}$ prescription for would-be supergiants above the H-D limit:

$$\dot{M}_{\rm LBV} = 0.1 \left(10^{-5} \frac{R}{R_\odot} \sqrt{\frac{L}{L_\odot}} - 1.0 \right)^3 \left(\frac{L}{6 \cdot 10^5 L_\odot} - 1.0 \right) M_\odot \ {\rm yr}^{-1} \tag{2.1}$$

which is added as a term onto the mass loss of all stars with $L/L_\odot > 6 \cdot 10^5$ and $10^{-5}(R/R_\odot)\sqrt{L/L_\odot} > 1.0$.

This has the desired effect of forcing massive single stars to evolve bluewards and reproduce the H-D limit. However, given the recent work on RSG mass-loss rate prescriptions, it is not deemed an entirely desirable model for the purposes of this thesis. This default BSE prescription is still used though in Section 3.1 for comparison to the results generated with the Beasor et al. (2020) $\dot{M}$ prescription, which is as follows:

$$\log_{10} \left(\frac{\dot{M}}{M_\odot \ {\rm yr}^{-1}} \right) = \left(-26.4 - 0.23 \frac{M_{\rm ini}}{M_\odot} \right) + 4.8 \log_{10} \left(\frac{L_{\rm bol}}{L_\odot} \right) \tag{2.2}$$

where $M_{\rm ini}$ is the initial mass of the star, and $L_{\rm bol}$ is the bolometric luminosity of the star.

The second exception to the use of BSE as it is distributed is a series of modifications regarding maximum stellar mass. BSE restricts the mass of stars to $100 M_\odot$ for both mass transfer and merger events. For some results regarding the most massive stars generated in this thesis this causes some problems, so this restriction is removed.

BSE is written in `FORTRAN`. `Python` is used in this thesis to write wrapper programs that provide input values to and extract output values from BSE. The role of the wrapper is to generate a synthesized population of stars with realistic parameters. These parameters can then be input into BSE to evolve every star from birth to the present day, at which point the set of outputs should faithfully resemble a real stellar population. The following subsections will explain how, and under what assumptions, these input parameters are generated.

2.1.1 Single star input parameters

For single star simulations the relevant input parameters are stellar mass M, stellar age t, metallicity Z – i.e., the fraction of the mass of all elements heavier than helium within a star and the total mass of that star – of the environment, and population size N.

For the purpose of comparison, N is normalized so that the output number of RSGs match Figure 4 in McDonald et al. (2022) which contains data from 415 RSGs in M31. However, many of these RSGs have relatively low luminosities, which overlap with luminosities of asymptotic giant branch (AGB) stars. Since differentiating AGBs and RSGs in the BSE output is not entirely straightforward, generating any AGBs at all is avoided by enforcing a minimum mass. The repercussion of this is that the number of generated RSGs towards the lower end of the luminosity spectrum will also be decreased and thus not match what is seen in McDonald et al. (2022). This is solved by not fitting the number of generated RSGs to the entire luminosity range covered by McDonald et al. (2022), but only to those with $\log_{10}(L/L_\odot) > 5.0$. This cut-off was chosen since investigations with the code revealed that no AGBs were generated above it. The population size N is found by guessing an input N_{guess} and then renormalizing using the output number of RSGs above the cut-off N_{RSGs} and the observed number of RSGs above the cut-off N_{obs}. I.e., $N = \frac{N_{\text{obs}}}{N_{\text{RSGs}}} N_{\text{guess}}$. In McDonald et al. (2022), $N_{\text{obs}} = 65$. In the single star case, $N = 6500$ is found.

M is generated by using a cumulative distribution function (CDF), from some minimum mass M_{min}, based on the Salpeter initial mass function (IMF), up to some maximum mass M_{max}, such that

$$M = \left[\left(M_{\text{max}}^{-1.35} - M_{\text{min}}^{-1.35} \right) C(M) + M_{\text{min}}^{-1.35} \right]^{-1/1.35} \tag{2.3}$$

where the CDF $C(M)$ is generated as some random value in the range $[0, 1]$. The minimum mass is chosen as to make it impossible to generate AGB stars, for the reasons explained above. By experimentation with the code it was found that choosing $M_{\text{min}} = 10 M_\odot$ comfortably excludes all AGBs. M_{max} is much more straightfoward by comparison. So long

as $M_{\max}$ is large enough to encompass the mass of RSGs at the H-D limit – $M \sim 60 M_{\odot}$ (Humphreys & Davidson 1979) – then the specific choice is quite arbitrary. Additionally, the Salpeter IMF deselects against stars with masses near $M_{\max}$. So even if stars with very high masses had a large effect there would be very few of them in the final synthesized population. By stochastically sampling the steep upper end of the IMF, $M_{\max} = 100$ was found to be a fitting upper limit.

When generating the age of the stars, t, a uniform star formation rate (SFR) is assumed. This is quantified as a flat probability distribution for t between zero and a maximum age $t_{\max}$. The maximum age is obtained by evolving a star with $M = M_{\min}$ in BSE. The evolution time of said star, i.e., the time it takes for that star to evolve into some type of compact object – white dwarf, neutron star, or black hole – is used as $t_{\max}$. Assuming a uniform SFR is motivated by the short timescales involved, as the evolution time for any star that may become an RSG is on the order of 10^7 years or less. This is brief on galactic timescales so the SFR is not expected to change much, if at all, during that time.

Finally the metallicity Z, which is kept constant for all stars in the population unlike M and t. The relevant metallicities to this thesis are those of M31, the LMC, and the SMC, see Table 1.1, which are approximated to be constant for each galaxy. However, to ensure accuracy of the results, the BSE simulations are not ran with exactly these metallicities. Hurley et al. (2000) fitted their formulae to evolutionary tracks with seven discrete values for Z. The closest available 'on-track' Z to the observed M31, LMC, and SMC metallicities are thus used, see Table 1.1 for the values chosen.

2.1.2 Binary input parameters

In the population synthesis of binary systems there are a few more parameters to consider than in the single star case. Masses and ages for both stars need to be considered, as well as orbital semi-major axis and eccentricity. Metallicity needs to be handled too, but this is done in exactly the same way as in the single star case and for the same reasons. Regarding the population size N, this is also largely done in the same way as for the single star case. When normalizing the sample to N_{obs}, N_{RSG} is counted as the number of systems where the primary is a RSG. A double RSG system would not count as two RSGs in this context. Systems where the secondary is a RSG but the primary is not are seen incredibly rarely in the results, so for the sake of simplicity these are not counted either.

The mass of the primary star, denoted M_1, is generated in the same way as the stellar mass in the single star case. The mass of the partner star, M_2, is then calculated from the fraction $q = M_2/M_1$. Observations have found the probability distribution of q to be flat for close binaries with O-type primaries, Moe & Di Stefano (2017), so q is generated as a random value between 0 and 1.0.

The age of both stars in the binary is assumed to be equal, and is generated following

the same approach as the single star age. The maximum age, $t_{\max}$, is however set to twice the evolution time of an $M_{1,\min}$ star. This is done as to not exclude interesting behavior of the partner star after the primary has fully evolved.

In determining the semi-major axis a of the binary orbit, a probability distribution based on very recent unpublished work by Maxwell Moe is used. The work suggests that at least 80% of all stars with masses greater than $25M_\odot$ are in close binaries with semi-major axes smaller than 10 a.u. The probability distribution of $\log_{10}(a)$ is taken to be flat between two points $\log_{10}(a_{\min})$ and $\log_{10}(a_{\max})$ which are individually determined for each binary system. $a_{\min}$ is determined as $a_{\min} = 3R_1$ where R_1 is the radius of the main star at zero age, while $a_{\max}$ is determined by the equation

$$a_{\max} = a_{\min}\left(\frac{10 \text{ a.u.}}{a_{\min}}\right)^{1/\left(0.8\sqrt{\frac{M_1}{25M_\odot}}\right)} \tag{2.4}$$

which is derived from a model fit based on the observed binary fraction of massive stars in the unpublished work by Moe, as well as the flat probability distribution of $\log_{10}(a)$. Using said probability distribution a value for $\log_{10}(a)$ can then be generated for each binary. The radius R_1 is found by a power law relationship

$$R_1 = 0.8276 \cdot M_1^{0.6579} \tag{2.5}$$

between mass and radius for main sequence stars at zero age. The power law is found by fitting to outputs from BSE at initial time with differing initial masses.

Eccentricity e does not make much difference beyond circularization and Hurley et al. (2002) states that an eccentricity distribution is unnecessary for a population of interacting binary systems. Moe's unpublished work is expressed in terms of the orbital radius at the closest point, r_{peri}, so for Equation 2.4 to be accurate we set $r_{\text{peri}} = a$ and thus $e = 0$.

Chapter 3

Results

3.1 Single stars

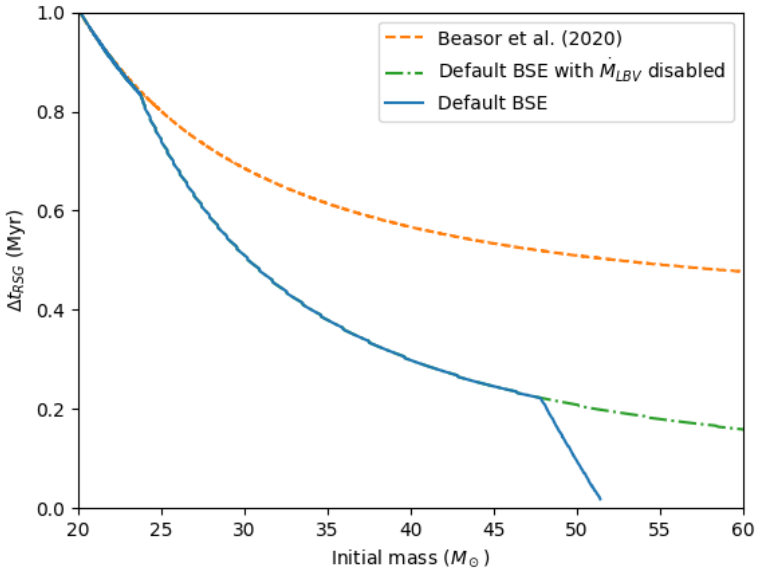

Figure 3.1: Comparison between different supergiant mass-loss rate prescriptions of the time a given star spends as an RSG as a function of initial mass.

Fully evolved stars were investigated to quantify the impact of different mass-loss rates upon the RSGs. In Figure 3.1 the amount of time stars spent as RSGs before fully evolving is plotted against initial mass. This is done for the default BSE mass-loss rate, both with its LBV-like mass loss enabled and disabled, and the Beasor et al. (2020) mass-loss rate prescription. For relatively low-mass stars the mass-loss rates display nearly identical behavior, while for high-mass stars the behaviors diverge significantly. The effect of the $\dot{M}_{\mathrm{LBV}}$ term is made clear by the "knee" in the default BSE rate at $\sim 50 M_{\odot}$ that is not

present when $\dot{M}_{\mathrm{LBV}}$ is disabled. However, the default BSE rate diverges from the Beasor et al. (2020) rate at another knee at $\sim 25 M_{\odot}$ regardless of if $\dot{M}_{\mathrm{LBV}}$ is enabled. In Section 7.1 of Hurley et al. (2000) the mass-loss rates used by BSE are outlined and it is stated that for the giant branch and above a rate from Kudritzki & Reimers (1978) is used, while for the most massive and luminous stars a rate from Nieuwenhuijzen & de Jager (1990) is used. The program decides on which rate to use based on which rate gives a larger $\dot{M}$ for any given star. So, the knee at $\sim 25 M_{\odot}$ is due to the rate for massive, luminous stars becoming dominant over the general giant branch rate.

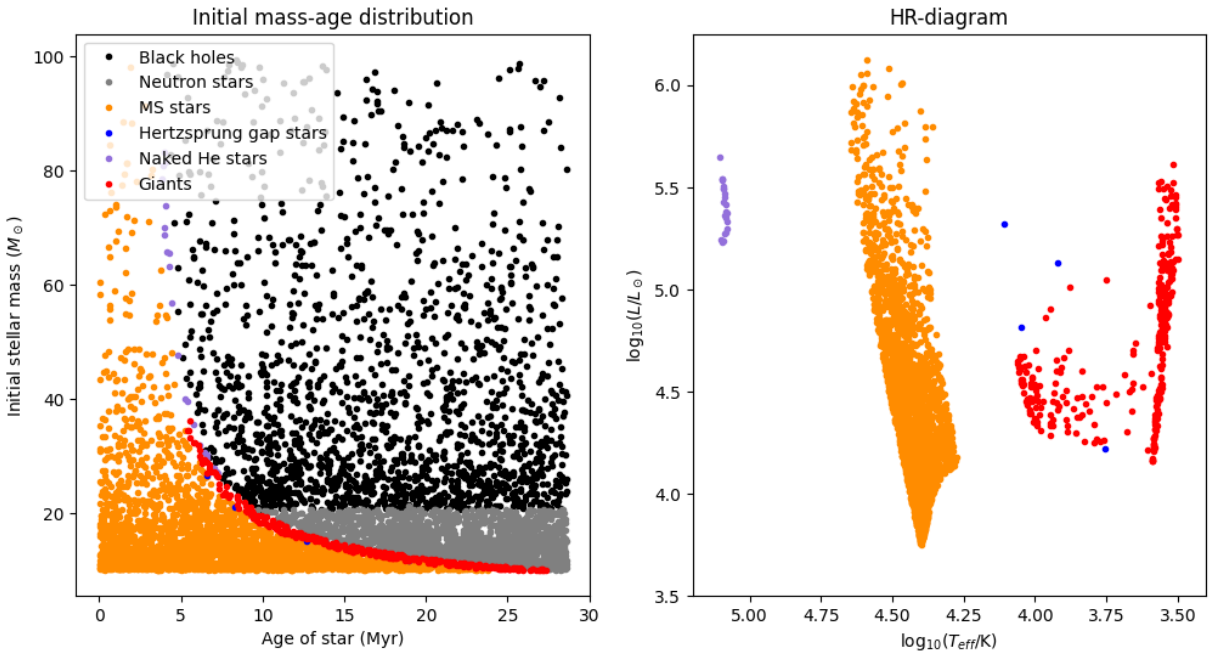

Figure 3.2: Characteristics of a synthesized population of 6500 single star systems with default BSE mass-loss rate prescription. M31 metallicity, $Z = 0.02$. Left: Stellar type as a function of initial mass and age. Right: HR-diagram, with stellar types. The legend applies to both plots, but compact objects have been excluded from the HR-diagram.

Figure 3.2 displays one synthesized population of stars generated according to the standard parameters outlined in Section 2.1.1. BSE evolves any given star i with initial mass M_i and age t_i individually so that its properties at time t_i – current day – can be extracted. The distribution of initial masses seen in the left plot follows the Salpeter IMF as intended. The vast majority of stars are found towards the lower end of the mass distribution while very massive stars are rare. The giants are, for any given initial mass, found in a narrow age range that grows even narrower the higher the initial mass. Note that BSE provides stellar type as an output parameter, and that sometimes this leaves some room for interpretation. "Naked helium" stars are likely BSGs or WRs and the set of stars labelled "Giants" consists of both yellow supergiants (YSGs) and RSGs. Drout et al. (2009) studied YSGs in M31 and defined $\log_{10}(T_{\mathrm{eff}}/\mathrm{K}) = 3.681$ as the dividing line between the two populations: hotter

stars than this were considered YSGs, and cooler stars RSGs. The same sorting criterion is used here. The RSGs are by this definition found in an almost-vertical line at the cooler end of the HR-diagram. At the very top of the RSG branch the H-D limit is reproduced as it was observed by Humphreys & Davidson (1979).

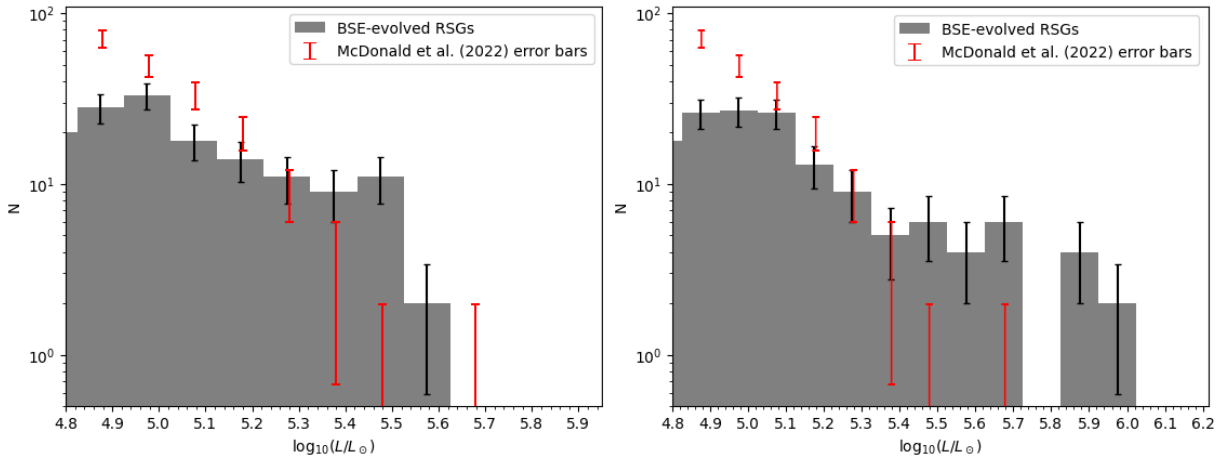

(a) Default BSE $\dot{M}$ prescription.

(b) Beasor et al. (2020) $\dot{M}$ prescription.

Figure 3.3: Histograms of the luminosities of all RSGs in two different synthesized populations. The populations were generated with the same population size, $N = 6500$, but with different $\dot{M}$ prescriptions, see subcaptions. The error bars are found as the square roots of the number of counts in each bin.

A histogram of the luminosities of the RSGs in Figure 3.2 is shown in Figure 3.3a. The figure is a reproduction of Figure 4 in McDonald et al. (2022) which contains luminosities of RSGs in M31. Figure 3.3a features 62 RSGs with $\log_{10}(L/L_{\odot}) > 5.0$ – Figure 4 in McDonald et al. (2022) has 65. The bins match those of the McDonald et al. (2022) figure and so the error bars from this paper can be directly applied onto the bins in Figure 3.3a. The error bars show general agreement for highly luminous stars but diverge at the lower end of the luminosity range. McDonald et al. (2022) found more RSGs in this range than what is seen in the synthesized population. This is a result of the minimum mass constraint: some of the less luminous RSGs in McDonald et al. (2022) likely have initial masses $< 10M_{\odot}$ and there are thus no corresponding stars in the synthesized population.

Figure 3.4 shows analogous plots to those in Figure 3.2 with the same metallicity and population size. However, the default BSE supergiant mass-loss rate prescription is replaced by Equation 2.2. The plots are largely similar in appearance to the ones in Figure 3.2; the only qualitative differences regard the supergiants. The RSGs occupy the same temperature regime as before, though with a slightly smaller spread, but there are now more luminous RSGs above the H-D limit. In addition, there are no BSGs or WRs at all. The Beasor et al. (2020) mass-loss rate is too low for single RSGs to evolve bluewards through loss of the envelope.

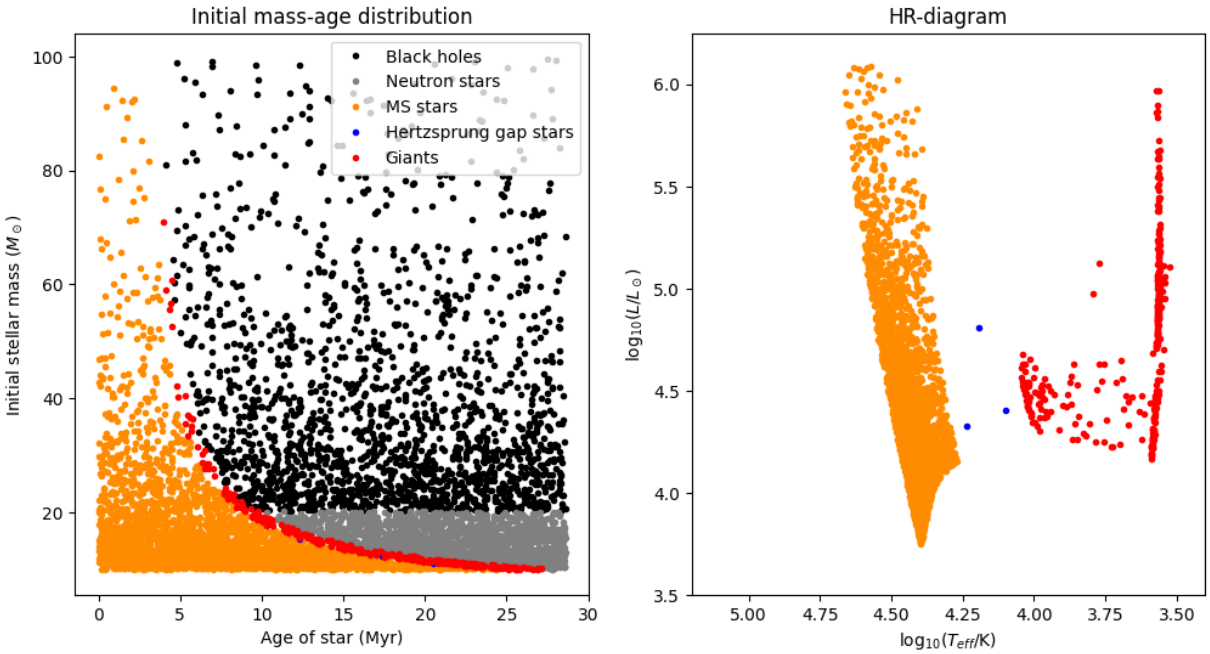

Figure 3.4: As Figure 3.2, but with Beasor et al. (2020) mass-loss rate prescription.

Figure 3.3b is the corresponding RSG luminosity histogram to Figure 3.3a. The RSGs reach higher luminosities than with the default BSE mass-loss rate prescription, and there are fewer of them at the lower luminosity end. Since there is a lack of strong winds, stars with highly massive cores evolve into RSGs, resulting in more luminous RSGs. There is no limiting factor for the RSG luminosities save for the maximum initial mass, and the maximum luminosity in Figure 3.3b is therefore not indicative of any physical H-D limit.

Since the histograms in Figure 3.3 are the results of single sample populations there is a statistical uncertainty on the luminosity of the most luminous RSG, L_{HD}. In order to quantify the uncertainty, the L_{HD} from 100 populations generated with the same parameters were collected. Figure 3.5 shows this for multiple mass-loss rates and at M31, LMC, and SMC metallicities for each rate. The nominal value of each H-D limit is the median value of the 100 maximum RSG luminosities of each data set and the errors are found as one standard deviation from the median. For comparison, two literature values of the H-D limit are included, here assuming complete metallicity-independence. The default BSE mass-loss rate consistently falls within the Humphreys & Davidson (1979) luminosity range. All of the prescriptions predict H-D limits higher than that found by McDonald et al. (2022). While the nominal luminosity values of all prescriptions seem to indicate a downwards trend with metallicity, the error bars for any given prescription overlap with each other at all metallicites. It is uncertain from this figure whether or not the predicted H-D limit is independent of metallicity, as the observed limit is. The figure also includes two additional mass-loss rates that are considered by Beasor et al. (2020): one originating from de Jager et al. (1988), and another from van Loon et al. (2005). Neither prescription

recreates the H-D limit and both display a metallicity-dependence. As both prescriptions produce values of the limit that are intermediate between the default BSE and Beasor et al. (2020) prescriptions and with larger uncertainties while not displaying more observationally motivated behaviors, these prescriptions are not studied further. The extreme cases that are the default BSE and Beasor et al. (2020) rates are chosen instead to represent the entire range of outcomes.

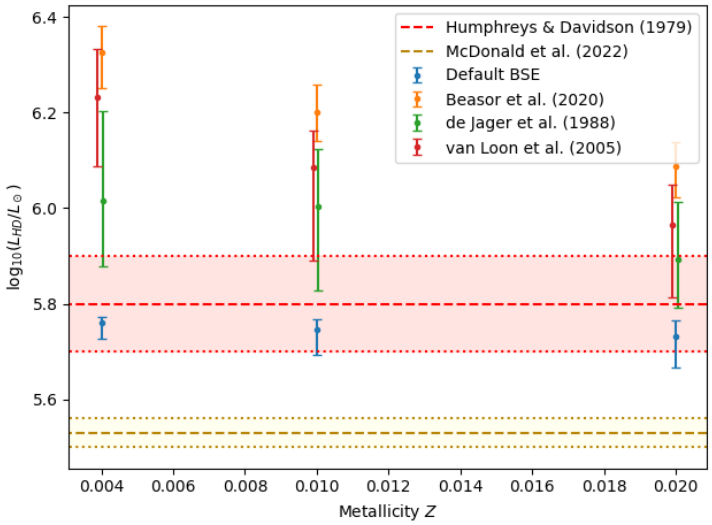

Figure 3.5: H-D limits at different metallicities and with different mass-loss rate prescriptions, based on the maximum RSG luminosities in 100 simulations per datapoint. Also included are the estimates for the H-D limit made by Humphreys & Davidson (1979) for the MW and the LMC, and by McDonald et al. (2022) for M31. The Davies et al. (2018) H-D limits for the LMC and SMC are consistent with the McDonald et al. (2022) value.

3.2 Binaries

Since the default BSE mass-loss rate is engineered to reproduce the H-D limit regardless of the presence of a partner star, it is not suited for use in studying the effect of binarity upon the luminosities of RSGs. For the rest of the thesis the Beasor et al. (2020) mass-loss rate will be assumed to be correct and exclusively used to investigate the binary star systems.

Figure 3.6 was made by altering the input parameters to BSE to reflect a population of binary stars as outlined in Section 2.1.2 and it shows the analogous binary star plots to Figure 3.4. The initial mass-age plot displays the same overall qualitative features as in the single star case, but the boundaries between different regimes of stellar types are less clearly defined. The HR-diagram has been cropped to fit the same limits as in Figure 3.4 and as a result some low-luminosity MS stars are cropped out. These are of little interest as they are low-mass companion stars that cannot reach the RSG phase. The most striking

feature of this HR-diagram is the large number of BSGs and WRs. These stars were not seen at all in Figure 3.4, confirming that binary star interactions can on their own make RSGs evolve bluewards. A larger N was required to fit to the same N_{RSGs} as in the single

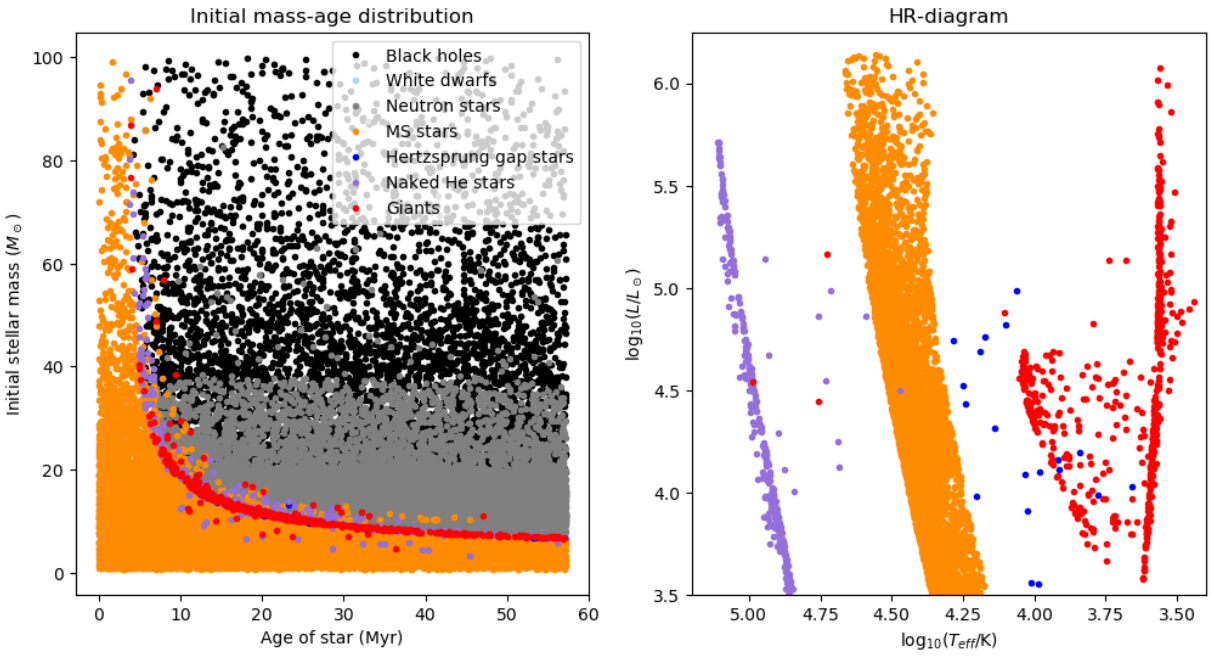

Figure 3.6: Characteristics of a synthesized population of 20800 binary star systems with Beasor et al. (2020) mass-loss rate prescription. M31 metallicity, $Z = 0.02$. Left: Stellar type as a function of initial mass and age. Right: HR-diagram, with stellar types. Both partner stars of the binaries are plotted separately in both plots.

star case, since a greater number of RSGs in Figure 3.6 evolve into BSGs or WRs than in Figure 3.2. The H-D limit is still not reproduced, though: the maximum RSG luminosity seen in the HR-diagram is greater than both the Humphreys & Davidson (1979) value, $\log_{10}(L/L_\odot) \approx 5.8 \pm 0.1$, and the values in Table 1.1.

Figure 3.7 displays the luminosities of RSGs in a much larger synthesized population than the one featured in Figure 3.6. The RSGs have been split into three categories based on certain sorting criteria. If the RSG is the result of a binary system coalescing into a single star, then it is denoted a merger system. If the RSG is not the result of a merger, and is thus still in a binary system, and is currently overfilling its Roche lobe it is denoted a Roche lobe overflow (RLOF) system. If neither of these conditions are true, then the system is denoted effectively single, since the primaries are not interacting with their companions in any meaningful way. The histogram of effectively single systems displays an abrupt fall-off at $\log_{10}(L/L_\odot) \approx 5.5$, which is a convincing recreation of the H-D limit. Most RLOF systems are also found below the limit, but the RLOF histogram lacks a drop-off feature like the effectively single histogram and some RLOF systems are found above the limit. The

merger histogram is very similar in shape and luminosity range as the single star histogram in Figure 3.3b. This is very reasonable as the merger systems are essentially single star systems evolving with the Beasor et al. (2020) mass-loss rate, just as the stars in Figure 3.3b. Taking all three categories together paints a clear picture: any RSGs born into a

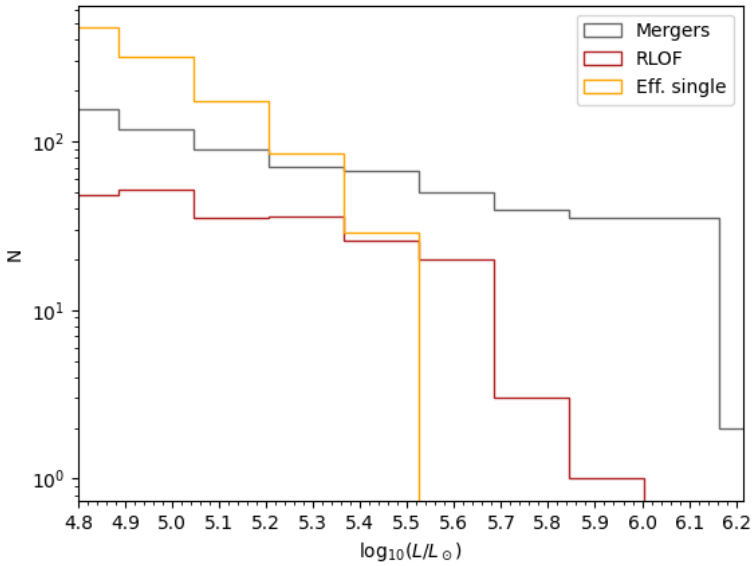

Figure 3.7: Luminosities of 2981 RSGs separated into three categories: RSGs that have merged with the secondary, RSGs with overflowing Roche lobes, and RSGs which have not merged and which are not transferring mass to a companion at current time. The RSGs were evolved with the Beasor et al. (2020) mass-loss rate.

close binary that have not merged or are filling their Roche lobes in the current day obey the H-D limit even with a very conservative mass-loss rate. Mergers and RLOF systems do not. Since the limit is empirical fact, some other mechanism must be preventing these systems from being observed as RSGs, or – if they are observed as RSGs – limiting their luminosities.

Figure 3.8 is similar in appearance to Figure 3.5 and serves a similar function: finding statistical values for the H-D limit at different Z by synthesizing many populations, here for the binary systems. All data in Figure 3.8 was however made with the Beasor et al. (2020) $\dot{M}$ prescription. The data sets are instead split into the three cases also seen in Figure 3.7, and the exact results are tabulated in Table 3.1. The maximum luminosities of the effectively single systems at SMC and LMC metallicites are found to overlap with the H-D limits for those galaxies as found by Davies et al. (2018), while the found result at M31 metallicity falls just below the McDonald et al. (2022) value. Systems with RSGs currently undergoing mass transfer due to Roche lobe overflow and merger RSGs produce maximum

luminosities above the modern values of the H-D limit. There is a downwards trend in maximum luminosity with metallicity if only the nominal values are considered, however all error bars for the effectively single systems overlap for some range of values and likewise for the RLOF systems. The error bars for the merger systems at SMC and M31 both overlap with the LMC error bar, but are themselves distinct from one another. It is inconclusive from this Figure if the simulated binary systems display a metallicity-dependence, as in the single star case.

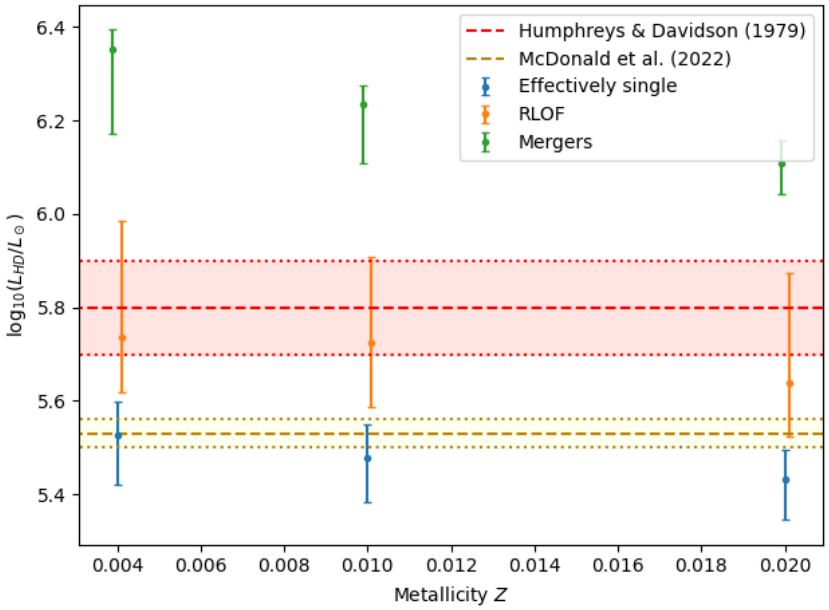

Figure 3.8: Maximum luminosities for each of the three categories of systems (effectively single, RLOF, and merger) at M31, LMC, and SMC metallicities. Each datapoint represents 100 synthesized populations with the nominal value representing the median maximum luminosity per category and the error bars one standard deviation from the median. Like in Figure 3.5, two literature values of the H-D limit are included.

Table 3.1: Figure 3.8 results, tabulated.

Z	$\log_{10}(L_{\text{HD, eff. single}}/L_\odot)$	$\log_{10}(L_{\text{HD, RLOF}}/L_\odot)$	$\log_{10}(L_{\text{HD, mergers}}/L_\odot)$
0.020	$5.43^{+0.06}_{-0.09}$	$5.64^{+0.23}_{-0.12}$	$6.11^{+0.05}_{-0.07}$
0.010	$5.48^{+0.07}_{-0.09}$	$5.73^{+0.18}_{-0.14}$	$6.23^{+0.04}_{-0.13}$
0.004	$5.53^{+0.07}_{-0.11}$	$5.74^{+0.25}_{-0.12}$	$6.35^{+0.04}_{-0.18}$

In order to characterize the differences between merger, RLOF, and effectively single systems Figures 3.9, 3.10, and 3.11 display corner plots of some initial and final properties of the different system types. The same data set is used for all three Figures in order to facilitate direct comparison and it was made very large to avoid statistical deviations dominating the results. The final luminosity histograms featured in all three Figures are

as the ones in Figure 3.7 though with this larger set of data.

The initial mass fraction distributions in the three Figures all have differing features. In the merger case, initial q is split into two peaks separated by a trough. The majority of merger systems have a small initial q with a relatively small initial primary mass and orbital semi-major axis, implying that most merger systems are the result of a primary merging with a very low-mass companion in a close orbit. There is a sizeable minority, though – the second peak in q – of merger systems that are the result of the merging of two similarly massive stars that are exclusively in very small orbits. The average final primary mass is lower than the average initial primary mass which may initially seem contradictory for the only category of systems where mass is actively transferred to the primary. However, since most mergers are with a low-mass star it is not expected that the merging event would have a considerable impact on the mass of the primary and while the Beasor et al. (2020) mass-loss rate is quite modest it is not zero. The mergers that are the result of similarly massive stars are seen as a small increase at high mass in the histogram of final masses.

As for the Roche lobe overflow systems, initial q is found in a quite flat distribution, though exclusively at $q \gtrsim 0.3$. Since the RLOF systems are found in a very narrow $\log_{10}(a/R_\odot)$ range that overlaps with the low-q merger systems the most likely explanation for the lack of low-q RLOF systems is that any system in that range which might have become a RLOF system merges instead. It can also be seen in all histograms in Figure 3.10 that the total number of RLOF systems is the lowest among the three cases. The final primary mass is on average lower than the initial primary mass for these systems. This is expected, as the systems are defined by an ongoing mass transfer from the primary onto its companion.

The effectively single systems have an initial q distribution that is very flat and uniform between 0 and 1 and are thus the most representative of the input parameters in this regard. The initial semi-major axis distribution is distinct from the merger and RLOF systems: there is no overlap in the $\log_{10}(a/R_\odot)$ range between them and the effectively single. The lower bound of this distribution marks a difference in evolutionary behavior, while the upper bound is an artefact of the input parameters. There is little difference between the initial and final primary mass distributions for these systems, but the final mass is on average slightly smaller due to the mild mass loss.

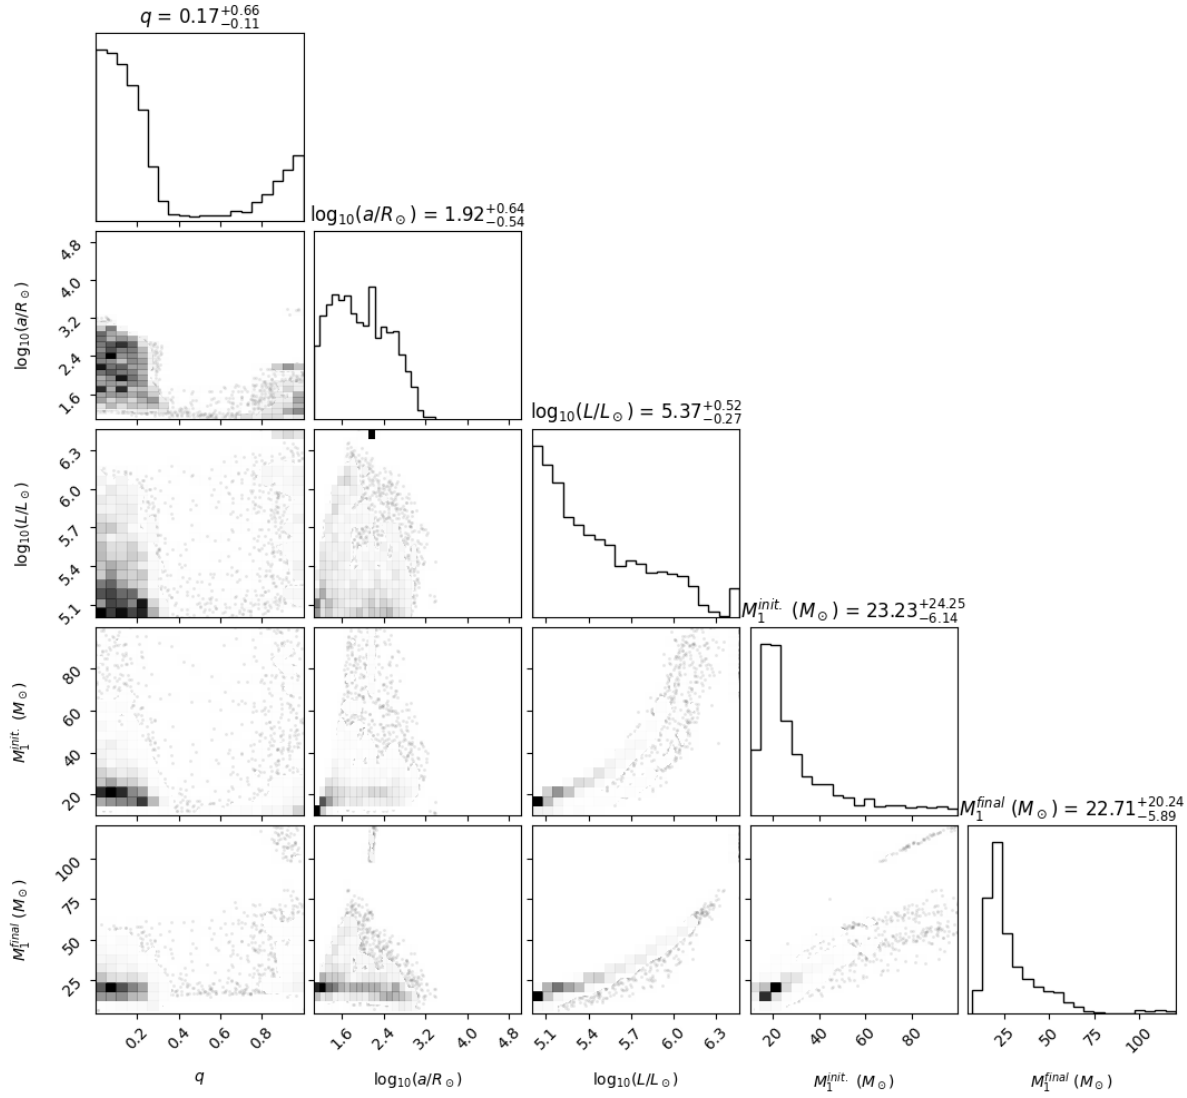

Figure 3.9: Corner plot of stellar parameter distributions for systems where the primary is an RSG that has merged with its companion: the initial mass fraction, initial semi-major axis, final luminosity, initial primary mass, and final primary mass. The data set consists of all such systems from 100 $N = 20800$ populations synthesized at $Z = 0.02$ with the Beasor et al. (2020) $\dot{M}$ prescription for RSGs.

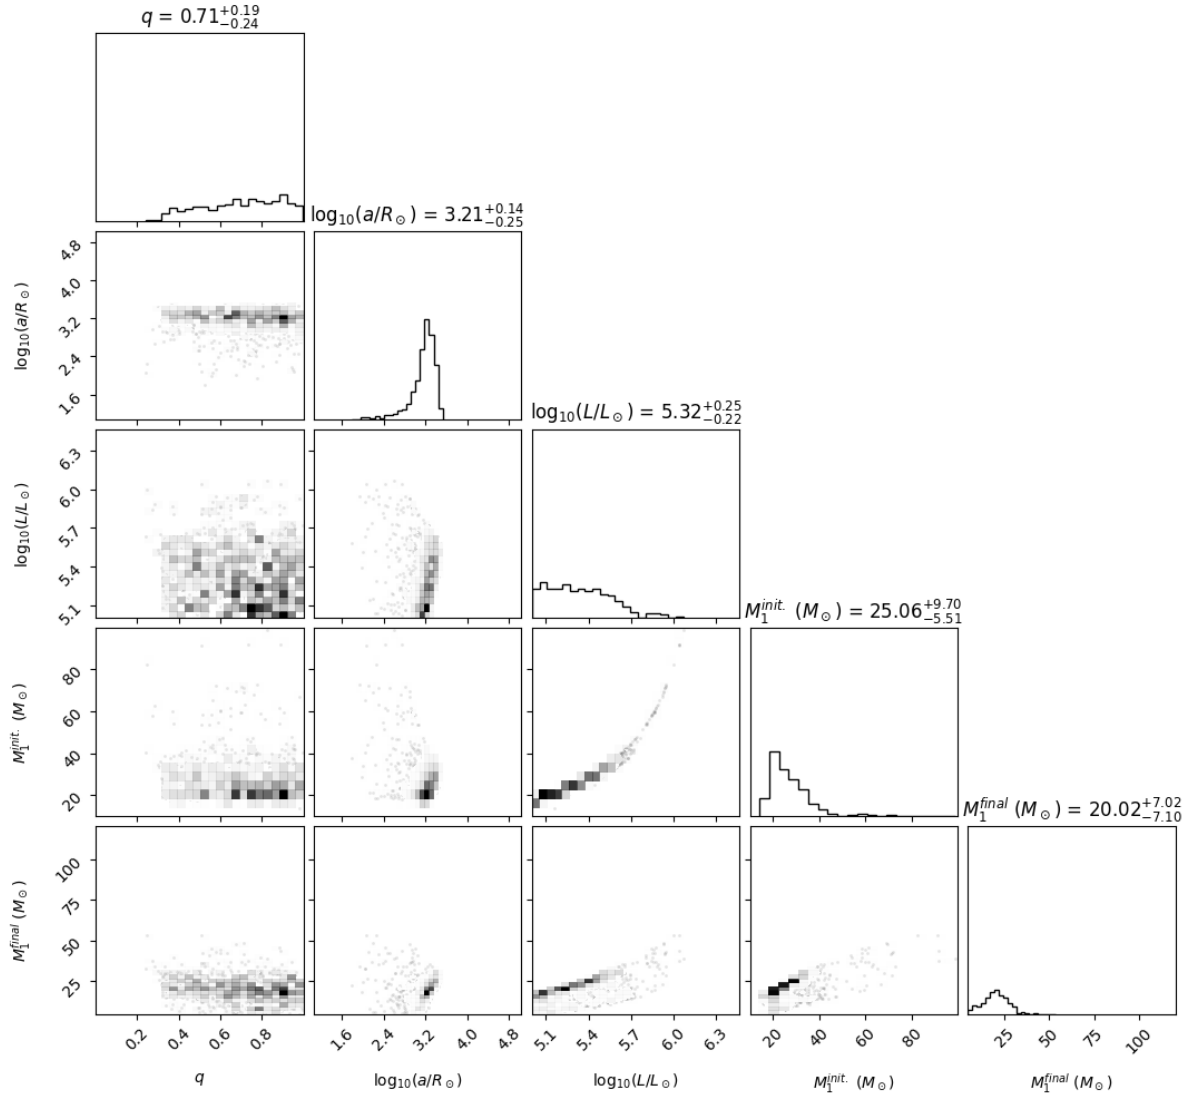

Figure 3.10: Corner plot of the same stellar parameter distributions as in Figure 3.9 for systems where the primary is an RSG that is currently overfilling its Roche lobe. The axes are the same as in Figure 3.9. The data set consists of all such systems from the same 100 populations used for Figure 3.9.

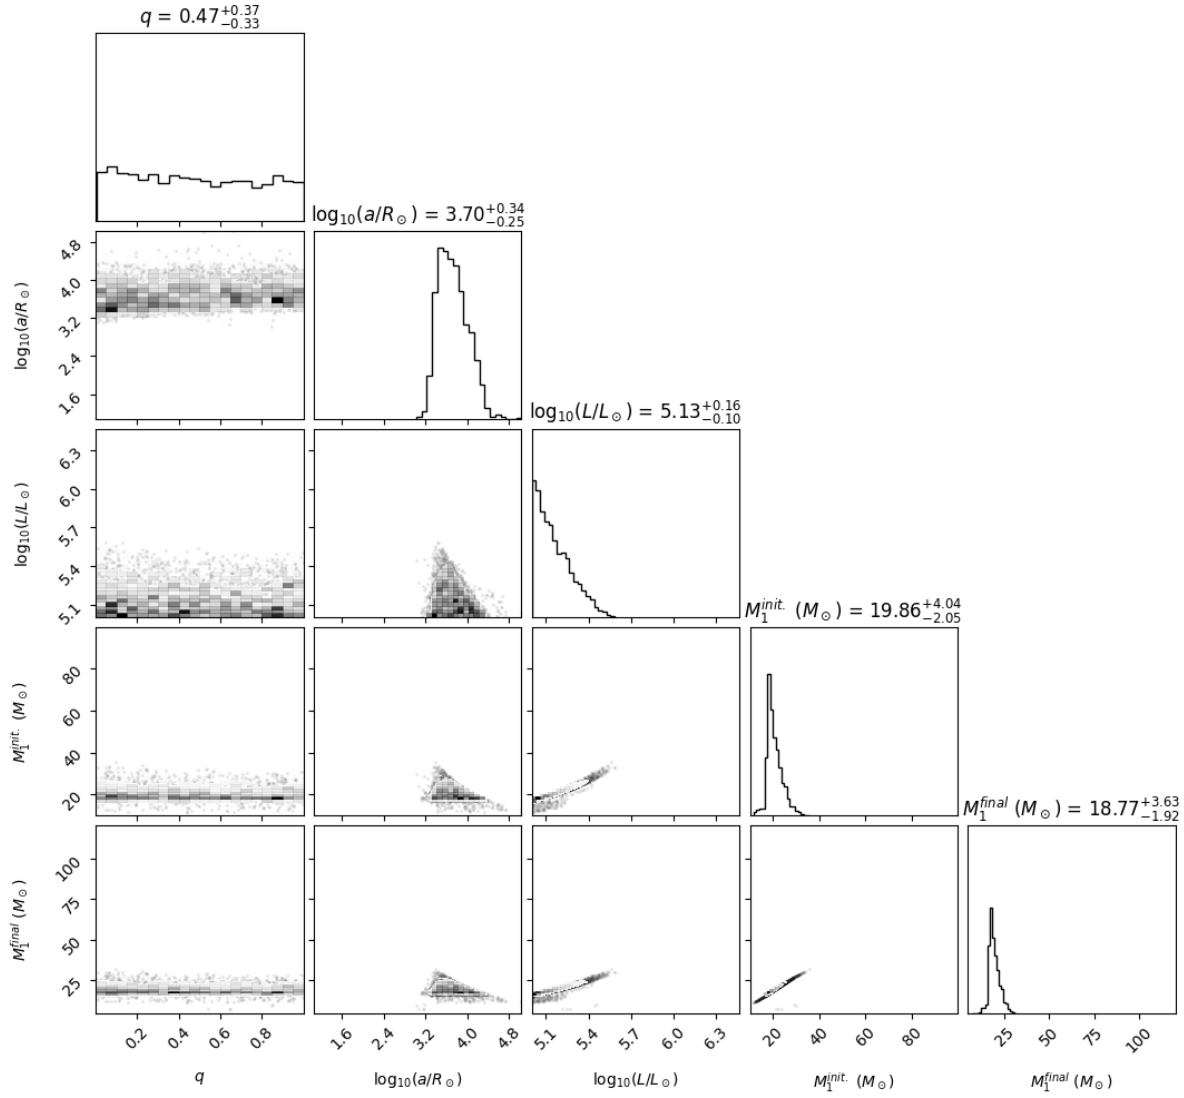

Figure 3.11: Corner plot of the same stellar parameter distributions as in Figures 3.9 and 3.10 for systems where the primary is an RSG that is not filling its Roche lobe and which has not merged with its companion. The axes are the same as in Figures 3.9 and 3.10. The data set consists of all such systems from the same 100 populations used for Figures 3.9 and 3.10.

Chapter 4

Analysis

4.1 Single stars

A foundational idea of the thesis is that the default BSE and Beasor et al. (2020) $\dot{M}$ prescriptions are distinct and describe two different models of RSGs. Figure 3.1 reveals this to be true for stars with masses in excess of $25M_{\odot}$ – which is the relevant mass regime for exploring the H-D limit.

BSE as it stands reproduces the Humphreys & Davidson (1979) value for the H-D limit for single stars as can be seen in Figure 3.3a. This is unsurprising. The $\dot{M}_{\mathrm{LBV}}$ term added onto large, luminous stars is specifically introduced by Hurley et al. (2000) in the context of the H-D limit and enforces the luminosity of RSGs to remain below the limit as it was found by Humphreys & Davidson (1979). Since the term is introduced to enforce this value of the limit it is expected that default BSE reproduces it but not the value found by McDonald et al. (2022). Disabling the $\dot{M}_{\mathrm{LBV}}$ term and applying the more recent and observationally motivated $\dot{M}$ prescription from Beasor et al. (2020) removes the limit from the single star results. If LBV eruptions are ubiquitous for massive stars at all metallicities – as BSE assumes by default – then single star evolution could still explain the H-D limit. If not, some other mechanism is needed – assuming that the Beasor et al. (2020) $\dot{M}$ prescription is correct.

4.2 Binaries

With populations whose properties are based on probability distributions comes uncertainty regarding whether any given run is a statistical anomaly or genuinely indicative of real behavior. Both Figures 3.6 and 3.7 display many RSGs above the H-D limit, but in order to test whether violation of the limit is a statistical certainty the number of RSGs per run that violate the limit was plotted in Figure 4.1. A population without RSGs above the limit is never seen. Figures 3.9, 3.10, and 3.11 convincingly show that the violators belong to the RLOF and merger systems, and the differences between the effectively single,

RLOF, and merger systems are discussed in the following subsections.

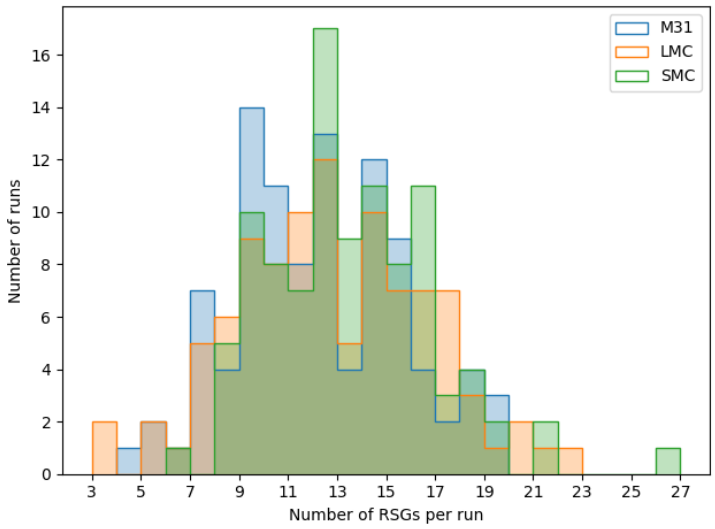

Figure 4.1: Histogram of number of RSGs that are above the H-D limit in each synthesized population, totalling 100 populations with $N = 20800$ per galaxy metallicity. See Table 1.1 for the Z used for each galaxy.

4.2.1 Effectively single systems

The maximum luminosities of the effectively single RSGs evolved with the Beasor et al. (2020) $\dot{M}$ prescription match the H-D limit. The hypothesis of this project is partially proven correct: binary star interactions can reproduce the H-D limit, but only for systems that do not merge or are currently overfilling their Roche lobes. Any star that remains non-interacting evolves into a RSG below the limit, so by logical inverse any star that would evolve into a RSG above the limit must interact with its partner star. The effectively single RSGs are constrained to $\log_{10}(a/R_\odot) \gtrsim 3.2$, see Figure 3.11, and are constrained from above by the assumption that all RSGs in binaries are in close binaries. This regime is mutually exclusive with the $\log_{10}(a/R_\odot)$ regimes for Roche lobe overflow and merger system, see Figures 3.9 and 3.10. Those evolutionary paths are closed to stars in the effectively single $\log_{10}(a/R_\odot)$ regime. This shows that such stars that would have become RSGs above the H-D limit must have interacted with their companion in a way that prevented them from evolving to RSGs at all. The vast number of BSGs and WRs in the binary sample, Figure 3.6, reveals what most likely happened to these stars: they evolve bluewards following mass transfer.

As outlined in the introduction it is observationally established that the close binary

fraction of massive stars approaches unity as mass increases, and the upper bound of $\log_{10}(a/R_\odot)$ follows from the assumption that all RSGs are in close binaries. The lower constraint on the semi-major axes seen in Figure 3.11, however, does not appear to be reflected in reality. The behaviors seen for smaller values of $\log_{10}(a/R_\odot)$ – Roche lobe overflow and merging – can not be dismissed and will now be discussed.

4.2.2 RLOF systems

The Roche lobe overfilling systems are very constrained in $\log_{10}(a/R_\odot)$ and are thus rare, reflected by the relatively low number of such systems seen in e.g. Figure 3.7. Any massive system within the proper regime will additionally only spend a small fraction of its lifetime in a RLOF state, as mass transfer occurs on dynamical timescales. Observing a binary system with an RLOF RSG might be a very rare occurence, but correctly identifying it as such might be even more challenging. By their very nature RLOF systems feature relatively short semi-major axes – though not so short as to facilitate merger – and at least one star with a very large radius relative the semi-major axis, with an ongoing trail of mass between the stars. A system like this is not resolvable with direct imaging at intergalactic distances. The angular size of an orbit with $\log_{10}(a/R_\odot) = 3.2$ is on the order of 10^{-4} arcsec for systems in the SMC and LMC and 10^{-5} arcsec in M31. Using current telescopes – like the *Spitzer* telescope that McDonald et al. (2022) based their findings on – an RLOF system would not be straightforwardly recognized as a binary system, and then the question becomes if the primary would register as an RSG. The companions of RLOF RSGs are typically rejuvenated MS stars, see Figure 4.2, and if both RSG and companion are believed to be one star then the color of that star may not be recognized as red enough to be a RSG. RSGs are found in a very constrained temperature regime, so if the color is being 'averaged out' by a MS companion then it would likely not make the color cut for an RSG. In that case the binary would be identified as some type of YSG. Quantifying this is outside the scope of this thesis, but it is certainly possible to estimate what the observed color of the RLOF RSG systems would be from the luminosity and temperature output from BSE. Inspection shows that in the majority of RLOF systems the luminosity of the primary and the companion are of the same order of magnitude, so it is reasonable that an unresolved companion could have a significant effect on the perceived color of an RLOF RSG.

4.2.3 Mergers

There are fundamentally two different types of merger systems revealed by Figure 3.9. On the one hand there are systems with low q, where the primary mass is typically $\sim 25M_\odot$. These are O-type stars with companions that are a few times more massive than the sun and orbital semi-major axis below $\log_{10}(a/R_\odot) = 3.2$. For the majority of such systems, where the primary mass is not much larger than $30M_\odot$, the H-D limit is not violated. The other type of system is that of two similarly massive stars with masses that make them both O-type stars. These systems are rarer than the first type, but almost all of them

produce a merged star that evolves into a RSG above the limit. They also have a lower upper limit for the orbital semi-major axis than the first type at $\log_{10}(a/R_\odot) = 2.4$.

Both types of systems have relatively small orbital semi-major axes, corresponding to orbital periods that would classify most as Case A mass transfer systems and some as early Case B. These systems interact – and presumably, merge – while still on the MS or shortly thereafter. Sana et al. (2012) found that $\sim 24\%$ of stars born as O-types and that are in binary systems merge with their partner star. Moe & Di Stefano (2017) found a flat q distribution for small orbital periods, supporting that both the low-q and high-q merger systems could exist in reality. It must be concluded that both types of merger systems must exist in reality and are common enough that they should be observed. For binary systems with an O-type primary and a companion in a close orbit, binary star interactions are not sufficient to solely explain the H-D limit.

4.3 Comparison to observations

Patrick et al. (2022) studied binary systems in the SMC with RSGs as primaries. The paper features an HR-diagram of the RSGs and their companions. Figure 4.2 is based on that diagram, and features the data points from Patrick et al. (2022) as well as RSGs and companions from one BSE population at $Z = 0.004$. The BSE binaries have been split into the usual categories, though merger systems are excluded on account of them lacking a companion star. The effectively single and Patrick et al. (2022) companions are mainly found along the MS and the populations agree very well with each other. The RLOF companions are also found on the MS but at higher luminosities and with a much later turn-off point. This is due to the MS companions receiving mass and being rejuvenated to resemble younger, more massive MS stars.

The BSE and Patrick et al. (2022) RSGs differ in their luminosity ranges, however. The minimum luminosity bound differing is a result of the minimum mass cut used for BSE sample and is not indicative of a real difference in behavior. The difference in behavior close to the maximum luminosities is more puzzling. The BSE sample produces a higher maximum luminosity regardless of whether RLOF systems and systems Patrick et al. (2022) would not detect are included, but this could be explained as a matter of probability and sample size if the trends were otherwise similar. They are not. The BSE RSGs are numerous all the way up to their maximum bound, while the Patrick et al. (2022) RSGs gradually decline in number as a function of luminosity. Figure 4.3 compares this slope with a corresponding slope fitted to the RSG source catalogue used by Patrick et al. (2022) – which includes SMC RSGs not in confirmed binaries – and corresponding slopes fitted to a large BSE sample of effectively single RSGs at SMC metallicity. The Figure shows that the slope of confirmed RSG binaries in the SMC is significantly steeper than the others. The BSE slopes are very similar to the source catalogue slope, suggesting that the choice of IMF and the basic stellar evolution in BSE are correct. Patrick et al. (2022) accounted

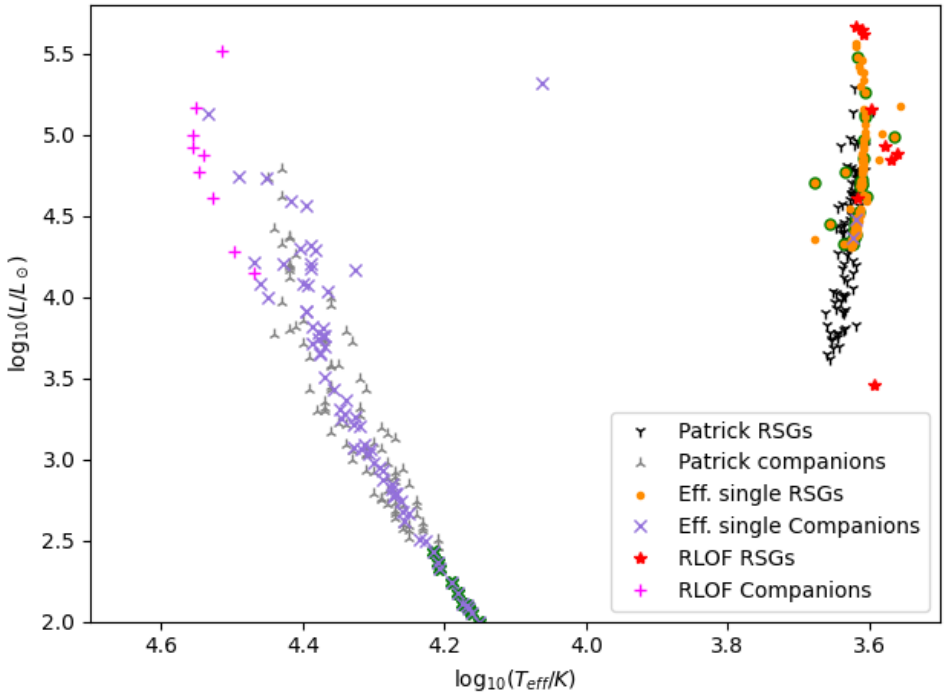

Figure 4.2: HR-diagram of RSGs and companions, that have not merged, from one synthesized population at SMC metallicity. The Figure is a based on Figure 5 in Patrick et al. (2022), though here the RSGs and companions have been split into effectively single and RLOF systems. Systems that are below the q/M_{RSG} detection limit from Patrick et al. (2022) are marked by a green outline. Included for reference are observed systems from Patrick et al. (2022).

for highly luminous RSGs 'hiding' faint companions with a detection limit characterized as a lower bound on q, and removing the systems from the BSE samples which fall below the detection limit makes no qualitative difference upon the BSE slope. It must be concluded that the luminosity function of RSGs in confirmed binaries in the SMC does not behave as the luminosity function of all known in the RSGs in the SMC – but that the BSE effectively single binary population does. Discovering the difference between the binary sample and full catalogue requires further work as it stands.

When comparing the BSE results to observations, the final RSG population is only one side of the coin. If it is believed that binary star interaction is the primary driving factor behind massive stars evolving towards higher effective temperatures, then the resulting BSG and WR population should also match observations. Massey et al. (2021) compared the number ratio of RSGs and WRs in multiple nearby galaxies to evolutionary models, and in order to do so compiled observational counts of RSGs and WRs in those galaxies. Figure 4.4a compares RSG/WR ratios for M31, the LMC, and the SMC calculated from data in Table 2 in Massey et al. (2021) to corresponding ratios in BSE populations. BSE

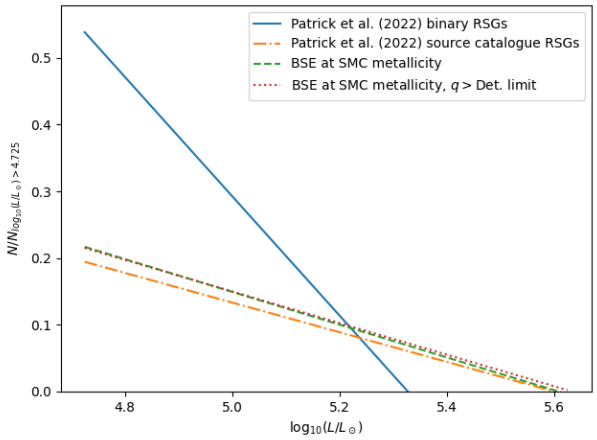

Figure 4.3: Slopes fitted to the bar heights of luminosity histograms of the RSG primaries in Patrick et al. (2022), the RSGs in the source catalogue for Patrick et al. (2022), and the effectively single RSGs from 100 BSE populations at SMC metallicity. The SMC population is fitted to twice: once as is, and once with systems undetectable by Patrick et al. (2022) removed – denoted by '$q >$ Det. limit'. Each histogram was normalized to the number of RSGs in that histogram with luminosities $\log_{10}(L/L_{\odot}) > 4.725$, the minimum luminosity at which BSE does not produce any AGBs assuming minimum mass $M = 10M_{\odot}$. Luminous AGBs may be observed above this threshold, but it is assumed the Patrick et al. (2022) sample is free of AGBs.

does not identify WRs directly, rather they make up some subset of the "Naked Helium" stars. Massey et al. (2021) cites Crowther (2007) who states an effective temperature cut of $T_{\text{eff}} > 30,000 - 100,000$ K and a surface mass fraction of hydrogen $X_H < 0.3 - 0.4$ as theoretical cuts for WRs. While the temperature cut is straightforwardly implemented, BSE unfortunately does not provide the surface fraction of hydrogen X_H for the stars it evolves. Massey et al. (2021) also mentions that Eldridge et al. (2017) uses a luminosity cut $\log_{10}(L/L_{\odot}) > 4.9$ to distinguish WRs. That cut is used here to provide another constraint.

While there is relatively good agreement between Massey et al. (2021) and the BSE populations for M31 and the LMC, there is a large discrepancy for the SMC. In Massey et al. (2021) there are few identified WRs in the SMC. In the BSE populations there are typically many more. If the BSE populations were to be restricted in size so that the numbers match there would likely be a larger spread in the BSE ratio values for the SMC, but it seems exceedingly unlikely that the increased spread would be able to alone account for the discrepancy. BSE seems to predict a RSG/WR fraction close to unity while observations point to an order of magnitude in difference. Figure 4.4b reveals that all three metallicities produce both WR temperature distributions and number of WRs above the threshold that are very similar.

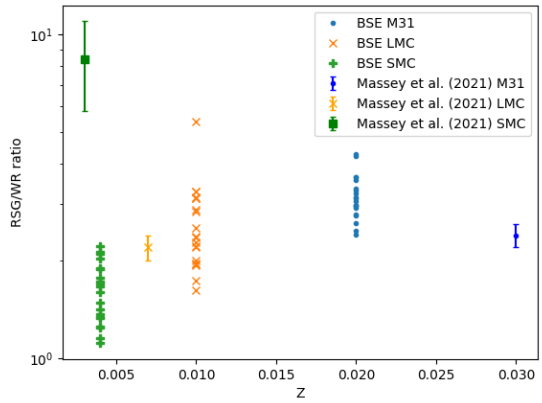

(a) Comparison of the Massey et al. (2021) RSG/WR ratios to the BSE ratios. The BSE ratios are located at the Z they were found at, and the paper ratios at the Z given in the paper for those galaxies.

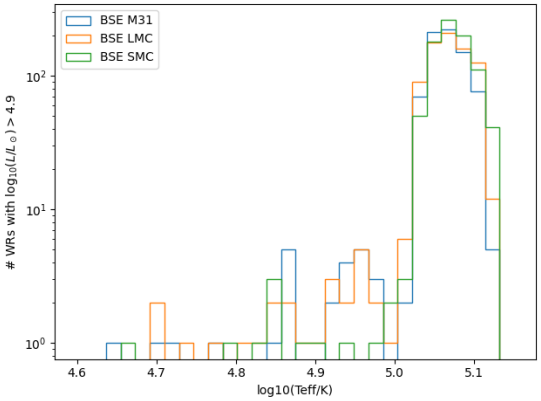

(b) Temperature distributions of all WRs above the luminosity threshold from all 20 runs per galaxy.

Figure 4.4: RSG and WR statistics from 20 BSE populations per value of Z (M31, LMC, and SMC, see Table 1.1) along with ratios from Massey et al. (2021) for comparison. The $\log_{10}(L_{\rm RSG}/L_\odot) > 4.8$ RSG count from the paper is used. For the WR stars, a luminosity cut $\log_{10}(L_{\rm WR}/L_\odot) > 4.9$ is used. The paper ratios and error bars are calculated from Table 2 in the paper and using the 'classical' procedure outlined in Section 3.3 of the paper.

There are two possibilities explaing this. Either the WR population should be approximatively independent of metallicity, as BSE predicts, and BSE generates too few RSGs at SMC metallicity. Or, the WR population should rapidly decrease at low metallicities which is not seen in the BSE result.

The former alternative for explaining the SMC RSG/WR ratio seems unlikely. Given that McDonald et al. (2022) found that the luminosity function of RSGs is independent of metallicity and the similarity of Figures 3.9, 3.10, and 3.11 with their LMC and SMC counterparts in Appendix A, it seems to be the case that RSG evolution is independent of metallicity. Also, Figure 4.3 implies strong agreement between SMC RSGs and the BSE population. An argument can be made for the latter alternative, though: the low hydrogen fraction in WR surfaces implies that most of the envelope has been lost. Stellar winds have mostly been ignored in this thesis on the grounds that Beasor et al. (2020) ruled them insufficient to explain the H-D limit, but a lower metallicity should correspond to weaker stellar winds and thus less envelope lost. It could be the case that while winds are unable to on their own make an RSG evolve towards the blue, they are still strong enough to dictate whether the final product becomes a WR or not. Eldridge & Stanway (2022) state that the wind strength of a WR is dependent on the initial iron abundance

of the star in their review on mass-loss rates in stellar winds – and that lower metallicity stars typically have weaker stellar winds. For future work, the metallicity-dependence of the strength of winds in WR stars should be investigated to quantify this possibility.

4.4 Discussion

The H-D limit is despite the findings regarding RLOF and merger systems empirical fact. If binary star interactions are unable to entirely account for it, then some other mechanism must be at work. Non-conservation of core mass may be such a mechanism. Stellar mergers in BSE are not universally mass conservative, but the core mass is always conserved. That is, the core mass of the final star is the sum of the core masses of the pre-merger stars. Since $L_{\mathrm{RSG}} \sim M_{\mathrm{core}}$ the final RSG luminosity is proportional to the sum of the core masses of the pre-merger stars. This is perhaps a sound argument: as the most dense parts of the pre-merger stars the cores are probably unlikely to be flung into space rather than sink into the core of the merged star. But it does seem worthwhile to investigate hydrodynamical simulations of massive O-type stars merging with companions – both low-mass and another O-type – to find some relation of core masses before and after merger. Figure 3.9 provides constraints for stellar and orbital parameters one could use in such an investigation. If it is found that core mass should not be conserved, then the luminosities predicted by BSE for merged RSGs would be too high. As is stands, this is a task for future work.

Another explanation as to why the highly luminous merged systems are not seen in reality could be that many massive stars are found in higher order multiple systems. Moe & Di Stefano (2017) found that the majority of O-type MS stars are in triple or quadruple systems, and it is binary systems of such massive stars that may produce merged stars that evolve to RSGs above the limit. One could imagine a triple system wherein the massive O-types merge into a star capable of violating the limit, but interactions with the tertiary star then prevents this just as for the non-merging binary systems. A quadruple system that contains two binary systems within it would reduce to a triple system, which could then interact. From the multiplicity statistics in Table 13 in Moe & Di Stefano (2017) a rising trend in multiplicity as a function of primary mass can be inferred, with primary mass $> 16 M_{\odot}$ corresponding to a triple or quadruple fraction of 0.73 ± 0.16. If this fraction continues towards unity as primary mass increases, then for the primary mass regime that produces H-D limit violating mergers, $M > 30 M_{\odot}$, essentially all systems should be higher order multiple systems. Interactions with tertiary or quaternary companions in systems where the central binary has merged should certainly be a focus for future work in this area.

Finally, rotation-induced mass loss may also contribute to solving the problem of the merged systems. Merger products should generally be rotating rapidly in order to conserve the angular momentum of both pre-merger stars and their orbit. Ekström et al. (2012) found that mass loss increases with rotation, and that mass loss caused by rotation becomes more dominant as stellar mass increases. This causes much angular momentum

to be lost from the star, lowering its rotational speed. Highly massive merger products should therefore experience a period of intense mass loss caused by their rapid rotation, before reaching a more stable final mass once the rotation has sufficiently decreased. Figure 3.9 reveals that all merger systems that surpass the H-D limit are very massive with final masses $> 30M_{\odot}$, and are thus eligible for this high mass loss. No quantitative analysis of this is performed here, but it is possible that rotation-induced mass loss causes merger products to always have final masses in the range correlated with luminosities below the H-D limit.

As mentioned in the introduction, McDonald et al. (2022) performed some numerical experimentation with the `BPASS` binary star evolution code. Below $\log_{10}(L/L_{\odot}) \approx 5.5$ the histogram in Figure 3.7 of effectively single stars has a very similar shape as the one suggested by the luminosity function of binaries in McDonald et al. (2022), which was generated by `BPASS`. Above $\log_{10}(L/L_{\odot}) \approx 5.5$, the RLOF histogram is more reminiscient of the `BPASS` result. The McDonald et al. (2022) function ultimately shows an H-D limit at $\log_{10}(L/L_{\odot}) > 5.95$. This is consistent with the RLOF systems in Figure 3.7. The `BPASS` luminosity function seems to fit very well to a superposition of the effectively single and RLOF histograms, but the very high luminosities seen as the result of certain mergers are not seen in the `BPASS` results. The `BPASS` results in McDonald et al. (2022) are not extensively commented upon, making it difficult to quantify from where this difference arises. Either `BPASS` handles mergers in a different way from BSE, or mergers are not a part of the McDonald et al. (2022) luminosity function.

It should be noted that BSE might not be the best choice for investigating stars at the upper end of the mass range considered here. The formulae in BSE are the results of fitting to evolutionary tracks between 0.5 and 50 $M_{\odot}$ (Hurley et al. 2000). Many – though admittedly not all – of the systems that end up with primaries above the H-D limit have initial primary masses higher than this. The authors restricting stellar masses to a maximum of $100M_{\odot}$ as a result of stellar interaction is another indication that the fits in BSE are not intended for use in this range, especially not for the merger of two massive O-types. BSE is a fast and easily usable evolutionary code that offers itself to modification and use in a wide range of metallicites and masses, but further work on massive O-types in binaries may consider a code that is explicitly written to handle that mass range. If BSE is used further, then the specifics of the mass transfer and merger physics employed by BSE should be investigated. It is possible that a different, more modern, treatment of the evolution of interacting binaries could alleviate the degree to which the merger systems violate the H-D limit. This is something to consider for future analysis.

4.5 Conclusions

The results of the single star simulations confirm the conclusions of Beasor et al. (2020): their derived mass-loss rate is insufficient for RSGs to evolve into BSGs or WRs as single

stars. Equivalent stars that with the default BSE mass-loss rate prescription evolved bluewards remained on the RSG branch when evolved with the Beasor et al. (2020) mass-loss rate prescription, and as a result reached luminosities above the H-D limit. Running binary system simulations with the Beasor et al. (2020) mass-loss rate prescription did not wholly resolve this. It was however discovered that any RSGs that remain in close binaries and that are not currently overfilling their Roche lobes do obey the H-D limit where similar single stars did not.

The RLOF systems seem to be able to surpass the H-D limit. They are, however, a rare type of system with a degree of uncertainty as to whether they would be observationally classified as RSGs, as parts of binaries or otherwise. It is possible that these superluminous RLOF systems exist in reality but that few have been observed and none identified as such. As for the systems that merge, there is no strong case against their existence with the properties that allow them to violate the H-D limit. Such systems probably do exist and do merge into extremely massive O-type stars, but some other mechanism keeps them from evolving into RSGs above the H-D limit. Likely candidates for this include higher order multiples, which are increasingly common at higher stellar masses, and mass loss caused by rapid rotation. Hydrodynamical simulations of mergers of such stars may also be of interest to fully characterize them.

The numerical results for the effectively single populations agree well with observations for the LMC and M31, but for the SMC there are discrepancies which are difficult to explain with the results in this thesis alone. The RSG/WR ratio is seemingly too low for SMC metallicity, possibly due to BSE generating too many WRs in this regime than what would be observed. The companions of RSGs in confirmed binaries in the SMC are seemingly well-represented by the BSE results, but the RSG primaries are not. Their luminosity function approaches zero at a steeper pace in reality than both the SMC RSG population at large and what BSE predicts – the reasons for this are unclear. The strong agreement between the total SMC RSG population and the synthesized effectively single binary population is intriguing and offers a strong argument for continuing to use binary systems to model massive stellar populations.

Ultimately, the H-D limit can not be fully explained by binary star interactions. Binary interactions can, however, reduce the problem of the limit from one encompassing all RSGs to one concerning a subset of RSGs consisting of some RLOF systems and the most massive merger events. These systems probably do exist in reality, but some as of yet unknown process acts on them to make their final characteristics like those of the effectively single population. For future work, the most immediately clear candidates for this process are higher order multiples and mass loss from rotation, and these should be investigated further. Though the mystery of the H-D limit has not been solved, multiple effects of binarity on the evolution of massive star have been explored and analyzed. If nothing else, an increased understanding of these important systems can be gleaned from the results.

Acknowledgements

Above all, Dr. Maxwell Moe of the Steward Observatory, University of Arizona, should be acknowledged, as the foundational idea of this thesis – that binary star interactions may be the cause of the H-D limit – originally came to my supervisor from him. Additionally, unpublished work by Moe on binaries was invaluable for setting up the parameter generation. I would also like to thank Dr. Jarrod Hurley for making BSE publically available, and Michela Mapelli at the University of Padova for her online available slides documenting the use of BSE in an easy-to-understand manner, without which the early parts of the project would have been considerably more taxing. In a similar vein, I thank my supervisor Dr. Ross Church of the Lund Observatory for all the time and effort he spent – much more than he needed to – on helping me progress on and understand my project on a deeper level. Finally, on a personal note, I want to thank all those without whose support I would not have made it so far as to finish a bachelor's thesis: Frida Artman, Melvin Tham, Benjamin Dahlén, Oskar Nilsson, Claudia Skoglund, Matilda Skantz, Viktor Lehmann, Johan Holmberg, Elliot Winsnes, Johann Schmand, and Filip Gustavsson. Thank you all.

Bibliography

Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016, Phys. Rev. Lett., 116, 061102

Beasor, E. R., Davies, B., Smith, N., et al. 2020, MNRAS, 492, 5994

Crowther, P. A. 2007, ARA&A, 45, 177

Davies, B., Crowther, P. A., & Beasor, E. R. 2018, MNRAS, 478, 3138

de Jager, C., Nieuwenhuijzen, H., & van der Hucht, K. A. 1988, A&AS, 72, 259

Drout, M. R., Massey, P., Meynet, G., Tokarz, S., & Caldwell, N. 2009, ApJ, 703, 441

Ekström, S., Georgy, C., Eggenberger, P., et al. 2012, A&A, 537, A146

Eldridge, J. & Tout, C. A. 2019, The structure and evolution of stars (World Scientific Publishing Europe Ltd)

Eldridge, J. J. & Stanway, E. R. 2022, ARA&A, 60, 455

Eldridge, J. J., Stanway, E. R., Xiao, L., et al. 2017, , 34, e058

Gull, T. R., Hillier, D. J., Hartman, H., et al. 2022, ApJ, 933, 175

Hilditch, R. 2001, An Introduction to Close Binary Stars (Cambridge University Press)

Humphreys, R. M. & Davidson, K. 1979, ApJ, 232, 409

Humphreys, R. M., Davidson, K., & Smith, N. 1999, PASP, 111, 1124

Hurley, J. R., Pols, O. R., & Tout, C. A. 2000, MNRAS, 315, 543

Hurley, J. R., Tout, C. A., & Pols, O. R. 2002, MNRAS, 329, 897

Kippenhahn, R., Weigert, A., & Weiss, A. 2012, Stellar Structure and Evolution (Springer-Verlag)

Kudritzki, R. P. & Reimers, D. 1978, A&A, 70, 227

Massey, P., Neugent, K. F., Dorn-Wallenstein, T. Z., et al. 2021, ApJ, 922, 177

McDonald, S. L. E., Davies, B., & Beasor, E. R. 2022, MNRAS, 510, 3132

Moe, M. & Di Stefano, R. 2017, ApJS, 230, 15

Nieuwenhuijzen, H. & de Jager, C. 1990, A&A, 231, 134

Patrick, L. R., Thilker, D., Lennon, D. J., et al. 2022, MNRAS, 513, 5847

Podsiadlowski, P., Joss, P. C., & Hsu, J. J. L. 1992, ApJ, 391, 246

Sana, H., de Mink, S. E., de Koter, A., et al. 2012, Science, 337, 444

van Loon, J. T., Cioni, M. R. L., Zijlstra, A. A., & Loup, C. 2005, A&A, 438, 273

Appendix A

Additional Figures

Though the discussion in the Analysis section is mainly based on the M31 metallicity results, this Appendix features corresponding plots at different metallicities for the sake of completeness.

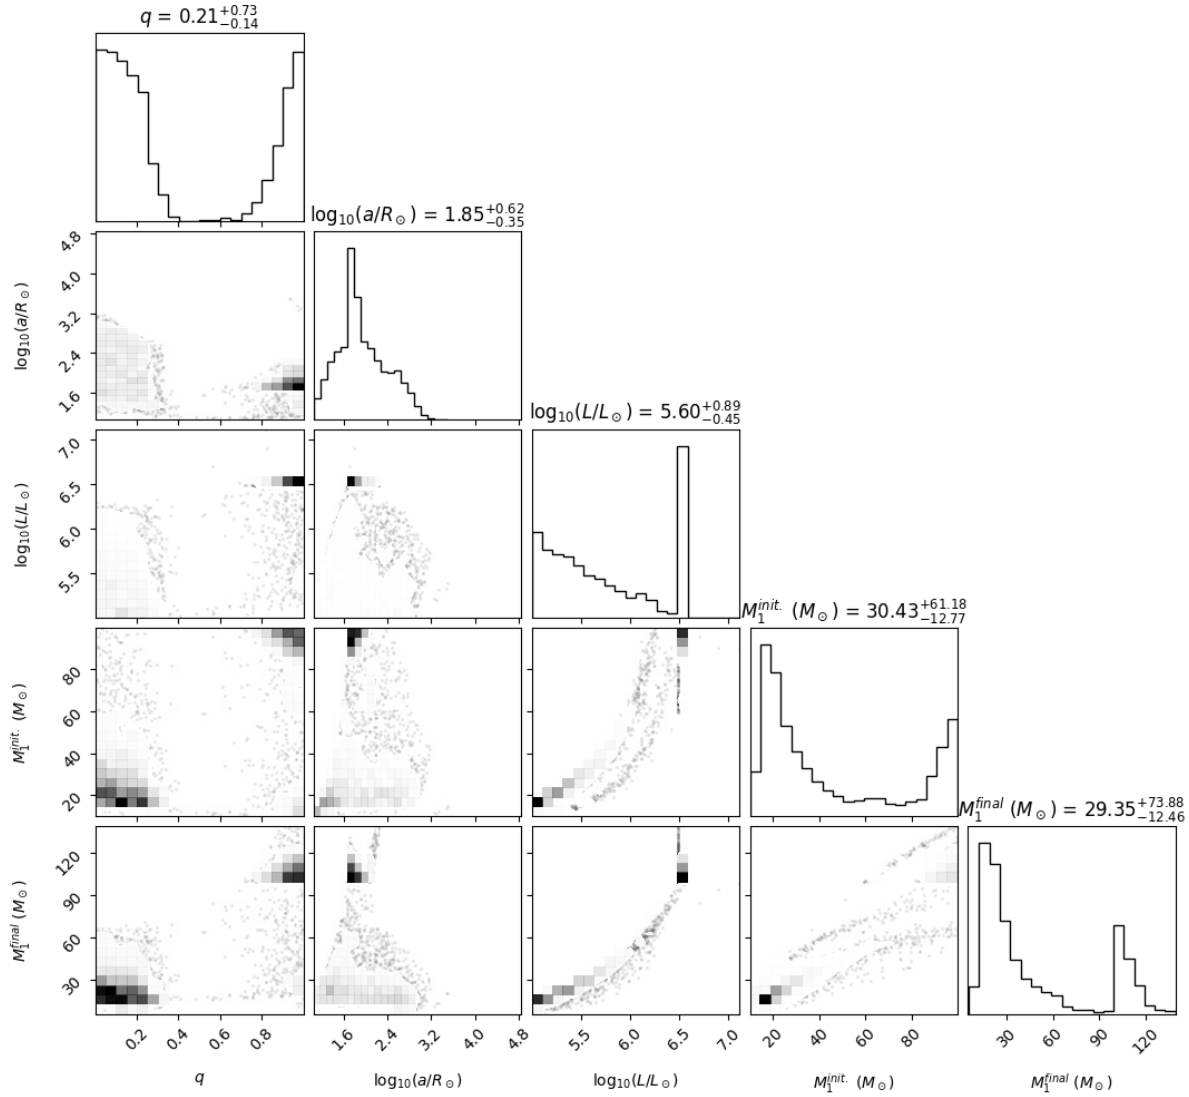

Figure A.1: Stellar and orbital properties of merger systems with $Z_{\rm LMC}$ metallicity. Corresponding M31 result is Figure 3.9.

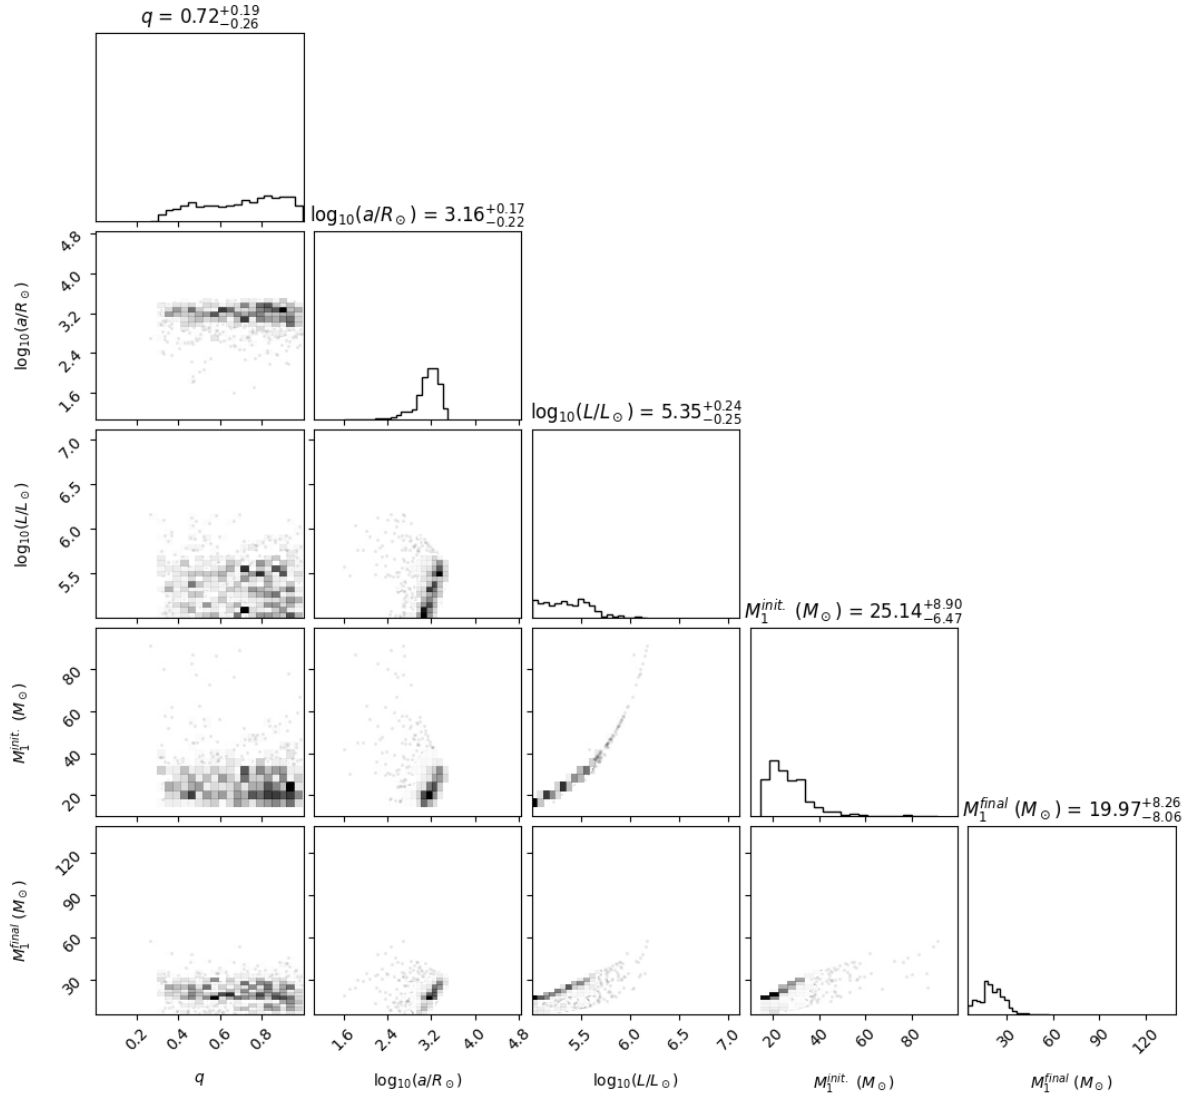

Figure A.2: Stellar and orbital properties of Roche lobe overfilling systems with $Z_{\rm LMC}$ metallicity. Corresponding M31 result is Figure 3.10.

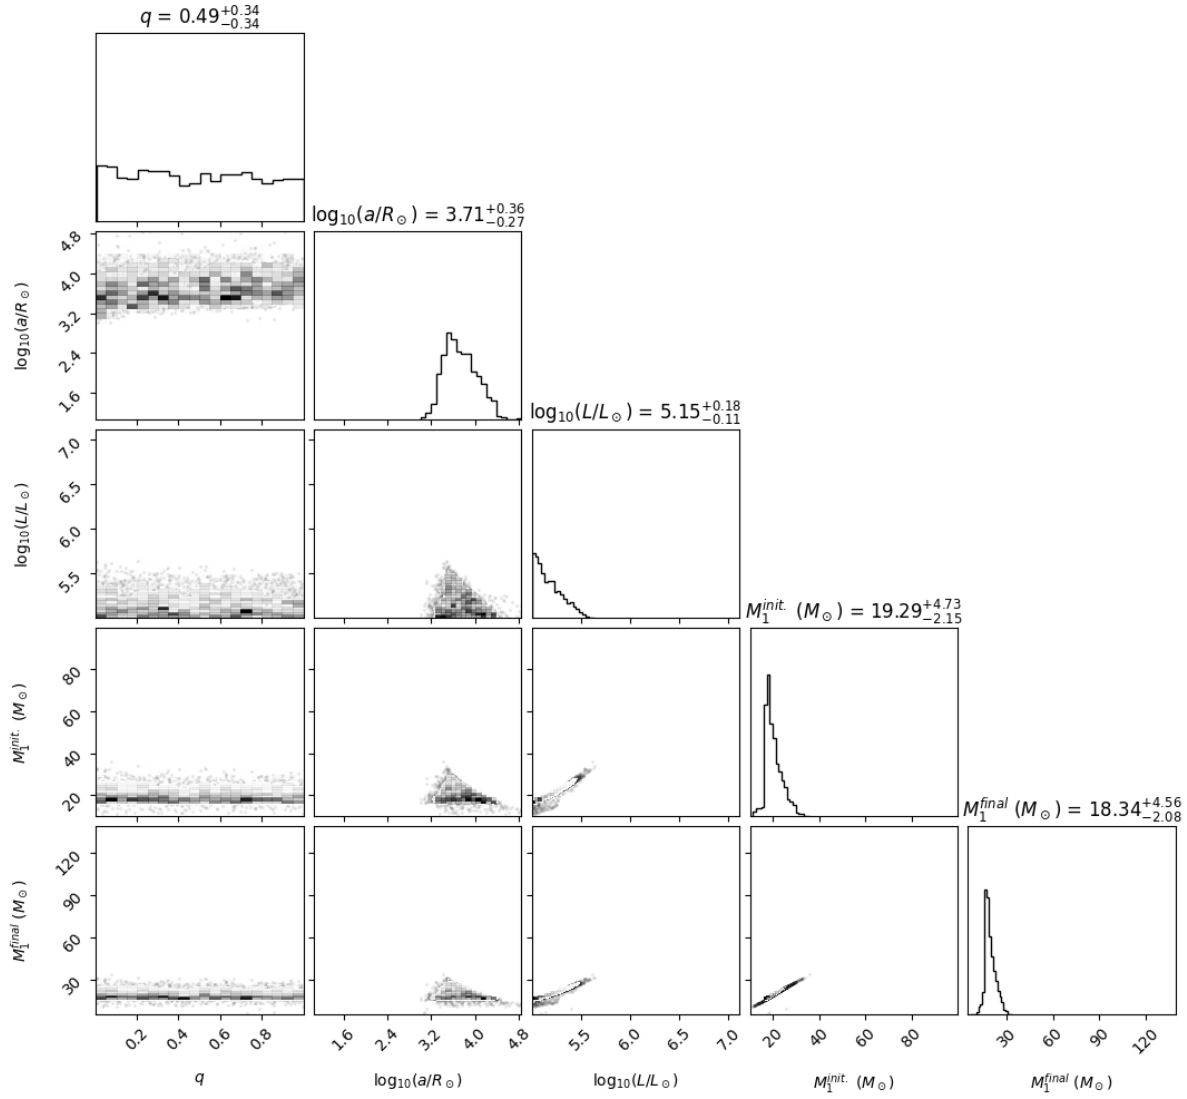

Figure A.3: Stellar and orbital properties of effectively single systems with $Z_{\rm LMC}$ metallicity. Corresponding M31 result is Figure 3.11.

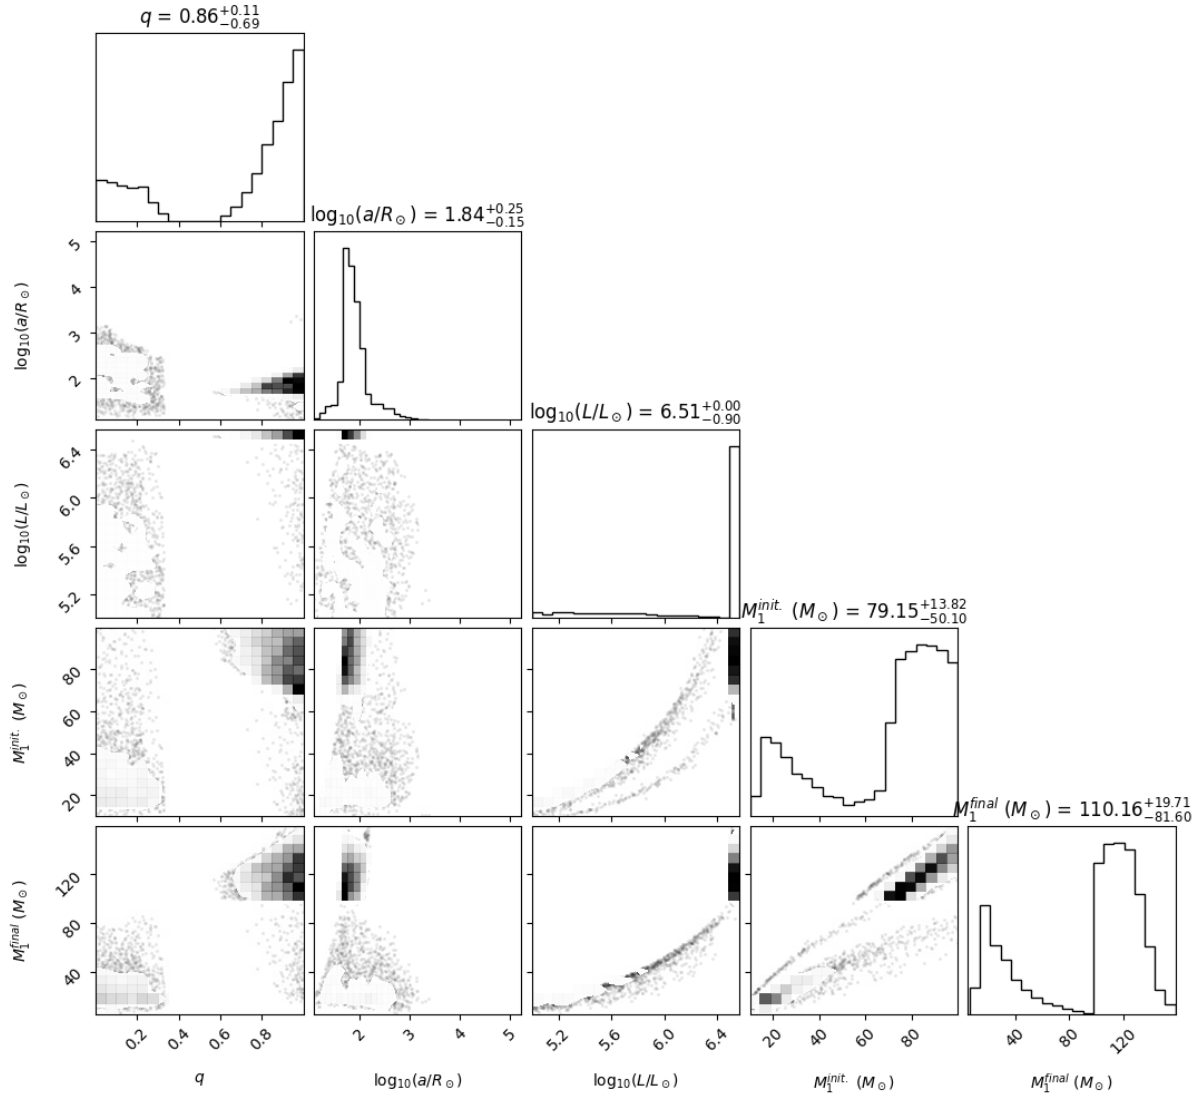

Figure A.4: Stellar and orbital properties of merger systems with $Z_{\rm SMC}$ metallicity. Corresponding M31 result is Figure 3.9.

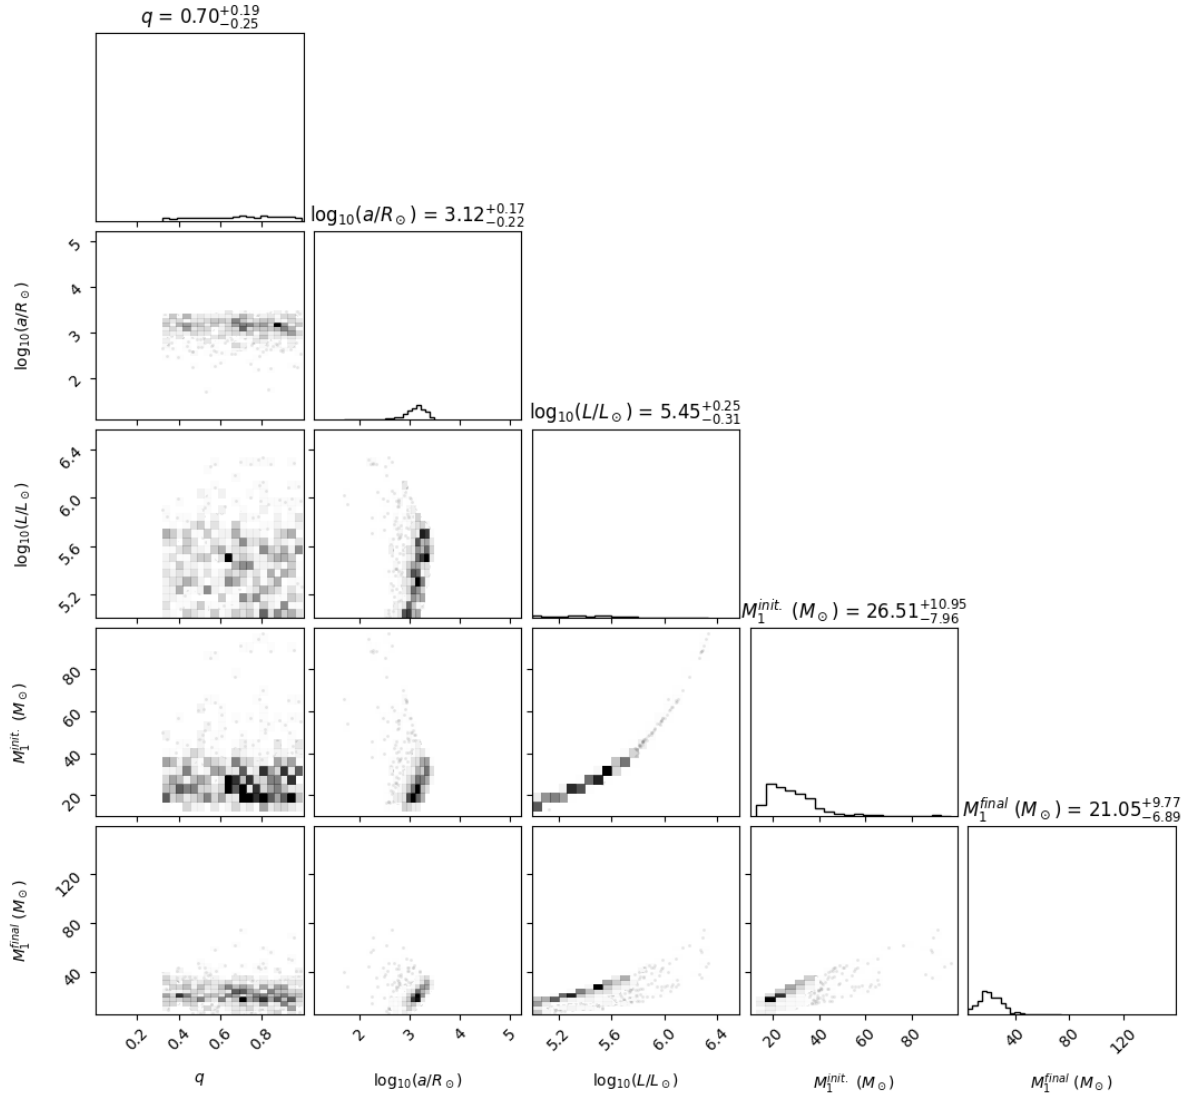

Figure A.5: Stellar and orbital properties of Roche lobe overfilling systems with $Z_{\rm SMC}$ metallicity. Corresponding M31 result is Figure 3.10.

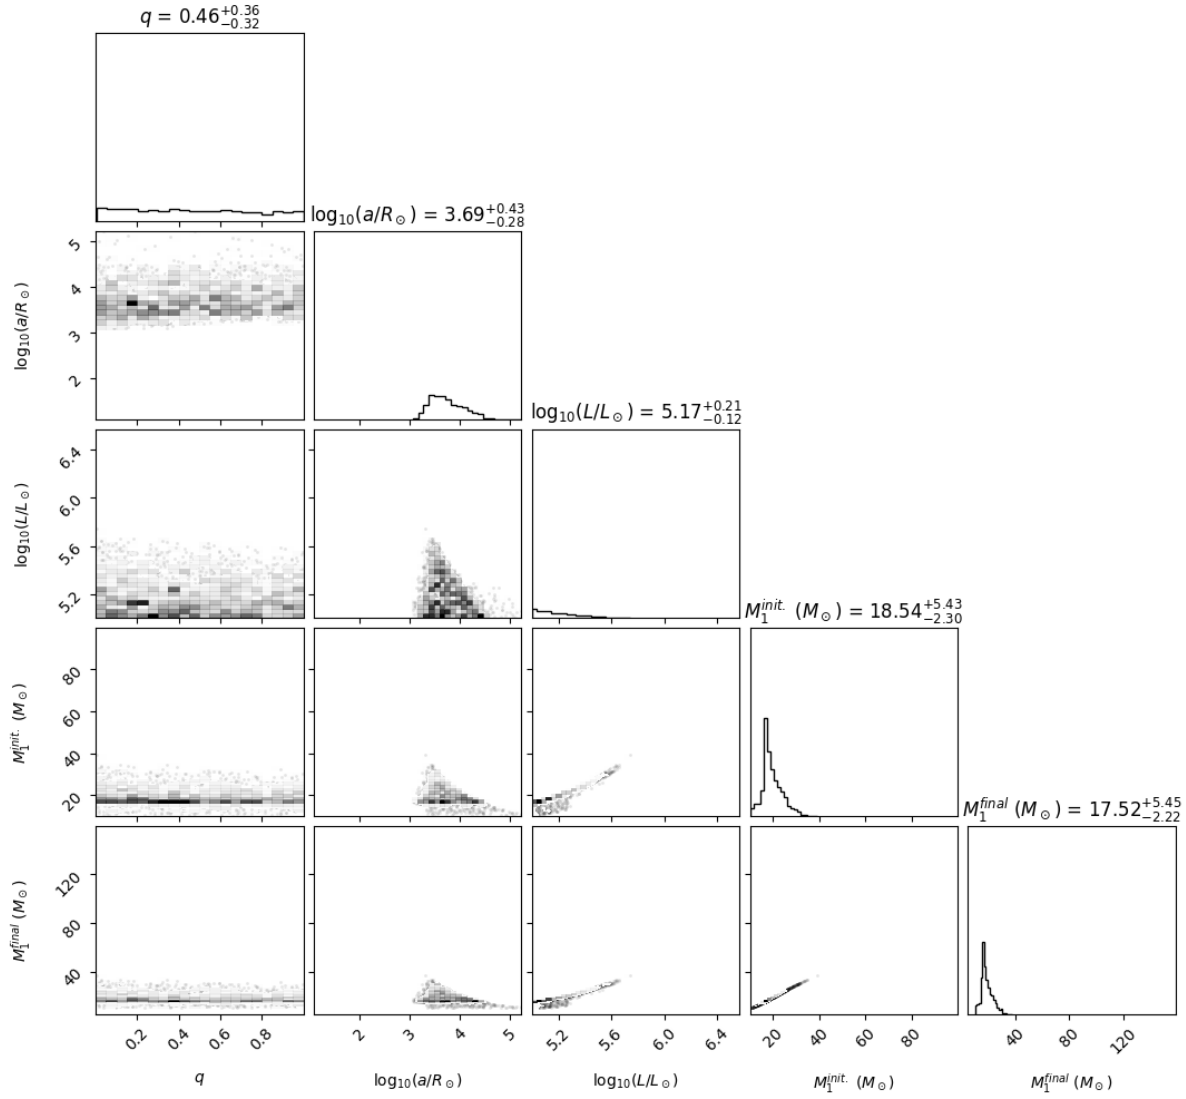

Figure A.6: Stellar and orbital properties of effectively single systems with $Z_{\rm SMC}$ metallicity. Corresponding M31 result is Figure 3.11.