

Dynamical Mass Loss from Unstable Giants



Matthew Clayton

Exeter College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Trinity 2018

Abstract

Giant stars are believed to lose significant fractions of their total mass over their lifetimes, but the mechanisms responsible for this are ill-understood. One possible mechanism is dynamical mass loss – a hydrodynamical process in which matter is ejected from the stellar surface in ballistic outflows.

In this thesis, dynamical mass loss is studied in three stellar regimes: common-envelope objects, asymptotic giant branch stars, and red supergiants. Using hydrodynamical simulations performed with the stellar evolution code MESA, we examine the dynamical behaviour and stability of stars in each of these regimes.

We examine the dynamical properties of common-envelope objects during the slow spiral-in phase using a parameterised 1-dimensional model of orbital dissipatory heating. We find that the envelope becomes unstable to high-amplitude dynamical pulsations that can lead to repeated mass-ejection events capable of removing the entire envelope and terminating the common-envelope phase. We estimate this process's α efficiency value and suggest how these results might be employed in parameterised common-envelope models.

We employ coupled evolutionary and hydrodynamical simulations of AGB stars to study their dynamical properties as they traverse the TP-AGB and examine their dependence on basic stellar properties and on the thermal pulse cycle. We find that these models experience large amounts of dynamical mass loss, and we construct a parameterised model to estimate its strength. We find that this model is successful at locating the termination of the AGB.

We apply a similar approach to a study of RSGs, and find that dynamical mass loss also emerges in this regime. We estimate the conditions under which this occurs and discuss how this mechanism may resolve theoretical problems relating to the Humphreys-Davidson limit and the progenitors of SNe II_n.

We conclude that dynamical mass loss is likely to form a vital part of the mass-loss histories of cool giant stars.

Statement of Originality

I declare that no part of this thesis has been submitted in support of another degree, diploma or other qualification either at the University of Oxford or elsewhere.

This thesis and its contents are entirely my own work except where otherwise stated. The stellar evolution code MESA was provided by the MESA collaboration. Data on the observationally inferred masses of white dwarfs were provided by Dr I. Pelisoli. Fig. 1.1 has been reproduced from Ivanova et al. (2013b).

Acknowledgements

Firstly, I would like to thank my supervisor, Philipp Podsiadlowski, without whom none of this would have been possible. His guidance, mentorship, insight, and tireless enthusiasm throughout my studies have been a vital part of the creation of this work.

I would also like to thank Natasha Ivanova and Stephen Justham for their excellent and dedicated advice and suggestions for improving the CE models and for understanding their implications.

I gratefully acknowledge the financial support of the Science & Technology Facilities Council through grant ST/K502236/6. I am grateful to Ingrid Pelisoli for providing the white dwarf mass data used in this work, and to Lorne Nelson for assisting in the provision of additional computing resources generously supplied by Westgrid and Compute Canada. I would also like to thank the MESA collaboration and user community, in particular Pablo Marchant, Josiah Schwab, and Bill Paxton, for their assistance in overcoming technical issues with the code. I am also thankful to Fabian Schneider and Ryosuke Hirai for helpful discussions about MESA and stellar evolution, and to Jack Setford for either useful or amusing comments on earlier drafts of this document.

The Oxford Astrophysics Subdepartment is a special place because of its people, its friendliness, and its sense of community, and it has been a pleasure to be part of it. Particular thanks in that regard should go to the denizens of office 650 for making it such an enjoyable place to work.

I am of course thankful to my parents for their unwavering support and confidence, and who, as physics teachers themselves, must shoulder some part of the blame for this work. In particular, I can still remember my father explaining to my sister and I how orbits work for the first time, no doubt in response to repeated demands to know why the moon didn't fall down, so any complaints related to the accuracy of the celestial mechanics in this document should be referred to him.

Contents

1	Introduction	1
1.1	Common-envelope objects	2
1.2	Asymptotic giant branch stars	9
1.3	Red supergiant stars	15
1.4	Dynamical mass loss	18
1.4.1	Some definitions	18
1.4.2	Survey of existing work	21
2	The simulation model	35
2.1	The basic problem	35
2.1.1	The equations of stellar structure	35
2.1.2	The equation of state	37
2.1.3	Energy transport	37
2.1.4	Opacity	39
2.1.5	Nuclear reactions	39
2.1.6	Chemical mixing	40
2.1.7	Boundary conditions	41
2.1.8	Mass-loss treatment	42
2.2	The method of solution	43
2.3	Hydrodynamics	45
3	Common-envelope objects	48
3.1	Simulation methodology	49
3.2	Results	54

3.2.1	Hydrostatic simulations	54
3.2.2	Hydrodynamic simulations	58
3.2.3	Nonlinear physics in the pulsation cycle	63
3.2.4	Dynamical ejections	66
3.2.5	Non-ejecting simulations	73
3.2.6	To eject or not to eject	77
3.3	Discussion	82
3.3.1	Numerical considerations	82
3.3.2	Dynamical instability and mass loss	82
3.3.3	Ejection efficiency	83
3.3.4	Limitations of our simulations	86
3.3.5	Observational signatures	87
3.4	Summary	89
4	Asymptotic giant branch stars	91
4.1	Simulation methodology	92
4.2	Results	94
4.2.1	Hydrostatic models on the AGB	94
4.2.2	Hydrodynamic simulations	95
4.2.3	Modelling nonlinear pulsations	108
4.3	Physical analysis	114
4.3.1	The pulsation cycle	114
4.3.2	Dynamical ejection	118
4.3.3	Self-limiting pulsations	119
4.4	The white dwarf mass distribution	125
4.5	Discussion	128
4.5.1	Dynamical instability as a driver of mass loss	128
4.5.2	Uncertainties in our models	130
4.5.3	Turbulent viscous damping	131
4.5.4	Observational signatures	133
4.5.4.1	Circumstellar relics	133

4.5.5	The possibility of direct observation	135
4.6	Summary	136
5	Red supergiants	138
5.1	Simulation methodology	138
5.2	Results	141
5.2.1	The Ya'ari effect	146
5.2.2	Modelling pulsation growth	149
5.2.3	Dynamical mass loss	150
5.3	Discussion	156
5.4	Summary	161
6	Conclusion	163
A	Mass-removal routine	169
	Bibliography	171

List of Figures

1.1	The phases of a common-envelope event	5
1.2	The thermal pulse instability	13
3.1	Hydrostatic evolution under uniform envelope heating	55
3.2	Hydrostatic evolution under base envelope heating	56
3.3	Hydrodynamical evolution under uniform envelope heating	59
3.4	Hydrodynamical evolution under base envelope heating	60
3.5	Some Hertzsprung-Russell diagrams of dynamical pulsations	64
3.6	Pulsation phases and nonlinear physics at large amplitude	67
3.7	Pulsation phases and nonlinear physics of a dynamical mass ejection	68
3.8	The structure of a CE dynamical mass ejection I	70
3.9	The structure of a CE dynamical mass ejection II	71
3.10	The structure of a CE dynamical mass ejection III	72
3.11	The structure of a CE cooling catastrophe I	74
3.12	The structure of a CE cooling catastrophe II	75
3.13	The structure of a CE cooling catastrophe III	76
3.14	A simple model for the arrangement of ejecting and non-ejecting pulsation regimes	78
3.15	A 1,500-year hydrodynamical simulation	79
4.1	The TP-AGB evolution of a $1 M_{\odot}$ star	96
4.2	The TP-AGB evolution of a $1.3 M_{\odot}$ star	97
4.3	The TP-AGB evolution of a $1.6 M_{\odot}$ star	98
4.4	The TP-AGB evolution of a $2 M_{\odot}$ star	99
4.5	The TP-AGB evolution of a $4 M_{\odot}$ star	100

4.6	The TP-AGB evolution of a $6 M_{\odot}$ star	101
4.7	Examples of hydrodynamical evolution regimes	102
4.8	A Hertzsprung-Russell diagram of a dynamical pulsation	103
4.9	A 3,700-year hydrodynamical simulation	109
4.10	Example growth rate fits	110
4.11	Best-fitting models for pulsation growth rate and dynamical mass-loss rate .	112
4.12	An evolutionary calculation using our parameterised model for dynamical mass loss	115
4.13	Pulsation phases and nonlinear physics of a dynamical mass ejection	116
4.14	The structure of an AGB dynamical ejection I	121
4.15	The structure of an AGB dynamical ejection II	122
4.16	The structure of an AGB dynamical ejection III	123
4.17	The white dwarf initial-final mass relation	124
4.18	The white dwarf mass distribution	129
4.19	The effect of turbulent viscous pressure	134
5.1	The starting locations of hydrodynamical simulations of RSGs in the Hertzsprung- Russell diagram	142
5.2	The hydrodynamical evolution of four $25 M_{\odot}$ simulations	143
5.3	Hydrodynamical simulations of RSGs in the L/M vs T_{eff} plane	145
5.4	The Ya'ari effect in radius and entropy	147
5.5	Pulsation growth-rate models	151
5.6	The distribution of simulations undergoing dynamical mass loss	153
5.7	Dynamical mass loss as a function of L/M	154
5.8	The structure of an RSG dynamical ejection I	157
5.9	The structure of an RSG dynamical ejection II	158
5.10	The structure of an RSG dynamical ejection III	159

Chapter 1

Introduction

The theory of stellar evolution has come an impressively long way since the philosopher Anaxagoras first suggested that stars were flaming rocks in the 5th century BC (Arny, 1990). The first mathematically supported stellar models were published by Helmholtz (1854) and Kelvin (1861), with the first true theory of stellar evolution following due to Ritter (18 papers, see Ritter 1898). With subsequent developments like the formulation of radiative transport theory (Eddington, 1926), the determination of hydrogen and helium composition (Payne, 1925), the incorporation of nuclear power sources (Gamow, 1938), the modelling of stellar convection (Cowling, 1934), the construction of giant models (Hayashi, 1949; Oke & Schwarzschild, 1952; Sandage & Schwarzschild, 1952), and the first computer calculations (Heney et al., 1955, 1964), we possess today a well-developed theory of stellar evolution able to reproduce observations well for most classes of stars and to explain their origins and their fates. However, there remain several sources of uncertainty in this theory, one of the largest of which is stellar mass loss.

Mass loss is important in determining many observable stellar properties, most clearly in the giant stages towards the end of stars' lives. For low- and intermediate-mass stars, it can be the primary factor determining the end of the star's life, and hence the distribution of white dwarf remnants and therefore of supernova type Ia progenitor systems, the total amount of visible and infra-red starlight from a stellar population, and the supply of enriched material to the interstellar medium; for more massive stars it can determine the relative rates of different type II supernovae; and for common-envelope systems it is a matter of life and death for a binary system, determining the merger rate and the period of the surviving binaries, and thereby controlling the formation rate of many of the most interesting interacting binary

systems.

Despite its importance, the regime of cool giant stars remains one of the areas in which our understanding of stellar mass loss is the most uncertain. It is stars in this phase which tend to have the highest luminosities and radii (and therefore that produce a large fraction of the total supply of observed starlight), and it is also these stars which are believed to undergo the highest (and least constrained) rates of mass loss, although the mechanisms by which this mass loss occurs remain highly uncertain (see the review by Willson 2000).

In general, mass loss from stars can be divided into two classes of mechanisms: wind mechanisms and dynamical mechanisms. Wind mechanisms remove matter from a star in (more or less) steady outflows driven by the radiation pressure exerted by a star on material in its atmosphere, or by Alfvén or acoustic waves, whilst the star as a whole remains in or near a state of hydrostatic equilibrium. Conversely, dynamical mass loss produces inertial, semi-ballistic outflows which are launched in violent ejection events, during which the star departs from hydrostatic equilibrium and evolves dynamically. In this thesis we will examine the possibility that dynamical mass ejection mechanisms play a significant role in the mass-loss histories of cool giants from three distinct stellar classes – common-envelope (CE) objects, asymptotic giant branch stars (AGBs), and red supergiants (RSGs) – and discuss whether a unifying model of dynamical mass loss can be constructed across all of these stars. In this chapter we will discuss the established evolutionary and structural characteristics of each of these stellar classes in turn, summarise our current understanding of their mass-loss properties, and then review the existing literature on dynamical mass loss as it relates to these classes of objects. We will then discuss the simulation model upon which the work reported in this thesis is based in Chapter 2, and then in Chapters 3, 4, and 5, we will apply this model to each stellar class under consideration in turn, before making some summarising remarks in Chapter 6.

1.1 Common-envelope objects

The majority of stellar remnants have evolved from stars several orders of magnitude larger than themselves. There are many observed classes of binaries containing such remnants whose progenitors would have had radii significantly greater than the current

orbital separation of the binary. The existence of these systems can be explained if they evolved from much wider binaries which have lost most of their angular momentum during their lifetimes. This is accomplished during a short phase of binary evolution known as the common-envelope phase.

During a CE event, the binary orbit is completely engulfed within an extended shared envelope of stellar matter (Paczynski, 1976). A CE configuration can arise through several processes, for example a giant star expanding and overflowing its Roche lobe will transfer mass to its companion; if this mass transfer occurs at a high enough rate that the companion cannot accrete all of this material (as is expected in cases of dynamical mass transfer when mass loss from the primary causes its radius to *increase* rather than decrease), then material will build up in the companion's Roche lobe also, eventually overflowing this too and entirely flooding the orbit with material. Another possible formation scenario is due to an orbital instability in the binary such as the Darwin instability (Darwin, 1879), which can bring initially separated binaries into contact and lead to runaway mass transfer. The CE configuration is in any case a binary pair, usually made up of the core of the giant and its companion (whether a compact object or a main-sequence star), orbiting within a shared envelope. This arrangement gives rise to drag forces which act on the binary components, transferring energy and angular momentum to the envelope and causing the binary to shrink. This clearly limits the lifetime of the CE phase: if the binary is to survive, this frictional dissipation must be brought to an end by the removal of the envelope. This may be accomplished if the amount of energy transferred is large enough to unbind the envelope from the system and eject it to infinity. The binary would in such a case survive with a shortened period. Conversely, if the envelope is not ejected, the binary will continue to shrink until the two stars merge.

In either case, the CE process causes dramatic changes in the structure of a binary system and is thought to be a major formation channel for many classes of late-stage stellar systems which are of great interest to astrophysicists; examples include X-ray binaries, cataclysmic variables, ms pulsars, SN Ia progenitors (in both single and double degenerate models), stellar-mass gravitational-wave sources, Thorne-Żytkow objects, and potential progenitors of both short- and long-duration gamma-ray bursts (see Ivanova et al. 2013b for an overview of the importance of CE events in forming stellar exotica). But CE events are by no means

only a feature of unusual stellar systems; rather, they are believed to be extremely common – it has been estimated that approximately 30 % of all binaries undergo at least one such event during their lifetimes (Han et al., 1995). Some systems, including some possible SN Ia progenitor formation channels, are expected to go through multiple CE events, as both binary components pass through their giant stages in turn.

The importance of common envelopes in the evolution of binary systems clearly makes a robust understanding of the process highly desirable, and a large amount of work has been done in this area. However, the CE process has proven remarkably difficult to constrain either theoretically or observationally. There are no clearly identified examples of ongoing CE events that have been observed, although there are a number of candidate systems and an associated class of transients which are believed may represent CEs (see below). For this reason, the majority of our knowledge about CE evolution comes from theoretical considerations and simulations.

Initial attempts to model the CE phase were performed using 1-dimensional simulation codes, with influential early studies reported by Taam et al. (1978), Meyer & Meyer-Hofmeister (1979), and Delgado (1980). Using this kind of 1-dimensional methods, three phases have been identified in the evolution of CE simulations (Podsiadlowski, 2001; Ivanova, 2002):

- **I: Loss of corotation** – Before the onset of the CE process, we expect, for most mass ratios, a giant in a close binary system to be in a circularised orbit and tidally locked in corotation with its partner (Portegies Zwart & Meinen, 1993; Soker, 1996). As the CE forms, however, this corotation will break down, and the orbit will cease to be circular. This phase may occur over a long timescale (depending on the nature of the onset of CE evolution), with the expected initial corotation of the system leading to low rates of energy dissipation.
- **II: Fast plunge-in** – As corotation breaks down and the binary begins to shrink significantly, the system enters the rapid plunge-in phase, in which the rate of dissipation is high and the binary orbit contracts on a dynamical or orbital timescale. The path of the binary components in this phase is not well described as a closed orbit, but is rather a pair of inward spirals. This dynamical process may lead to a prompt ejection

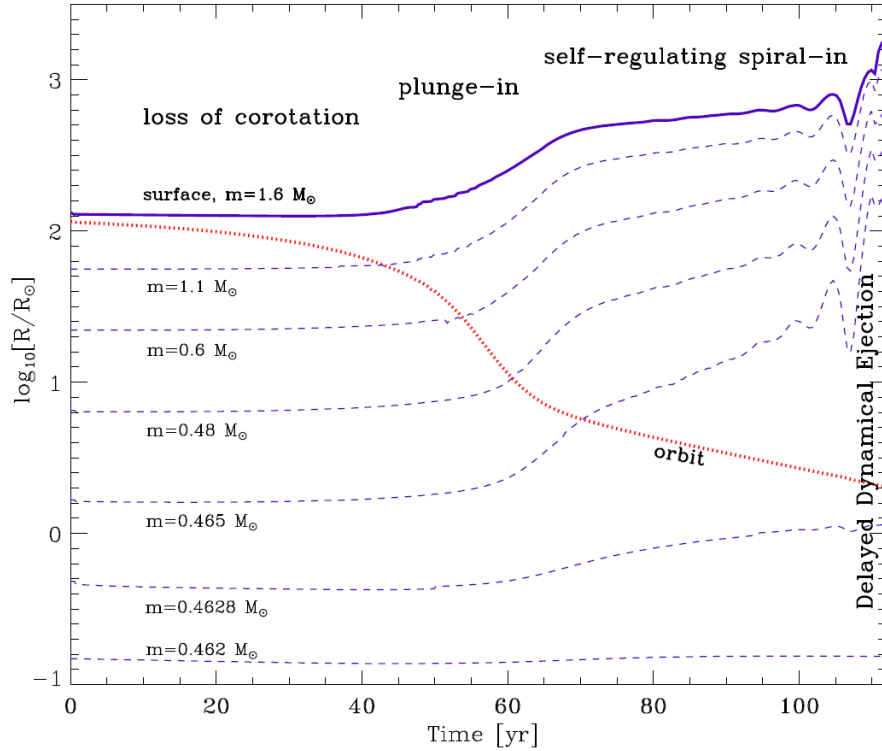


Figure 1.1: The common phases of CE evolution as they appear in a simulated CE event involving a $1.6 M_{\odot}$ giant and a $0.3 M_{\odot}$ white dwarf. The dashed lines represent Lagrangian mass coordinates within the giant, and the red dotted line shows the radius of the secondary's orbit. This figure is reproduced from Ivanova et al. (2013b) and is based on simulation data from Ivanova (2002). Republished with permission of Springer Verlag Berlin/Heidelberg, from *Common envelope evolution: where we stand and how we can move forward*, N. Ivanova, S. Justham, X. Chen, O. De Marco, C. L. Fryer, E. Gaburov, H. Ge, E. Glebbeek, Z. Han, X.-D. Li, G. Lu, T. Marsh, P. Podsiadlowski, A. Potter, N. Soker, R. Taam, T. M. Tauris, E. P. J. van den Heuvel, R. F. Webbink, 21:59, 2013; permission conveyed through Copyright Clearance Center, Inc.

of the envelope, but if this does not occur, the expansion of the envelope caused by the strong deposition of orbital energy as heat reduces its density and leads to a reduction in the rate of orbit contraction.

- **III: Self-regulating spiral-in** – This phase (sometimes referred to as the slow spiral-in) returns to a longer timescale, as frictional dissipation in the envelope, rather than dynamical effects, determines the rate of energy loss from the binary. The expansion of the envelope driven by heat deposition reduces the drag force on the binary and represents a negative feedback process controlling the rate of orbit contraction. For ejection to occur in this phase, the mechanism must be dependent not on the dynamical effects of the spiral in, but rather on the thermodynamic and hydrodynamic structure of the envelope and its thermal properties. However, the exact nature of such delayed ejections remains ill-understood. Not all CE events feature this third phase, as the envelope may have been ejected during the plunge-in, or the binary may merge promptly. For certain initial configurations containing compact secondaries, the secondary can enter more dense layers below the convective envelope, leading to an increase in dissipation, and it is possible that the binary will enter the plunge-in phase for a second time which may lead to a delayed dynamical ejection event.

For more information on our current understanding of CE events, see the review by Ivanova et al. (2013b) as well as (Izzard et al., 2012; Taam & Sandquist, 2000; Iben & Livio, 1993).

The variation in scales during the three phases of the CE process is the main reason that it is difficult to model. The timescales which it is necessary to resolve range from the orbital period, likely of order hours or days, and the dynamical timescale of the interaction between the secondary and the surrounding fluid (which may be of order seconds if it is a compact object), to the thermal timescale of the envelope and possible duration of the event, which may exceed 1000 years. The spatial scales range from the size of the secondary, which may only be 10 km for a neutron star, to the radius of the shared envelope which will likely be of order $1000 R_{\odot}$. For these reasons, despite the exponential growth in available computing power since the CE process was first proposed, we remain unable to model the entirety of the CE event in a physically self-consistent way, especially during the slow phases.

Although several groups of authors have performed 3-dimensional hydrodynamical simulations of CE events (recent examples include Ricker & Taam, 2012; Passy et al., 2012; Nandez et al., 2015; Ohlmann et al., 2016; Staff et al., 2016; Nandez & Ivanova, 2016; Iaconi et al., 2017), the limitations on the physics accessible to 3-d codes, and their high computational costs, have largely restricted these simulations to the fast, dynamical plunge-in phase of CE events, in which the system evolves almost adiabatically and the evolutionary timescales are short. For the foreseeable future, 3-d codes are not likely to be able to simulate the self-regulated spiral-in phase, as the simultaneous lengthening of the evolution timescale and the shortening of the binary orbit renders the timescales unfeasible to simulate, although 3-d codes have begun to be able to follow the transitional period into the beginning of this phase (Ivanova & Nandez, 2016). In addition to timescale constraints, the slow spiral-in phase is sufficiently long that thermal flows through the envelope must be taken into account, as well as the nuclear energy generation of the giant’s core. For these reasons, 1-d simulations remain our most effective tools for studying the slow phases of the CE event, as they are able to include more of the relevant stellar physics whilst also being able to cover the full range of timescales involved in the CE process.

Much of the theoretical work that has been done on CE evolution has centred around attempts to construct a simple formalism for predicting the outcomes of CE events from energy considerations. In the simplest case, we can equate the energy used to expel the envelope to the energy lost from the binary orbit, modified by some efficiency factor. This energy formalism, often referred to as the α formalism after the efficiency parameter introduced by Iben & Tutukov (1984), Webbink (1984), and Livio & Soker (1988), can be formulated for example as

$$\alpha\Delta E_{\text{orb}} = E_{\text{bind}} = \frac{GM_{\star}M_{\text{env}}}{\lambda R_{\star}}, \quad (1.1)$$

where α is a parameter of order 1 encoding the efficiency of the CE ejection process, ΔE_{orb} is the change in orbital energy in the binary over the course of the CE event, E_{bind} is the binding energy of the giant envelope at the onset of the CE phase, G is Newton’s gravitational constant, M_{\star} is the mass of the giant, M_{env} is the mass of the envelope, R_{\star} is the radius of the giant and λ is some parameter which encodes the structure of the envelope over which the binding energy is integrated.

This formalism is often used by binary population synthesis (BPS) codes to parameterise the entire CE process, using some chosen value or range of values of the product $\alpha\lambda$ (see De Marco & Izzard 2017 for an inventory of BPS codes and their capabilities). Unfortunately, this formalism has a number of weaknesses:

- It relies on having accurate expressions for all the major sources and sinks of energy present in a CE system, which are not yet fully established, and may not be possible to describe in a simple manner.
- It applies only to prompt dynamical ejection of the envelope during the plunge-in phase, because if the system undergoes a self-regulating spiral-in the lifetime of the CE event will not be short compared to the thermal timescale of the envelope and it cannot be treated under the assumption of adiabaticity implicit in the α formalism.
- It assumes that the kinetic and thermal energy of ejected material is zero at infinity, whereas simulations suggest that as much as half of the released orbital energy may persist in the ejecta at infinity (Nandez & Ivanova, 2016).
- It requires some clear definition of the base of the envelope – what is envelope and must be ejected and what counts as part of the core and can be retained in a post-CE binary. There is no clear point to use for this, and the value of the binding energy can depend sensitively on where one chooses to integrate from (Ivanova et al., 2015).
- It assumes that there is some universal value of α which is the same across different classes of giants and secondaries. There is no strong reason to believe that this is the case, and attempts to calibrate α have proven very difficult (see Ivanova et al., 2013b, for discussion).

A significant source of uncertainty remains the identification of available energy sources. In particular, there is strong evidence that an important part of the energy budget can be provided by the recombination energy stored in ionized material, which can be accessed by an expanding envelope to help unbind it from its star (see, for example Han et al., 1995, 2002, 2003; Ivanova et al., 2015). The addition of recombination energy release to a 3-d hydrodynamical model recently resulted in the first simulation to show a complete envelope

ejection (Nandez et al., 2015). Nonetheless, the details of how such liberated recombination energy may affect the ejection dynamics remain unclear (see discussion in Ivanova et al. 2013b and Ivanova 2017). Importantly, Ivanova et al. (2013a) argued that the properties of a class of optical transients would be well explained if a large fraction of the hydrogen recombination energy in the ejected material was often radiated away rather than used to help the ejection; however the fraction of CE ejections to which that conclusion applies is uncertain, and their argument does not restrict the usefulness of the helium recombination energy.

The most likely observational candidates for ongoing CE events seem to be the class of transients known as luminous red novae (LRNe) (see Ivanova, 2017). These bright, cool transients have been associated with dynamical events during the CE process, including the prompt ejection of the envelope, but are also expected to be seen in the case of a merger. In a few cases these transients have been fitted with CE ejection or merger models (Tylenda et al., 2011; Ivanova et al., 2013a; MacLeod et al., 2017).

In this work, we will examine the possibility of dynamical mass ejections occurring during the self-regulating spiral-in phase. In this phase, the binary orbit does not evolve quickly, and so we cannot expect it to drive dynamical processes in the envelope directly. Considering this fact, it is highly uncertain how the envelope can be removed in this phase. In Chapter 3, we will consider the possibility that ejection of the envelope can occur in this phase due to instabilities inherent in the envelope itself and examine whether these instabilities can give rise to sufficient dynamical mass loss to clear the envelope and end the CE phase.

1.2 Asymptotic giant branch stars

Asymptotic giant branch stars are low- and intermediate-mass stars in the fourth major stage of their evolution. AGB stars are giants which have passed from the hydrogen-core-burning main-sequence, along the hydrogen-shell-burning Red Giant (RG) branch, through the helium-core-burning red clump or horizontal branch phase, and entered a phase in which helium is burned in a shell around an inert carbon/oxygen core, and hydrogen is burned in a second shell around the helium shell. For this reason, they are sometimes known as double

shell stars or second ascent giants¹. The AGB stage is the final stellar phase in the life of stars that are massive enough to achieve the core temperatures required to ignite helium burning ($\sim 10^8$ K) but not massive enough to ignite carbon burning ($\sim 5 \times 10^8$ K), corresponding to masses between approximately 0.47 and 8 solar masses. As a star progresses along the AGB, its inert carbon/oxygen core grows and its total mass shrinks due to strong surface mass loss. Eventually, the entire envelope is lost and the core cools to become a white dwarf (WD), possibly producing a visible planetary nebula (PN). It is believed that it is the surface mass loss which is of primary importance in terminating the AGB phase, although the mechanisms by which mass is lost from these stars, particularly towards the end of the AGB phase, are not well understood. For a full summary of the properties and evolution of stars on the AGB, see Kippenhahn et al. (2012), Habing & Olofsson (2003), and reviews by Iben & Renzini (1983), and Herwig (2005).

Our limited knowledge of the mass-loss mechanisms and histories of AGB stars is one of the largest sources of uncertainty in low- and intermediate-mass stellar evolution theory. For example, stars in the AGB phase are believed to be the biggest source of IR emission in many galaxies (Villaume et al., 2015), among the most important producers of interstellar dust (see, for example, Tielens, 2005; Matsuura et al., 2009), the source of much of the heavy element enrichment of the interstellar medium (Karakas & Lattanzio, 2014), and of course the progenitors of WDs and PNe. The length of time a star spends in the AGB phase is therefore information of vital importance when performing calculations in all these areas, and because the lifetime of the AGB is primarily determined by how long it takes AGB stars to lose their envelopes via surface mass loss, rather than by nuclear evolution, it is highly desirable that we understand the mechanisms through which this mass loss occurs.

Evidence that AGB stars undergo rapid mass loss is extensive, with estimated rates reported which range between 10^{-8} and above $10^{-3} M_{\odot} \text{ yr}^{-1}$ (Höfner & Olofsson, 2018), and includes direct observational evidence such as studies of circumstellar material (see,

¹For the avoidance of confusion, we note that the term “red giant” can be ambiguous, as some authors use it to refer only to first ascent giants, that is, giants with helium cores and only one burning shell, whereas others also include AGB stars in this definition. The inclusion of AGB stars under the red giant label is reasonable when authors discuss their envelopes as in that context both classes of stars are similar, in contrast to evolutionary and chemical contexts in which their differences are acute and the term can be confusing. We shall employ Red Giant with capitals, abbreviated RG, to refer to first-ascent giants, and in cases where we wish to describe both RGs and AGBs, we shall use the lower case red giant, which we shall not abbreviate.

for example, Deutsch, 1956; Ramstedt et al., 2009; Riebel et al., 2012), and also theoretical evolutionary arguments, for example: the masses of AGB cores at the end of nuclear burning determine the WD mass distribution, which is seen to be strongly peaked within a small mass range in the vicinity of $0.6 M_{\odot}$ (e.g. Koester et al., 1979; Limoges et al., 2015; Kepler et al., 2016); that this distribution differs strongly from the predicted mass distribution of stars entering the AGB phase necessitates significant mass loss during the AGB. In a similar vein, clusters such as the Hyades are seen in which main-sequence stars significantly out-mass the WDs thought to result from coeval stars of greater mass (Auer & Woolf, 1965), requiring strong mass loss during the stars' lifetimes. Furthermore, it has long been believed that PNe are made up of material lost from late-stage giants in the final periods of their lives, and therefore represent visible records of their mass-loss histories (Shklovsky, 1956).

There has been extensive work studying the mechanisms by which AGB mass loss occurs (for recent reviews, see Willson 2000, Höfner & Olofsson 2018), with interest in both wind mechanisms and dynamical mechanisms. Wind loss mechanisms have received the most attention recently, with much interest in dust-driven winds specifically (e.g. Winters et al., 2000; Jeong et al., 2003; Höfner et al., 2003; Eriksson et al., 2014; Liljegren et al., 2016) which depend to differing extents on the pulsations of the star to provide energy (e.g. Mattsson, 2016; Miller Bertolami, 2016) or to move material into dust-forming regions. We shall discuss the dust-driven wind model in more detail below. There has also been considerable interest in the possibility that dynamical mass loss may be important on the AGB, an idea which has a long history in this regime. This work is summarised in detail in Section 1.4.

Pulsations are an important feature of AGB envelopes, and a complex series of pulsation sequences has been identified in observed stars (Wood, 2015), covering periods in the range of 10-1000 days. As AGB stars evolve, they become larger and more luminous, and their pulsation properties change, with more evolved stars becoming Miras, long-period variables that exhibit large-amplitude fundamental mode pulsations. Mira pulsations are believed to be driven by the κ mechanism operating in the extended partial-ionisation zones of these stars, although this remains uncertain, and seems to be more of a default assumption rather than an established theory (dating back to at least Langer 1971).

This evolution is complicated by the presence of thermal pulses: strong cyclical variations in nuclear luminosity caused by a secular instability in the dual burning shells, which develops once stars have progressed sufficiently far along the AGB and repeats on timescales of thousands to tens of thousands of years, depending on stellar mass. The thermal pulse cycle consists of a long ($\sim 10,000$ years) hydrogen burning phase, in which the star’s helium burning shell is inactive, and a short (~ 100 years), extremely luminous helium burning phase (often called the helium shell flash), during which the heat generated by helium burning expands the helium layer outward, suppressing hydrogen burning at its outer edge due to temperature reduction. After the expansion of the helium layer has reduced its temperature enough to restore stability, helium burning ceases, leading to an extended trough in luminosity, and hydrogen burning steadily resumes as the region surrounding the core cools and contracts, causing luminosity to gradually recover. The intermittent nature of both burning shells can be seen in Fig. 1.2. For the avoidance of doubt, we will use the term “pulsation” to refer only to instabilities arising in the envelope, and “pulse” to refer only to thermal pulses. Stars whose thermal pulses are fully developed are known as TP-AGB stars.

Although we will not discuss extensively the prevailing wind models in modern AGB theory (see Höfner & Olofsson 2018 for such a discussion), it will be informative to summarise the current “standard” wind model. The present leading model for AGB envelope loss is the pulsation-enhanced dust-driven wind (sometimes called PEDDRO)². In the PEDDRO model, AGB stars experience mass loss on scales of $10^{-7} - 10^{-5} M_{\odot} \text{ yr}^{-1}$ in the form of a wind driven by radiation acting on dust and accelerating material to speeds of a few tens of kilometres per second. A recent review of this model is given in Höfner (2015).

Unlike in hot luminous stars, where opacity is believed to be generated by atomic transition lines (see, for example Owocki, 2010; Smith, 2014), there is insufficient opacity generated by the gas in the atmospheres of AGB stars to drive significant winds, so the PEDDRO model invokes dust formation in the atmosphere to provide this opacity. However, dust is only able to form at temperatures and densities located high up in the atmosphere

²Confusingly, there is also a pulsation-driven, dust-enhanced wind which appears in the literature (Mattsson, 2016; Miller Bertolami, 2016; McDonald et al., 2016). The distinction between dust driving and pulsation driving lies in the relative amount of energy supplied to the acceleration of outflows from either radiation pressure acting on dust or hydrodynamical acceleration due to pulsations, with the more significant contributor said to be driving the wind.




Figure 1.2: A Kippenhahn diagram showing the 6th and 7th thermal pulses of a $1.6 M_{\odot}$ stellar model. Areas of strong burning, where the nuclear energy release rate exceeds 10^5 , 10^6 or 10^7 ergs $\text{g}^{-1} \text{s}^{-1}$ are shown in increasingly darker blue, while convective regions are marked by green hatching. Also shown are the approximate boundaries of the helium and carbon cores in dashed red and dashed green. Model number is used as a proxy for time (see Chapter 2) to enable both hydrogen and helium burning phases to be seen clearly. This figure was created from a simulation performed by the author using the stellar evolution code MESA (see Chapter 2).

of AGB stars, and material must be elevated to this region before dust can form. This elevation is provided, in this model, by pulsations of the AGB envelope. Steady, sinusoidal pulsations of the envelope below the photosphere are generally assumed, which then steepen into shocks as they propagate outwards. These shocks act to “levitate” material – successive outward-propagating shocks lift it upwards until it reaches high enough altitudes that dust can form. This dust is then accelerated outwards by the star’s radiation field due to its high opacity, and it can carry gas (which still constitutes the majority of the mass) along with it in an outflow due to collisional coupling.

Observational estimates of the mass-loss rates from AGB stars can be obtained through studying the material in the outward-flowing circumstellar envelope. Estimates of the density in this region, the outflow velocity, and the mass-loss rate can be obtained with the assistance of radiative transfer codes from molecular emission lines, most often CO (e.g. Schöier & Olofsson, 2001; Ramstedt et al., 2009), from OH maser emission (e.g. Baud & Habing, 1983; Engels & Bunzel, 2015), or from continuum absorption by dust grains (e.g. Srinivasan et al., 2009; Groenewegen et al., 1998).

Importantly, there is evidence that AGB mass loss exhibits strong variation on timescales of a few hundreds of years. A small number of AGB stars have been seen to have distinct geometrically thin shells of material surrounding them (see Höfner & Olofsson 2018 for a recent summary of these findings), indicating that short (hundred-year) increases in mass loss by multiple orders of magnitude have occurred in these stars’ recent past. It has been claimed that variability of this form is a common feature of stars at the end of the AGB phase (Olofsson et al., 1990). It has also been argued for some time that very large mass-loss rates (in excess of $10^{-5} M_{\odot} \text{ yr}^{-1}$) are required at the end of the AGB in what is usually termed a “superwind” phase (Renzini & Voli, 1981; Vassiliadis & Wood, 1993), with extreme mass loss leading to heavy dust-shrouding, which causes stars to appear as OH/IR stars, so named because their emission is dominated by OH masers in their circumstellar envelopes, while the stars themselves are heavily extinguished (see Justtanont et al., 1996, 2013). The exact nature of this superwind phase and the mechanisms driving it remain under debate (see, for example Engels & Jiménez-Esteban, 2007; Justtanont et al., 2013; de Vries et al., 2014). There are also features seen in PNe that appear to represent episodes of strong mass loss separated by periods of a few hundred years in the final stages of the AGB; we discuss these

features in greater detail in Chapter 4, but it is important to note that this timescale is both considerably longer than the known pulsation timescale (100-1000 days), and considerably shorter than the thermal pulse timescale (a few $\times 10,000$ years).

In this thesis, we will examine the possibility that dynamical mass loss forms an important part of the mass-loss histories of AGB stars, and consider whether this mechanism can help explain the strongly episodic mass loss seen in some of these stars and the extreme mass-loss rates that are theoretically required at the ends of their lives. There has been considerable interest in this field in the past, and we will examine the literature on dynamical mass loss in stars of this kind in detail and discuss where the present work fits in to this tradition in Section 1.4.2.

1.3 Red supergiant stars

Red supergiants are evolved massive stars with initial masses of between approximately 8 and $40 M_{\odot}$ which are undergoing advanced nuclear burning and have developed giant convective envelopes, causing them to migrate to the red region of the Hertzsprung-Russell (HR) diagram (we will discuss RSGs as an evolutionary class rather than as an observational one – for example certain AGB stars are observationally classed as supergiants but we exclude them from this discussion)³. These stars have been identified as the progenitors of type II-P SNe (Smartt et al., 2004, 2009) and are expected to be the final stage of stellar evolution for moderately massive stars, while more massive stars are expected to lose their entire hydrogen envelopes and end their lives as Wolf-Rayet stars or yellow supergiants, either after a period as an RSG or directly from the MS, eliding the RSG stage completely (see e.g. Levesque, 2017). The upper mass limits for RSGs, both the limit for stars experiencing core collapse as RSGs and the limit for stars appearing as RSGs at all, remain uncertain, mostly because of uncertainties in treatments of the strong mass loss appearing in these stars. Observational estimates of the upper limit on the initial mass of RSG progenitors for SNe II-P ($\sim 17 M_{\odot}$) have for some time been substantially lower than the evolutionary limit ($25\text{--}30 M_{\odot}$) (see Smartt et al., 2009; Smartt, 2015; Davies, 2017), which has been variously

³There is another evolutionary RSG regime, occupied by very massive metal-free stars, which are believed to evolve to the red before undergoing pair-instability supernovae, and it is even possible that these stars undergo dynamical pulsations (Moriya & Langer, 2015), but we will not discuss them here.

ascribed to underestimated RSG mass loss (Georgy, 2012), the formation of black holes without bright SNe (Smartt et al., 2009), or systematic underestimation of RSG masses (Davies & Beasor, 2018). A better understanding of mass loss on the RSG would clearly be helpful in resolving this debate.

The mass-loss rates of RSGs are particularly poorly constrained, even by the standards of cool giants. No quantitative physical model of their winds or the principal driving mechanism exists (see Bennett 2010 for a discussion), and although empirical estimates can be made by measuring extinction by circumstellar dust (Jura & Kleinmann, 1990), and in a few cases gas densities (e.g. Bowers & Knapp, 1987; Keene et al., 1993), recent estimates such as those by Maeron & Josselin (2011) and van Loon et al. (2005) show a scatter of nearly 3 orders of magnitude between 10^{-7} and a few times $10^{-4} M_{\odot} \text{ yr}^{-1}$.

RSGs are known to exhibit pulsational variability, generally on timescales of a few hundred to a few thousand days. The nature of these pulsations is thought to be complex, and they often display considerable irregularity. In particular, interaction between pulsations and large convective cells is considered likely to excite a large number of both radial and non-radial pulsation modes (Kiss et al., 2006). It is currently unknown to what extent RSG pulsations affect their mass-loss histories, although it has been claimed that RSG pulsations may be analogous to those observed in AGB stars and may therefore drive winds by similar mechanisms (Heger et al., 1997).

The possibility of high-amplitude pulsations in RSGs driving mass loss has been investigated recently by Yoon & Cantiello (2010), following work by Heger et al. (1997). Those authors investigated the possibility that such pulsations could drive a superwind to enhance mass loss in a similar way to the superwind phase that it has been claimed occurs in AGB stars. Yoon & Cantiello studied the effect that such mass-loss enhancement would have on the evolutionary estimates for the upper RSG mass limit, and concluded that it was able to reduce the evolutionary upper initial mass limit for core collapse in RSGs to 19–20 M_{\odot} . They did not, however, examine the possibility that pulsation-driven mass loss could be dynamical in nature and, because their code was incapable of resolving shocks, they terminated their simulations before any dynamical ejection events might occur.

An important constraint on RSG evolution is provided by the Humphreys-Davidson (HD) limit (Humphreys & Davidson, 1979), an empirical upper luminosity limit beyond

which cool supergiants are not observed. This limit is believed to encode information about the mass-loss histories of massive stars, with stars above a certain critical mass (and therefore luminosity) undergoing sufficiently strong mass loss to strip their hydrogen envelopes before they evolve to the RSG phase, instead evolving directly to Wolf-Rayet stars. As such, evolutionary estimates for this limit can be calculated using evolutionary models and mass-loss prescriptions. Evidence has recently been reported that the HD limit is likely to be lower than previously believed (Davies et al. 2018 claims a value of $10^{5.5} L_{\odot}$, reduced from an older estimate of $10^{5.8}$ by Humphreys & Davidson 1979), implying stronger mass loss, while at the same time it has been claimed that RSG wind losses have probably been *overestimated* (Beasor & Davies, 2018); this contributes to a developing discrepancy between theory and observation. Davies et al. (2018) postulate that this discrepancy can be resolved by the addition of a short phase of enhanced mass loss sufficiently strong to cause RSGs in this phase to present instead as OH/IR stars, with a mass-loss rate of approximately $10^{-4} M_{\odot} \text{ yr}^{-1}$. Importantly, the HD limit does not appear to be different between the Large and Small Magellanic Clouds, suggesting that the dominant mass-loss mechanism at work is not strongly metallicity-dependent.

Another consideration is the likely presence of circumstellar mass shells surrounding RSGs at the ends of their lives. The presence of narrow- and intermediate-width hydrogen emission lines in the spectra of SNe IIn is generally interpreted as being due to the shock-heating of circumstellar material (CSM) surrounding the star at the time of core collapse (e.g. Schlegel, 1990; Chugai & Danziger, 1994; Filippenko, 1997)⁴. The amount of mass present in these circumstellar shells is difficult to constrain and dependent on the theoretical model used, but may be as high as of order $10 M_{\odot}$ (van Marle et al., 2010), requiring either extreme steady mass-loss rates that may be greater than $10^{-1} M_{\odot} \text{ yr}^{-1}$ (Kiewe et al., 2012), or cataclysmic eruptive events. Less extreme objects, however, are seen to have more moderate amounts of circumstellar matter of order $0.1\text{-}1 M_{\odot}$ (e.g. Chugai et al., 2004b), more in line with existing models of envelope-driven dynamical mass-loss processes in cool giants. It has even been inferred that RSG CSM can be present in a nested shell structure indicative of repeated episodes of strong mass loss in the run-up to core collapse, as observed in the

⁴Care must be taken because SNe that have been classified as IIn do not form a homogeneous class, and some are believed to be thermonuclear SNe of type Ia (e.g. Chugai et al., 2004a) and therefore unrelated to RSGs.

case of SN 1979c (Weiler et al., 1992).

The evidence that strong mass-loss enhancement occurs in RSGs suggests strongly that dynamical mass loss may occur in this regime when considered with the fact that nonlinear pulsation models of these stars have been found by Heger et al. (1997) and Yoon & Cantiello (2010) to develop high-amplitude pulsations in which velocities reach at least the sound speed. In Chapter 5, we will examine this possibility and report on simulations extending the work of Heger et al. (1997) and Yoon & Cantiello (2010) into the supersonic regime.

1.4 Dynamical mass loss

In this section we will discuss the physical nature of dynamical instability and review the existing literature on dynamical instability and dynamical mass loss in giant stars.

1.4.1 Some definitions

We will define here some quantities that will be useful in the following discussion.

We define the dynamical or freefall timescale of a star, the timescale on which it adjusts hydrodynamically, as

$$\tau_{\text{dyn}} \approx \sqrt{\frac{R^3}{GM}}, \quad (1.2)$$

where R is the surface radius of a star, M is its total mass, and G is Newton's gravitational constant. This definition is applicable to the dynamics of the envelope only – the core is likely to have a much shorter dynamical timescale, which can be calculated by replacing R and M with the radius and mass of the core alone. The dynamics of the core has little effect on envelope pulsations, and so wherever we refer to simply the dynamical timescale, this refers to the envelope only. Typical values are a few seconds for a white dwarf, half an hour for the Sun, and a year for a red giant.

We define also the Kelvin-Helmholtz timescale, (also the radiative cooling timescale or thermal timescale), the timescale on which the envelope adjusts thermally, as

$$\tau_{\text{KH}} \approx \frac{GM M_{\text{env}}}{2RL}, \quad (1.3)$$

where M_{env} is the mass of the envelope, and L is the surface luminosity. Typical values are a billion years for a white dwarf, 16 million years for the Sun, and 100 years for a red giant.

It is worth defining exactly what is meant by dynamical mass loss: dynamical mass loss is any process which removes material from a star in inertial ejection events, in which material is accelerated to high velocities and escapes the star on (semi-) ballistic trajectories governed mainly by its inertia and gravity rather than by the stellar radiation field. We will not discuss dynamical mass transfer onto a binary companion, although this might reasonably come under this general definition of dynamical mass loss.

A star that becomes unstable to dynamical motion, in which the inertia of bulk motion within the star is significant in determining its behaviour, is liable to experience dynamical mass loss. There has been a significant amount of work considering the possibility that dynamical mass loss is important in stars, most of which deals with AGB stars and the evolution of low- and intermediate-mass stars from red giants into white dwarfs. This work dates back to before computer simulation tools were available, and the history of the idea is not straightforward, with our understanding of this mass-loss mechanism remaining fragmentary.

We will deal a lot with the concept of “dynamical stability”. As a rough definition, a star is dynamically unstable if there are states adiabatically accessible from the state of hydrostatic equilibrium which are of a lower energy and are not themselves in hydrostatic equilibrium: if perturbed from equilibrium, the star will accelerate further from equilibrium on the dynamical timescale. More formally, for stellar matter moving radially such that the mass element at initial radius r_0 is displaced by some oscillatory distance $\delta r(r_0)e^{i\omega t}$, the value of the complex frequency ω determines the stability of the star: if $\omega^2 > 0$ the star is dynamically stable, and perturbations do not grow, whereas if $\omega^2 < 0$, perturbations grow exponentially in time and the star is unstable (that ω^2 is real is guaranteed by the form of the linear adiabatic radial stellar wave equation, see Kippenhahn et al. 2012). For adiabatic pulsations, this stability condition corresponds to dynamical stability. An exact expression for ω^2 which is fully general for linear adiabatic radial pulsations was given originally by Ledoux (1945), and in a slightly clearer form by Stothers (1999):

$$\omega^2 = \frac{-\int_0^R r^3 \frac{d}{dr} [(3\Gamma_1 - 4)P] \frac{\delta r}{r} dr}{\int_0^R r^4 \rho \frac{\delta r}{r} dr}, \quad (1.4)$$

where P is local pressure, ρ is local density, and Γ_1 is the first adiabatic index $(\partial \ln P / \partial \ln \rho)_S$. This expression represents the formal condition for dynamical stability in stars. For pulsations which are confined to above some mass coordinate m_{base} (as is expected for envelope pulsations in giants), we can perform an approximate integration to obtain

$$\omega^2 \approx \frac{(3\langle \Gamma_1 \rangle - 4) \int_{r(m_{\text{base}})}^R P dV}{\int_{r(m_{\text{base}})}^R \frac{1}{3} P dV}, \quad (1.5)$$

where we make use of the pressure-weighted, volume averaged value of Γ_1

$$\langle \Gamma_1(m_{\text{base}}) \rangle = \frac{\int_{r(m_{\text{base}})}^R \Gamma_1 P dV}{\int_{r(m_{\text{base}})}^R dV}. \quad (1.6)$$

Since pressure is everywhere positive, it is clear from Equation 1.5 that $\langle \Gamma_1(m_{\text{base}}) \rangle < 4/3$ is, in this approximation, a sufficient condition for dynamical instability.

A physical interpretation of this result can be found by recalling that hydrostatic equilibrium in stars is maintained by the balance of the forces of gravity and pressure. When an element is perturbed away from hydrostatic equilibrium, both of these forces change, and how strongly the pressure changes is determined by the value of Γ_1 . For an element displaced vertically outward, the force of gravity acting on it will drop due to its new position at a higher radius, and its pressure will drop due to its expansion. If the force acting on the element due to the pressure gradient drops more slowly than the force of gravity acting on it, the element will be accelerated upward away from its equilibrium position and the arrangement is therefore dynamically unstable. This argument applies in the vicinity of the element in question; to examine the stability of the entire star (or the entire envelope) a suitable average must be used, which is defined in Equation 1.6.

Importantly, this condition is derived under the assumption that the motion of the star is adiabatic, in which case growing or decaying pulsations cannot occur. More accurate models

are provided by non-adiabatic theory, in which fully analytic results are no longer available, but approximate conditions for dynamical stability can be identified which reduce to versions of the above condition (for developments of adiabatic and non-adiabatic pulsation theory, including an impressively varied set of derivations for this condition, or approximations to it, see Kippenhahn et al. 2012, Unno et al. 1989, Cox 1980, Cox & Giuli 1968, Ledoux & Walraven 1958). As we shall discuss later, the formal concept of dynamical stability is of only limited usefulness for estimating how stars will actually behave when simulated – motions that are strongly dynamical can arise without this condition being met – but it provides the historical starting point for the theory.

1.4.2 Survey of existing work

The vast majority of the existing work on dynamical instability and associated mass loss in cool giants deals with AGB stars specifically, although some of it, especially the earlier work, deals with generic cool giant envelopes and so is not associated with any particular stellar evolutionary stage. For this reason, this review will focus on the AGB literature, with relevant results from CE and RSG theory discussed elsewhere.

The problem of how stars shed their envelopes to leave the AGB and form planetary nebulae has been studied since Shklovsky (1956), but to this day the mechanism by which this occurs remains uncertain. The first authors to discuss dynamical instability in this context were Lucy (1967), Roxburgh (1967), and Paczyński & Ziółkowski (1968), following on from work by Biermann (1938). These authors proposed that the evolution of giants would naturally lead them into a state of dynamical instability, whereupon their envelopes would be dynamically ejected. Importantly, they imagined that this ejection would occur as a single outflow event and performed calculations to support this idea under the assumption of adiabaticity. When initial hydrodynamical simulations of these envelopes were performed by Keeley (1970), however, they instead showed that the instability that emerged was pulsational in nature – no single ejection occurred. Keeley also showed the motions of the envelope to be highly non-adiabatic, with energy transport and radiative losses highly important, and so calculations performed by previous authors under the assumption of adiabaticity were not likely to be accurate.

In Keeley’s simulations, large-scale pulsations emerged, which were sufficiently violent

to cause small amounts of dynamical mass loss (although this mass loss was possibly a result of modelling limitations in this case). However, Keeley concluded that the pulsations were not strong enough to be able to explain envelope loss in AGBs. Further studies by Wood (1974, 1975)⁵ and Tuchman et al. (1978, 1979) found that large-scale pulsations could in fact result in repeated mass-loss events of significant size and suggested that this phenomenon could be the mechanism responsible for envelope loss and PNe formation.

The techniques used in the early papers discussed here are based on simulation of the giant envelope alone, with some set of arbitrary core boundary conditions (usually stationary in time) imposed below it, and a simple atmosphere prescription (or none) applied above. The early papers by Wood (1974) and Tuchman et al. (1978, 1979) are the most relevant to this thesis, as they look directly at dynamical mass ejection by large amplitude pulsations as an envelope removal mechanism for AGBs and for the formation of planetary nebulae. We shall discuss them in some detail here.

In their 1978 paper, Tuchman, Sack & Barkat (who will be referred to hereafter as TSB) specifically consider the link between formal dynamical instability of envelopes and the presence of mass ejection events, to the point of assuming that formally dynamically unstable envelopes eject mass, while formally stable ones do not. We will examine whether this assumption is reasonable in practice. They perform non-adiabatic hydrodynamical simulations of dynamically unstable giant envelopes representing AGB stars. They reason that these stars may end up in dynamically unstable configurations either as a result of

⁵These older papers refer to their pulsations as “relaxation oscillations”, an unfortunate piece of terminology which comes down to us through, for example Smith & Rose (1972), from Rose (1966, 1967) and previously from Schwarzschild & Härm (1965). The original meaning of the term “relaxation oscillation”, as initially used by Rose and other authors in the 1960s, refers specifically to a solely thermal instability that emerges in the burning shells of an AGB star and is characterised by an oscillation in nuclear burning rate (this is the same instability that we now know to be responsible for the thermal pulses, known at the time as double-shell flashes). Rose speculated that this thermal instability could trigger response pulsations in the envelope, perhaps leading to the emergence of dynamical instability as a result. Smith & Rose (1972), however, discuss pulsations in the envelope caused by an impossibility of maintaining both thermal and dynamical stability simultaneously, which they also term “relaxation oscillations”, although their oscillations are not related to burning rate variability and it is not clear that their use of the term is justified. Smith & Rose also claim that there is a clear distinction between their oscillations, which they consider to be thermal in character, and those observed by Keeley, which they recognise correctly as dynamical pulsations. In this context, however, this distinction is not meaningful, because in the case of AGBs the similarity of the dynamical and thermal timescales of the envelope does not allow the existence of fully separated thermal and dynamical oscillation modes (see below). The later employment of the term “relaxation oscillations” to describe pulsations which are dynamical in nature and driven by instabilities intrinsic to the envelope, rather than as a response to changing luminosity at its base due to variability in the burning shells, is confusing, and we will call these dynamical pulsations instead.

luminosity perturbations emerging from thermal instabilities in the burning shells or by their normal, steady evolution. TSB’s dynamically unstable models experience repeated mass-ejection events as a result of fallback shocks, which they theorise may result in complete envelope ejection over a short period of time.

TSB differentiate between their results, simulating dynamically unstable stars, and those of Keeley (1970) and Wood (1974), which they term merely “pulsationally unstable” – lacking formal dynamical instability – and they assess such “pulsationally unstable” models as describing Mira pulsations. However, it is not clear that such a distinction is physically meaningful, and no bifurcation in behaviour is seen between the two situations, a fact TSB themselves admit when considering the possibility that “pulsationally unstable” models may also pulsate violently enough to lead to mass ejections despite lying outside the boundary of formal dynamical instability, as suggested by the results of Wood in 1974 (more discussion of dynamical vs pulsational instability can be found below).

Wood’s 1974 paper reported non-adiabatic hydrodynamical simulations of an AGB envelope associated with a $0.9 M_{\odot}$ star at different core luminosities. Wood reports large-amplitude pulsations with potentially very high growth rates, and associated shock-driven mass ejections. However, these mass ejections are small, and Wood does not attach great significance to them or suggest that they are likely to lead to the removal of the whole envelope.

Importantly however, Wood concludes that violent pulsations are expected from all fundamental-mode pulsations in realistic Miras, and that, therefore, stable Mira pulsations must be overtones (this is basically in conflict with Keeley’s results, which show a stable fundamental-mode pulsation). As we now believe that the fundamental is in fact the likely dominant pulsation mode for Miras, later work by Wood has involved calculations of fundamental mode pulsators with additional damping mechanisms, suppressing their pulsation amplitudes so that steady Mira models can be constructed.

TSB’s 1979 paper builds on their earlier results to address the main objection, as they saw it, to the use of dynamical mass ejections to remove the envelopes of AGBs: that mass loss caused by such ejections could return an envelope to a state of dynamical stability after some but not all of its mass had been removed. They also investigate the region of parameter space in which large-scale pulsations occur and become violent, but within

which the envelope is not formally dynamically unstable. They simulate envelopes from two different epochs during the AGB for 4 different stellar masses. Though they do not make attempts to follow the evolution of their stars in detail, in particular they perform no simulations during thermal pulse shell flashes (or DSFs – Double Shell Flashes – as they were known in this era) or the accompanying luminosity troughs, which have an enormous effect on the luminosity and temperature of the stars and, as we will show, on their pulsation properties. TSB’s results lead them to the conclusion that formal dynamical stability does not represent a bifurcation between mass-losing and non-mass-losing models.

TSB obviate their problem from their previous paper, in which they worried that the mass loss from an envelope would have a stabilising effect, preventing further mass loss before the entire envelope had been ejected, in the following manner: they perform repeated simulations in which they arbitrarily vary the mass of the envelope, whilst keeping a constant L at the inner boundary (to represent constant core mass). They find that mass ejections continue to occur when the envelope is reduced to $0.1 M_{\odot}$ in mass, and even when it is only $0.002 M_{\odot}$. Thus, they reason that ejections are not likely to stop as the envelope loses mass, and their main objection to dynamical mass ejections as a method to form PNe is removed.

TSB also use their 1979 paper to speculate on the effects of TP shell flashes (which were not well-understood at the time). Of note is their suggestion that if the high- L stage of the flash lasts much longer than the interval between successive mass ejections, the envelope is likely to be removed entirely during the first shell flash a star experiences. They suggest that the shortening in duration of the high-luminosity state that occurs as stellar mass increases may mean that AGB stars of lower masses ($<2 M_{\odot}$) lose their envelopes early due to shell flashes, whereas higher mass stars live longer as their shell flashes are too short to remove the entire envelope (or, possibly, any of it) early on the AGB. This idea is useful to TSB because they calculate that there are more Mira variables predicted by contemporary stellar evolution theory than are observed, so a method of “removing” them early on the AGB works towards solving this problem.

The deemphasising of the formal dynamical stability criterion in Tuchman et al. (1979), supported by Wood (1974), is not inconsistent with the mathematical theory of stellar stability, although it may appear so at first. Formal dynamical instability as a concept assumes that there exists a bifurcation between stability or instability against pulsations which both

oscillate and grow on the dynamical timescale. The idea of this bifurcation is only valid, however, under certain conditions. The dynamical stability condition discussed above is derived from adiabatic pulsation theory, whereas when non-adiabatic effects are considered it is possible, in various developments of the theory (see Cox, 1980), to identify approximate stability conditions for dynamical pulsation modes which grow on the dynamical timescale, pulsational (or vibrational) modes⁶ which grow on the thermal timescale (although they still pulsate on the dynamical timescale) and secular modes which both grow and pulsate on the thermal timescale. In this context, dynamical stability remains meaningful, and the condition described above still applies (if only approximately). However, this separation of pulsation types is only valid for stars whose thermal timescale is much greater than their dynamical timescale. This is not the case in very expanded giant envelopes, where the dynamical and thermal timescales are of comparable length, and in fact the thermal timescale of the envelope may be shorter than its dynamical timescale under certain circumstances. There therefore exist no purely adiabatic pulsation solutions in this regime, and the separation of pulsations into dynamical or pulsational/vibrational is inapplicable. Thus, the fact that formal dynamical stability does not represent a bifurcation in the behaviour of pulsations is not theoretically problematic in this regime, and the result that “dynamically unstable” envelopes behave similarly to merely “pulsationally unstable” ones when simulated represents the fact that there is no clear difference between dynamical and pulsational modes in these envelopes. We will use the term “dynamical pulsations” to refer to large-amplitude pulsations in which dynamical effects are important, although such pulsations are not distinct from pulsational/vibrational pulsations in the regimes we consider.

It is worth noting that there *were* two articles published reporting single-event ejections of the entire envelope based on non-adiabatic hydrodynamical simulations. The first of these was by Kutter & Sparks (1974), although those authors artificially inserted an evacuated layer into their giant at the base of the unstable envelope, so the fact that they report shell ejection of all exterior material is not surprising, and it is not clear if this result can be relied upon generally. The second, by Stry (1975), was a calculation performed with a core luminosity equal to the Eddington luminosity for electron scattering, and for an envelope of

⁶This terminology is confusing, as it implies the possibility of non-pulsational pulsations. For some reason this was considered reasonable at the time.

only 1/6th the mass of the star. It is unlikely that a giant will reach such an extreme scenario without having already encountered instability to dynamical pulsations and associated mass loss, and it is therefore possible that Stry's ejection may be analogous to the final mass ejection event in a sequence, or that realistic stars may lose their envelopes completely before they reach this regime. Notably, Stry's envelope does experience a pulsation before its ejection event, and its behaviour is consistent with a very high growth rate fundamental mode pulsation which one would expect if extrapolating the results of TSB and Wood to high enough luminosity-to-mass ratios.

After 1980, the astronomical community seems to have undergone something of a crisis of faith, and interest in dynamical mass ejections was lost, for reasons that are not immediately clear. In discussing the literature of this period in their 1994 paper (which we will discuss below), Wagenhuber and Weiss simply say "interest in the hydrogen recombination instability [their name for the phenomenon] ceased."

The source of this crisis can perhaps be traced back to Fox & Wood (1982), who performed linear non-adiabatic calculations of pulsation periods and growth rates in red giants and supergiants. Fox and Wood use their results to claim that "the dynamical instabilities found in adiabatic models (e.g. Paczyński and Ziolkowski 1968; Tuchman, Sack, and Barkat 1978) do not occur in more realistic models which include nonadiabatic effects." This statement is based on Fox and Wood having calculated that, for the two model envelopes they examine, the two highest luminosity models they consider have infinite periods for fundamental mode adiabatic pulsations (which is equivalent to showing that they are formally dynamically unstable) but have finite periods when non-adiabatic effects are included and, in both cases, the very highest luminosity model has a slightly lower period than the next most luminous. This seems to reverse the trend generally seen that higher luminosities lead to longer periods (and thus probably larger and more nonlinear pulsations), and motivates the authors' claim that "instabilities ... do not occur". However:

1. The models reported by Tuchman et al. (1978) are not adiabatic, Fox and Wood claim this incorrectly.
2. The conclusion reached by TSB in their 1979 paper, which is also supported by Wood's own paper from 1974, is that formal dynamical instability is not necessary

for dynamical mass ejections to occur, so the finite values of the fundamental mode pulsation periods that Fox and Wood recover are not evidence against mass ejections occurring (and, indeed, the values of the finite periods they calculate for these high- L models are similar or actually longer than those reported for similar stars in the simulations of Tuchman et al. 1978, 1979 and Wood 1974 which do exhibit mass ejections).

3. The decreasing period at high luminosity reported by Fox and Wood is extrapolated in both cases from only one value difference.
4. The predictions of linear stability analysis calculations such as those performed by Fox & Wood (1982) for the periods of nonlinear pulsations are not likely to be reliable in any case, due to the effect that the presence of high-amplitude pulsations has on the structure of the envelope (see Ya'ari & Tuchman 1996 and discussion of this effect in Section 5.2.1). Nonlinear calculations such as those performed by TSB are likely to be considerably more reliable in this regime.

Unfortunately, the presence of this paper has contributed to the literature surrounding dynamical mass ejections from that period being termed “inconclusive” (e.g. by Wagenhuber & Weiss 1994 and Han et al. 1994).

Another factor which led to a loss of interest in dynamical mass ejections is the rise in popularity of “superwind” theories of mass loss from AGB stars. Although superwind models have been around for a long time, with the term dating back to Iben & Renzini (1981), the superwind model did not reach its present state of popularity until later decades, when observations consistent with superwind-style mass loss were repeatedly reported in AGB and OH/IR stars (for example Knapp & Morris 1985, Bedijn 1987, Wood et al. 1992). With the presence of this theory, it was perhaps considered that dynamical mass ejection was less evolutionarily necessary than previously believed, although current models of AGB winds remain unable to explain some of the mass-loss variability that can be deduced from circumstellar material surrounding AGB stars and PNe.

After a dearth of interest in the 1980s, the idea of dynamical envelope ejection experienced a minor renaissance in the 1990s, with papers by Wagenhuber & Weiss (1994)

and Han et al. (1994) considering the possibility that dynamical ejection events may be responsible for envelope ejection and planetary nebula formation.

Published within a month of one another, these papers approached the problem with very different strategies. Han et al. (1994) performed an analysis based not on detailed simulation of stars as they traverse the TP-AGB or calculations of their pulsational stability properties, but rather on energy considerations and on comparison to observed populations of white dwarfs and the nuclei of planetary nebulae. Han et al. performed stellar evolutionary calculations of stars on the AGB which were designed to “smear” the evolution of those stars over time, therefore avoiding the infamous numerical problems associated with simulating stars on the TP-AGB whilst allowing them to study a kind of “average evolution” of the stars. By calculating the binding energy of their stars’ envelopes, Han et al. were able to calculate the expected populations of white dwarfs and PNNi (Planetary Nebula Nuclei) if it were the case that envelope ejection occurred approximately when the envelope’s total binding energy became positive. A comparison to data for these objects produced a good level of agreement, suggesting that envelopes were removed by some fast mechanism at or near the point at which envelope binding energy becomes positive. Their result is consistent with both a sufficiently fast superwind (although less so at higher masses), and with mass loss due to some dynamical effect (either as a single event or from multiple ejections repeated on a short timescale), the main requirement being that mass loss on the AGB is in some way “switched on”.

Wagenhuber and Weiss, conversely, performed hydrostatic stellar evolution calculations of stars with a variety of masses traversing the AGB and noted that these simulations were terminated by a runaway recombination of hydrogen during the luminosity peak of a thermal pulse. Because they carried out their calculations using a stellar evolution code, Wagenhuber and Weiss were able to argue that their instability was expected to occur in actual stars, in a way which isolated envelope calculations performed for some assumed core values (which all previous work had been) are unable to fully describe.

Wagenhuber and Weiss argue that the instability their models experience which they call the “hydrogen recombination instability”, because they observe a runaway recombination of hydrogen, occurs when the envelopes of their models become dynamically unstable. However, because they reach this conclusion by performing hydrostatic simulations, they

are not able to make reliable statements about the true nature of the instability. Their claim that the instability occurs at the point of dynamical instability is not in conflict with the conclusion of Tuchman et al. (1979) that strong pulsational instability is sufficient for mass loss and formal dynamical stability may be unnecessary, as linear non-adiabatic pulsation calculations also performed by Wagenhuber and Weiss in that paper found that a long-period, fast growing pulsational instability is present in their envelopes (though of course it does not appear in their stellar evolution calculations due to their hydrostatic nature).

The natural continuation of Wagenhuber and Weiss's work is to examine the instability with hydrodynamical simulations and determine whether it is indeed pulsational in nature and similar to the results of Tuchman et al. (1978, 1979) and Wood (1974), or whether a single outflow event is in fact experienced by realistic AGB stars as they progress along the TP-AGB. However, no work of this kind has ever been published.

In the recent couple of decades, there has been significant interest in simulating AGB envelopes, in particular pulsating ones, and although most of it is of only limited relevance to the work described in this thesis, we shall briefly summarise the work that has been done in this area.

Recent simulations of AGB envelope dynamics have been focussed into three streams:

1. Simulations of AGB atmospheres studying wind-loss mechanisms and wind driving.
2. Simulations of AGB envelopes and convection performed in multiple dimensions with dedicated hydrodynamics codes
3. Simulations of AGB envelopes performed with 1-d pulsation codes (that is, envelope-only simulations) studying pulsation properties.

Stream 1:

A large volume of work has been performed in recent decades on wind loss from AGB stars, concentrating on dust-driven wind models.

This mechanism has been studied from an observational point of view, with attempts to reproduce observed Doppler line profiles (and therefore velocity structures) surrounding AGB stars, as well as an evolutionary one, deriving mass-loss rates to compare to evolutionary models and population synthesis studies.

Observationally, it is possible to directly study the structure of AGB atmospheres with both spectroscopy and interferometry, allowing us to both probe the velocity structure and dust distribution of the atmosphere, and to fit models to these observables. Dynamical models of AGB atmospheres that attempt to reproduce spectroscopic data include Willson et al. (1982), Hinkle et al. (1982), Scholz & Wood (2000) and Nowotny et al. (2010), and studies of the induced mass-loss rate in by the PEDDRO model include Winters et al. (2000), Jeong et al. (2003), Höfner et al. (2003), Eriksson et al. (2014), Bladh et al. (2015) and Liljegren et al. (2016). These studies focus on the dynamics of the AGB atmosphere itself and are adapted to deal with its specific requirements (most importantly, it is optically thin). They do not simulate the envelope below the photosphere, which renders them incapable of simulating the pulsation driving region in the partial-ionization zones of the envelope. The pulsation of the envelope constitutes an input, or assumption of these models, and generally a sinusoidal variation in radius and luminosity is assumed, possibly with artificial phase shifts and asymmetries added (for example in Liljegren et al. 2016). However, we also note that there are models available of pulsation-driven winds that do not rely on prescriptions of envelope-driven pulsations, but instead on oscillations (the term “pulsation” is perhaps inappropriate) that emerge within the atmosphere themselves, as a result of feedback between radiative heating and opacity variation – the so called “exterior kappa mechanism” in analogy to the κ mechanism which drives pulsations within giant envelopes. Studies in this regime include Fleischer et al. (1995), Höfner et al. (1995) and, notably in 2-d, Woitke (2006). However, it is not clear that neglecting envelope pulsations is reasonable in this context.

The studies of pulsations in AGB atmospheres in this stream, although of undeniable importance in the context of winds, do not have much overlap with the study of dynamical pulsations or associated mass loss, in that they do not study pulsation driving or properties, and they assume steady, limited pulsations in which dynamical effects (in the envelope) are small. They are best viewed as modelling an alternative (and possibly complementary) mass-loss regime resulting from a less violent pulsation regime.

Stream 2:

There has been much interest in studying the processes of convection that occur within

giant envelopes. Stellar evolution calculations and pulsation models are both highly dependent on parametrised models of convective energy transport and mixing to describe what is a fundamentally multi-dimensional process in 1-d, so ensuring we can accurately describe convection in giants is highly valuable. The greater part of our knowledge of convection in stars, and therefore the best calibrations of our 1-d prescriptions, comes from the Sun, and it is important to understand how the larger-scale and lower-density convective processes in giant envelopes compare, so that we can ensure we model convection correctly in that regime. The literature on envelope hydrodynamics simulations reflects this: early simulations examined convective processes in the Sun, and it has recently been possible to extend this work to giants.

Initial simulations were limited by computational power considerations to studying only small patches of convection in Sun-like stars, for example Steffen et al. (1989), Stein & Nordlund (2000), Gadun et al. (2000) (in 2-d), Vögler (2004). Later work expanded this to slices through stars and to limited simulations of giant envelopes, for example Asida & Tuchman (1997) (in 2-d), Robinson et al. (2004) (in subgiants), and eventually moved on to models of entire giant envelopes, such as those reported by Jacobs et al. (1998), Freytag & Höfner (2008) and Freytag et al. (2017).

While these studies have made considerable improvements to our understanding of convection in giants, in particular regarding 3-d effects and non-sphericity of envelopes due to giant convective plumes, these simulations can only operate for a few dynamical timescales, and therefore cannot study stellar pulsations which grow on the thermal timescale, and have not been adapted to study the bulk flows of the envelope inherent in pulsational motion. Studies of stellar pulsation in giants therefore remain in 1-d.

Stream 3:

The pulsation behaviour of red giants is complex, and there has been a lot of effort put into understanding it. In particular, these stars are known to occupy multiple distinct sequences in the Period-Luminosity plane. Originally, 5 sequences were identified by Wood et al. (1999) using data from the MACHO project's observations of the LMC (see Alcock et al. 1992). A sixth sequence was later identified by Ita et al. (2004), a seventh by Soszyński et al. (2004) and perhaps an eighth by Soszyński & Wood (2013), with two sequences

possibly being split into three sub-sequences each (see Soszyński 2007). Most of these sequences have been identified with different pulsation modes, and combinations thereof, with additional sequences appearing due to duplicity, although there remains considerable debate regarding which modes match to which sequences. For a recent summary, see Wood (2015).

Clearly, the pulsation of AGBs is important to understand if we want to make predictions of mass-loss histories with pulsation-enhanced, dust-driven wind models. In particular, it has been suggested (see Jones et al. 1983) that the onset of superwinds is due to AGB stars switching from pulsating in overtones to the fundamental mode.

Calculations of red giant pulsations that have been performed in this context fall into two categories: linear pulsation calculations which can be used to estimate pulsation periods and growth rates in the linear regime, and nonlinear simulations which include a larger range of physics. Linear pulsation analyses (e.g. Fox & Wood 1982, Ostlie & Cox 1986, Xiong 1989, Gautschi 1999, Lebzelter & Wood 2005, Wood 2015) are performed using a static model envelope structure and do not model time evolution, whereas nonlinear pulsation models explicitly follow the time evolution of pulsating envelopes. It is only these nonlinear simulations which are able to access the regime in which dynamical pulsations occur as they are highly nonlinear once they attain large amplitudes.

Nonlinear pulsation codes perform 1-dimensional simulations of the envelopes of AGB stars, relying on simple atmospheric models for their outer boundary conditions and fixed core solutions at the inner boundary. These models have been used to study pulsation periods, growth rates and limiting amplitudes of non-dynamical pulsations in AGB stars; comparatively recent examples include Bessell et al. (1996), Ya'ari & Tuchman (1996), Ya'ari & Tuchman (1999), Olivier & Wood (2005), Ireland et al. (2008).

Whereas the pioneering studies into dynamical pulsations and associated mass loss (e.g. Tuchman et al. 1978, 1979) were performed with codes that fall into this category, they have not been used for this purpose in recent decades. Indeed, simulations performed with these models have deliberately suppressed the emergence of dynamical pulsations through the use of scalable damping parameters (see Olivier & Wood 2005, Wood 2006) which may be physically reasonable in some cases, but which rely on multiple free parameters which must be chosen arbitrarily. Recent studies have focussed on the smaller, self-limiting

pulsations that occur earlier in the AGB (and in higher-order oscillation modes), rather than the violent, dynamical, fundamental-mode pulsations examined by TSB, which are likely to occur in extreme AGBs and OH/IR stars.

It is also important to note that whilst nonlinear pulsation models are able to model the emergence of envelope pulsations in a self-consistent manner, they suffer from the fact that they simulate only the stellar envelope. This allows them to introduce more complex parameterisations of envelope physics, in particular experimental energy transport models, but means that they do not necessarily match to core models in a consistent way, and makes it difficult for them to be applied to studying the changes in pulsation properties that occur as individual stars evolve.

To summarise:

Early papers in the study of unstable giants by Lucy (1967), Roxburgh (1967) and Paczyński & Ziółkowski (1968) proposed that the loss of dynamical stability of giant envelopes would trigger the ejection of the envelope and the end of the AGB phase, and showed that such instabilities were to be expected in late-stage giants. However, hydrodynamical simulations of such envelopes suggested that direct ejections did not occur, but instead the instability was pulsational in nature (Keeley, 1970). This was confirmed by subsequent numerical studies performed by authors such as Wood (1974) and Tuchman et al. (1978, 1979), with those authors finding that a series of repeated mass-loss events occurred, dynamically ejecting shells of mass. The natural conclusion, which Tuchman, Sack and Barkat drew, was that since Miras are thought to be stars on the extreme AGB close to losing their envelopes, the pulsational mass loss they saw represented the mechanism by which the AGB is terminated. However, their ideas lost popularity in the 1980s, and their results were consigned to being called “inconclusive”, a state of affairs which continues to this day.

Although more recent work such as by Wagenhuber & Weiss (1994) and recently by Lau et al. (2012) has argued that dynamical instabilities do emerge at the end of the AGB, recent work on AGB simulations has focussed on other things, and no more hydrodynamical simulations of dynamical, pulsation-driven mass-loss events have been published since the second paper by TSB in 1979.

There is therefore a clear rationale for carrying out new hydrodynamical simulations of

dynamical pulsations in AGB stars: the scientific community at large has not been convinced that the phenomena reported by TSB occur in actual stars. However, the arguments that were advanced against their work at the time are weak, and subsequent research by different groups has concluded that dynamical instabilities are in fact likely to emerge. With the use of modern simulation methods, we are in a position to re-examine and extend the work of Tuchman, Sack and Barkat's work with calculations that are both considerably more advanced in terms of the physics included and more wide-ranging in terms of the evolutionary phases under examination. We shall also discuss the process of dynamical mass loss as a possibly universal property of unstable cool giants and consider the unifying characteristics of this mechanism across CE objects, AGB stars and RSG stars.

Chapter 2

The simulation model

*Computers are like Old Testament gods;
lots of rules and no mercy*

JOSEPH CAMPBELL

In order to investigate the dynamical behaviour of stars at different points in their lives, we need to be able to calculate both their long term nuclear-timescale evolution and their short-term hydrodynamical evolution. To achieve this, we shall employ the popular stellar evolution code Modules for Experiments in Stellar Astrophysics (MESA), created by the MESA collaboration¹. This code allows us to perform simulations of stars over their entire lifetimes and is flexible enough to allow us to investigate the specific dynamical regimes we are interested in. A full description of this code and the physics encoded in it is available in articles by Paxton et al. (2011, 2013, 2015, 2018). In this section we shall discuss the physical problem that must be solved in order to simulate a star, how MESA goes about solving it, and how we utilise the code's capabilities in this work.

2.1 The basic problem

2.1.1 The equations of stellar structure

In order to simulate the evolution of a star, it is necessary to solve a set of coupled partial differential equations governing the internal structure of a spherically symmetric star known as the stellar structure equations (for a full development of the theory, see a reference work such as Kippenhahn et al. 2012 whose formulation of the equations we follow here, Cox

¹<http://mesa.sourceforge.net>

& Giuli 1968, or Schwarzschild 1958). The stellar structure equations are the equation of mass continuity

$$\frac{\partial r}{\partial m} = \frac{1}{4\pi r^2 \rho}, \quad (2.1)$$

the radial momentum conservation equation for hydrostatic equilibrium

$$\frac{\partial P}{\partial m} = -\frac{Gm}{4\pi r^4}, \quad (2.2)$$

the energy conservation equation²

$$\frac{\partial l}{\partial m} = \epsilon_n - \epsilon_\nu - c_P \frac{\partial T}{\partial t} + \frac{1}{\rho} \left(\frac{\partial \ln \rho}{\partial \ln T} \right)_P \frac{\partial P}{\partial t}, \quad (2.3)$$

the temperature gradient equation

$$\frac{\partial T}{\partial m} = -\frac{GmT}{4\pi r^4 P} \nabla, \quad (2.4)$$

and the composition evolution equations

$$\frac{\partial X_i}{\partial t} = \frac{m_i}{\rho} R_i - \frac{\partial X_i}{\partial m} \zeta, \quad (2.5)$$

where r is the radius of a given spherical shell, m is the mass contained within a spherical shell (the mass coordinate), ρ is the local density, P is the local pressure, G is the gravitational constant, l is the local luminosity – the energy flux outward through a spherical shell, ϵ_n is the rate of energy generation by nuclear reactions per unit mass, ϵ_ν is the rate of energy production as neutrino emission per unit mass, c_P is the heat capacity at constant pressure, T is the local temperature, t is time, ∇ is the temperature gradient ($d \ln T / d \ln P$), X_i is the local mass fraction for the chemical species i , m_i is the mass of one nucleon of species i , R_i is the net rate of production or destruction of species i through nuclear reactions, and ζ is the local chemical diffusion coefficient.

Clearly, the model given above is valid only for spherically symmetric stars as it operates in only one spatial dimension, and it also assumes that stellar material is optically thick, and hence radiative transport of energy is only local, and that the stellar interior is in a state of local thermal equilibrium and global hydrostatic equilibrium.

²Although the recombination energy of ions does not appear directly in Equation 2.3, it appears via the equation of state through the penultimate term ($c_P \partial T / \partial t$) which describes the change with time of the internal energy of stellar material.

In order for this system of equations to be solvable, we require formulae for the functions that encode information about the equation of state of stellar matter (ρ , c_P , P), energy transport (∇), ongoing nuclear reactions (ϵ_n , ϵ_ν , R_i), and chemical mixing (ζ). We also require appropriate surface and centre boundary conditions and a treatment of stellar mass loss. We shall discuss each of these briefly in turn.

2.1.2 The equation of state

In the formulation of the stellar structure equations given above, we have used the thermodynamic variables P and T as primary variables, which has led to ρ appearing on the right hand side of the equations. In practice, we have considerable freedom in our choice of variables, and can choose any pair of state functions to perform this role. In all cases, some quantity will appear, in our case ρ , which requires knowledge of the equation of state (EOS) of stellar matter to calculate from the primary variables. MESA provides a module, `eos`, for performing these calculations based on a set of precalculated EOS tables which blend values from the OPAL EOS tables (Rogers & Nayfonov, 2002), with values from Saumon et al. (1995), Timmes & Swesty (2000), and Potekhin & Chabrier (2010). Full details of this blend are available in Paxton et al. (2011).

2.1.3 Energy transport

The value of ∇ depends on the dominant heat transport mechanism at a given location in the star. This is determined by whether or not the star is convectively unstable at any given location. For a convectively stable star, energy transport is dominated by heat diffusion by radiation, so we write

$$\nabla = \nabla_{\text{rad}} = \frac{3}{16\pi acG} \frac{\kappa l P}{m T^4}, \quad (2.6)$$

where a is the Stefan radiation constant, c is the speed of light in a vacuum, and κ is the Rosseland mean opacity (see e.g. Weiss et al. 2004 for a thorough derivation of this result). Note that this result is averaged across radiation frequency, for which the Rosseland mean gives the correct weighting over power per unit frequency. A value for this mean opacity as a function of temperature, density and chemical abundances is computed from tabulated values by the MESA `kap` module. If a star is convectively unstable, however, convective

energy transport is likely to dominate, in which case we must estimate ∇ using a convective model. In a 1-dimensional model it is impossible to treat convection natively, so instead we must use an approximate prescription, for which MESA makes use of the mixing length theory (MLT) due to Prandtl (1925) (see Cox & Giuli 1968 for a complete formulation). The MLT provides a value of ∇ for convective regions, but the user must provide a value of a free parameter of order 1, α_{MLT} , which is the ratio of the characteristic mixing length of convective motions to the pressure scale height of the star:

$$\alpha_{\text{MLT}} = \frac{l_{\text{mix}}}{H_P}, \quad (2.7)$$

Where l_{mix} is the mixing length that convective elements are approximated as travelling in a convective eddy, and H_P is the pressure scale height $-P(dr/dp)$. There is no single clear calibration for α_{MLT} (for examples of attempts using different methods see Noels et al., 1991; Ferraro et al., 2006; Trampedach et al., 2014; Jiang et al., 2015), and commonly used values are usually between 1 and 2. The values we adopt for this parameter for any given simulation, and those of other parameters it is necessary to specify, are discussed in the respective chapters, as the most appropriate values can vary with the stellar regime being investigated.

To determine which temperature gradient is appropriate for a given region of the star, we require a convective stability condition. In regions without chemical gradients this is the Schwarzschild stability criterion,

$$\nabla_{\text{rad}} < \nabla_{\text{ad}}, \quad (2.8)$$

where ∇_{ad} is the temperature gradient of material which is changing depth within the star (and therefore pressure) adiabatically $(d\ln T/d\ln P)_s$. This is simple enough to understand: if a packet of material is displaced vertically upwards and does not equilibrate thermally with its surroundings, its temperature will change adiabatically as its pressure changes. If this adiabatic change leads the packet to be at a higher temperature than its surroundings (which change temperature according to ∇_{rad} in a radiative region), then the packet will be less dense than its surroundings and will accelerate upwards under a net buoyancy force. However, many regions within stars have molecular weight gradients due to differences in chemical composition, in which case the formal convective stability condition is the

Ledoux criterion (Ledoux, 1947), which also considers the potentially stabilising effects of composition gradients,

$$\nabla_{\text{rad}} < \nabla_{\text{ad}} + \frac{\phi}{\delta} \nabla_{\mu}, \quad (2.9)$$

where we have defined

$$\delta = - \left(\frac{\partial \ln \rho}{\partial \ln T} \right)_{P, \mu}, \quad \phi = \left(\frac{\partial \ln \rho}{\partial \ln \mu} \right)_{P, T}, \quad \nabla_{\mu} = \frac{d \ln \mu}{d \ln P}, \quad (2.10)$$

and μ is the average molecular weight of local stellar material. The size of a region of convective instability in the presence of a stabilising chemical gradient is generally smaller according to the Ledoux condition than it is according to the Schwarzschild condition. Material which is unstable according to the Schwarzschild criterion but stable according to the Ledoux criterion is unstable against an oscillatory double-diffusive instability called semiconvection (Kato, 1966) which, while inefficient at energy transport, may be important for chemical mixing and is discussed further below. Whilst we have specified the function ∇ , we have introduced a dependence on an additional function, the opacity κ .

2.1.4 Opacity

We require values for the Rosseland mean opacity, κ , for which tabulated values are provided by the MESA module `kap` based on a blend of the OPAL opacity tables (Iglesias & Rogers, 1993, 1996) with the opacity tables given by Ferguson et al. (2005) at low temperatures and densities. These tables are able to account for variations in metallicity as well as in carbon and oxygen abundances independently, which enables accurate calculation of opacities even at later evolutionary stages. Similarly, full details of this blend are available in Paxton et al. (2011).

2.1.5 Nuclear reactions

The functions ϵ_n , ϵ_ν , and R_i encode the nuclear reactions occurring within the star. In practice, the full set of all elements and isotopes appearing in a star and the reactions linking them is too complex to follow, and we must restrict ourselves to following a limited set of species and reactions known as a reaction network, which provides approximate

values for the nuclear energy generation and molecular densities throughout the star. MESA provides a set of such models in its `rates` module, which provide a network suitable for the specific stage of stellar evolution being simulated.

2.1.6 Chemical mixing

Equation 2.5 includes a term for diffusive chemical mixing. Detailed chemical mixing in stars remains one of the most ill-understood and controversial subjects in stellar evolution theory and the value taken by the diffusion parameter ζ depends on which treatments are chosen. Convection is a highly efficient mixing mechanism, and although it is advective rather than diffusive, it is sufficiently effective that in practice convective areas become fully chemically homogeneous, which is equivalent to the effect of a diffusive mixing process with a sufficiently high diffusion coefficient, which is how it is modelled in MESA. MLT is able to produce an estimate for the equivalent diffusive strength of convection, which MESA calculates in its `mlt` module. Numerous other mixing processes may be relevant in stars, of which the most important is likely to be convective overshooting. The convective stability conditions discussed above are conditions for the acceleration of convective elements, and elements at the convective boundary will have acceleration zero. However, for elements already moving as part of a convective plume, when they encounter a convective boundary their inertia is expected to carry them some distance beyond it. To simulate this, MESA incorporates convective overshoot mixing, which extends convective mixing to regions that are formally convectively stable but which are near to a convective boundary. The specific prescriptions applied to simulate this process are discussed in the respective chapters.

In comparison, radiative layers of stars have very little mixing; although there are various mechanisms that can cause mixing in these regions, including diffusion, gravitational settling of heavy nuclei, thermohaline mixing and an impressive list of rotational instabilities (see Heger et al. 2000 for numerous examples). The most relevant additional mechanism in our simulations is semiconvection. We make use of a diffusive semiconvection model in our simulations of RSGs, which is discussed in Chapter 5. Additional mixing due to rotational instabilities is not considered in our models, and the thermohaline instability does not occur in the regimes we examine so it has also been neglected³.

³Thermohaline mixing is a diffusive instability which emerges when material has a convectively stable

2.1.7 Boundary conditions

Partial boundary conditions can be provided both at the centre and the edge of the star. At the centre, we have $r = m = l = 0$, while at the stellar surface we have to fit a model for the stellar atmosphere. MESA provides a number of potential atmosphere models through its atm module. Except where otherwise indicated, we use a prescription which places the outer boundary of the simulation at the photosphere at an optical depth τ of $2/3$, where optical depth is a measure of the proportion of radiative flux which is absorbed between a given point in the star and infinity,

$$\tau(r) = \int_r^\infty \kappa \rho dr'. \quad (2.11)$$

At the photosphere, we can provide a boundary condition using the Stefan-Boltzmann law

$$L = 4\pi\sigma R^2 T_{\text{eff}}^4, \quad (2.12)$$

where L , R , and T_{eff} are the values of l , r , and T at the photosphere, and σ is the Stefan constant. To provide a second boundary condition, specifying the pressure, we apply the Eddington Grey Atmosphere model, performing a numerical integration of the hydrostatic balance equation assuming the mass of the atmosphere is negligible,

$$\frac{dP}{d\tau} = \frac{GM}{R^2\kappa}. \quad (2.13)$$

The value of κ in this equation is estimated implicitly using the Eddington approximation (Eddington, 1926)

$$T^4(\tau) = \frac{3}{4} T_{\text{eff}}^4 \left(\tau + \frac{2}{3} \right). \quad (2.14)$$

Now that we have provided boundary conditions for two variables at the centre and two at the surface, the system is properly defined and can in principle be solved.

temperature stratification but an unstable chemical stratification – denser species lying above less dense ones. This situation can emerge in stars due to accretion of material transferred from a more evolved companion – it does not appear in our models. The process occurs in Earth’s oceans, from where it acquires its name (haline referring to differing levels of salt content in seawater).

2.1.8 Mass-loss treatment

One important factor remains to be considered, and that is the loss of stellar mass over time. As discussed above, this process is highly uncertain, and different prescriptions must be adopted for different stellar regimes. For the AGB models discussed in Chapter 4, we make use of the wind prescription given by Schröder & Cuntz (2005),

$$\dot{M} = \eta \frac{LR}{M} \left(\frac{T_{\text{eff}}}{4000 \text{ K}} \right)^{3.5} \left(1 + \frac{g_{\odot}}{4300g} \right), \quad (2.15)$$

where \dot{M} is the mass-loss rate in units of solar masses per year, η is a calibrated constant with value $8 \times 10^{-14} M_{\odot} \text{ yr}^{-1}$, and g is the acceleration due to gravity at the stellar surface GM/R^2 . As may be evident from the form of this equation, it represents a physically motivated semi-empirical correction to the well-known Reimers law (Reimers, 1975).

For the RSG models in Chapter 5 we adopt the ‘‘Dutch’’ wind prescription following Glebbeek et al. (2009). This prescription uses a theoretical mass-loss model given by Vink et al. (2001) for effective temperatures in excess of 10^4 K

$$\begin{aligned} \log(\dot{M}) = & -6.688 + 2.210 \log(10^{-5} \times L/L_{\odot}) \\ & - 1.339 \log(M/M_{\odot}/30) \\ & - 1.601 \log(v_{\infty}/v_{\text{esc}}/2) \\ & + 1.07 \log(T_{\text{eff}}/20000) \\ & + 0.85 \log(Z/Z_{\odot}), \end{aligned} \quad (2.16)$$

where $v_{\infty}/v_{\text{esc}}$ is the ratio of the terminal wind velocity to the surface escape velocity of the star and Z is the surface metallicity. We have multiplied this wind by an additional scaling factor of 0.3 to correct for the significant overestimation of observed mass-loss rates by theoretical prescriptions in massive stars (see Puls et al. 2008 for a discussion). At effective temperatures below 10^4 K , the well-known empirical formula due to de Jager et al. (1988) is applied, which can be found tabulated in full in that paper.

The mass loss rate \dot{M} does not appear in the stellar structure equations and must be added separately (see below).

2.2 The method of solution

MESA solves the stellar structure equation using a technique known as the Henyey Method (Henyey et al., 1964; Hofmeister et al., 1964; Kippenhahn et al., 1967). MESA’s formulation of the system operates in two dimensions: time and a radial spatial dimension, which we have chosen to parameterise by mass coordinate. Following the Henyey method, the solution is discretised in time, breaking it into a sequence of individual, static stellar models separated by some distance in time (the timestep), and in mass, with each model being made up of order 1000 numerical values of the 5 principle variables on a “grid” or “mesh”⁴. The size of the timesteps between models and the mass differences between meshpoints within a model are both adaptive and are refined or expanded as necessary to ensure sufficient resolution in both dimensions.

Individual models in the sequence are constructed by performing a simultaneous rootfind for solutions of all the structure equations at every meshpoint. Recalling that Equation 2.5 is in fact a series of equations, one for each chemical species being considered, this results in a simultaneous rootfind being performed for typically of order 100,000 coupled equations. This is accomplished using a multi-dimensional extension of the Newton-Raphson method.

If we form a vector \vec{x} from the values of all structure variables at all meshpoints in some candidate solution for which the residuals from the structure equations are $\vec{F}(\vec{x})$, we can perform the iterative procedure

$$\vec{x}_{n+1} = \vec{x}_n - \mathbf{J}_F^{-1} \vec{F}(\vec{x}_n), \quad (2.17)$$

where \mathbf{J}_F is the Jacobian matrix $\mathbf{J}_{Fij} = (\partial F(\vec{x})_i / \partial x_j)$. In practice, we can construct the Jacobian much more easily than its inverse, so to avoid having to invert an order 100,000 element-per-side matrix, MESA instead solves the system of equations

$$\mathbf{J}_F(\vec{x}_{n+1} - \vec{x}_n) = -\vec{F}(\vec{x}_n), \quad (2.18)$$

for the correction vector $\vec{x}_{n+1} - \vec{x}_n$ which must be added to our candidate solution to produce a (hopefully) better solution. This iterative procedure has the advantages of

⁴These terms continue to be used despite the fact that this grid or mesh is 1-dimensional, somewhat weakening the metaphor.

the Newton-Raphson method – efficient (quadratic) convergence to a solution – but also the disadvantages, most significantly that the candidate solution we start with must be sufficiently close to the true solution or the method may fail to converge entirely. To solve this problem, the first candidate solution for each stellar model in an evolutionary sequence is based on the converged solution of the previous model in the sequence. If the timestep is small enough, these models are usually similar enough to enable the solver to function effectively; if not, if the residuals of the structure equations are too high, the timestep is reduced. Having constructed a stellar model in this way, time-dependent adjustments are made, principally to the stellar mass to account for mass loss, and another model can be constructed after the next timestep has elapsed, with the necessary changes to mass and composition. In this way a complete evolutionary sequence can be constructed.

It should be noted that the form of the structure equations used by the MESA Newton solver is different from those given above. The numerical stability and computational efficiency of the system is improved by using finite-difference forms of the equations (see Sugimoto et al., 1981). The selection of different algebraically-equivalent formulations can also be more or less suited to different stellar regimes. Full details of this scheme can be found in Paxton et al. (2011, 2013).

The initial models from which evolutionary sequences begin are constructed as uniform composition, non-burning “pre-main-sequence” models with polytropic density profiles, which contract from some arbitrary initial radius to form model zero-age main-sequence stars. This method is used entirely for computational convenience and is not intended to reproduce actual pre-main-sequence evolution (see Stahler et al. 1980, Palla & Stahler 1993, Bernasconi & Maeder 1996).

As mentioned above, changes to the mass of the stellar model are done separately from solving the stellar structure equations. The process of changing the mass of the model due to mass loss occurs before the stellar structure equations are solved for a particular timestep. The mass values of each point in the grid are adjusted to enact this change. A full description of this process is available in Paxton et al. (2011) and Paxton et al. (2013).

2.3 Hydrodynamics

The treatment given above describes a star under the assumption of hydrostatic equilibrium. Clearly we cannot learn much about the dynamical behaviour of stars under such an assumption. We can generalise the stellar structure equations to the hydrostatic case by adding an additional structure variable for radial velocity v . We add an acceleration term to Equation 2.2:

$$\frac{\partial P}{\partial m} = -\frac{Gm}{4\pi r^4} - \frac{1}{4\pi r^2} \frac{\partial v}{\partial t}, \quad (2.19)$$

and an additional equation

$$\frac{\partial r}{\partial t} = v, \quad (2.20)$$

which is trivially the definition of velocity. All derivatives that appear in the stellar structure equations should now be interpreted as Lagrangian derivatives. It is not necessary to add another term to the energy conservation equation (Equation 2.3), as the final term in that equation already encodes work done in accelerating material. We add also an additional boundary condition, which is that at the centre of the star the velocity is zero. These changes provide us with a hydrodynamic stellar evolution code, and MESA has the capability to operate in this way, in what we will call its “hydrodynamics mode”. In order to ensure we obtain meaningful results from hydrodynamical simulations performed in this way, however, it is necessary to make some additional changes. Importantly, we must restrict the timestep between models to one sufficiently short to resolve dynamical behaviour – typically, this requires timesteps considerably shorter than the dynamical timescale. Choosing too long a timestep results in artificially damping the hydrodynamical motion of the star. Of course, because the timescales on which dynamical behaviour occurs are so much shorter than the nuclear evolutionary timescales of stars, it is not feasible to compute entire stellar lifetimes in this manner. We therefore adopt a simulation scheme using both hydrostatic and hydrodynamic simulations, relying on hydrostatic simulations to follow the long-term nuclear-driven evolution of the star, and performing hydrodynamical simulations at individual epochs along this evolution to sample the star’s dynamical behaviour. This does however introduce a risk that the hydrostatic evolutionary simulations will enter states

which the true hydrodynamical system would not reach due to dynamical instabilities. This concern, along with other details of this scheme, is discussed in detail in the following chapters.

Violent hydrodynamical behaviour is a basic feature of dynamically unstable stars, and it is this very violence which can lead to dynamical mass loss. In practice, this means that hydrodynamic shocks are likely to appear in these stars. Simulating these shocks in a meaningful way is not trivial, because the spatial scale of shock fronts is in reality likely to be considerably smaller than the resolution of our models. In addition, very strong gradients of structure variables (especially pressure, velocity and density) are likely to develop across shocks, which can cause numerical instability in the Newton solver. In order to treat shocks in a meaningful way, MESA employs an artificial viscosity treatment which adds additional viscosity in regions of compression in a manner originally suggested by Von Neumann & Richtmyer (1950), and uses the formulation due to Weaver et al. (1978). This viscosity has the effect of spreading the shock front over a larger radius interval so it can be resolved effectively. A dynamical viscosity (units $\text{g cm}^{-1} \text{s}^{-1}$) is added with a value given by

$$\eta = 3l_\eta^2 \pi \rho^2 r^4 \max(0, \nabla v), \quad (2.21)$$

where l_η is a dimensionless constant describing the desired width over which the shock front should be spread (see Dorfi, 1998). Throughout this work, we adopt a value for l_η of 0.05 in our hydrodynamical simulations of CE and AGB stars, which we found to significantly increase numerical stability without having a strong effect on our results. In our models of RSG stars, we use a value of 0.02, which provided the same numerical advantages in the massive regime despite the reduced value.

To simulate dynamical mass-loss events in which matter is ejected from the stellar surface, we need to be able to remove the ejected material from the simulation, otherwise the code will continue to treat it as part of the star and the stellar surface will follow this matter out to infinity. This is clearly undesirable, as ejected material ceases to affect the star and does not need to be simulated. In addition, as ejected material expands its temperature and density decrease continually and it will become optically thin (at least until it starts to form dust). As this happens, the energy transport assumptions and outer boundary prescriptions of our model become increasingly inappropriate.

To avoid this problem, we employ a custom mass-loss routine to remove unbound material which is expanding at above the local escape velocity: whenever a contiguous layer develops at the surface of the model containing matter which exceeds the local escape velocity at every point, the material in that layer is removed from the model’s surface. We achieve this in MESA by implementing a wind-loss scheme which removes this unbound matter from the surface, causing the unbound mass to decrease exponentially with a time constant of 0.01 years. At each timestep, we find the mass m of the ejecting layer (if one is present), and apply mass loss at a rate of

$$\dot{M} = 100/\text{yr}^{-1} \times m. \quad (2.22)$$

This mass-loss rate was chosen to ensure that the timescale for mass removal is at least an order of magnitude below the dynamical and thermal timescales of the star at all times, whilst not being so fast as to introduce artefacts into the dynamics of the star. In comparative tests, we found that varying this rate did not have a significant effect on the total amount of mass removed in a given ejection event. In practice, the removal of an entire escaping layer is accomplished by this method over approximately a month of star time. The algorithm used to generate this mass loss is available in full in Appendix A.

Selecting the most appropriate set of parameters for a given stellar evolution simulation is difficult (to illustrate, MESA version 8845 has 1,700 individually-adjustable parameter settings), and the required parameters vary with the regime under study; a detailed description of the simulation parameters in use in a given simulation can be found in the chapter in which that simulation appears.

Chapter 3

Common-envelope objects

The work described in this chapter has been published as

Episodic mass ejections from common-envelope objects

Clayton M., Podsiadlowski Ph., Ivanova N., Justham S.; 2017, MNRAS, 470, 1788

The slow spiral-in phase of the CE process remains less well studied than the dynamical plunge-in due to the difficulties in simulating it with 3-d hydrodynamics codes. Although the capabilities of those simulations continue to increase, and recently it has been possible to follow the evolution of a CE object into the beginning of this phase (Ivanova & Nandez, 2016), the relatively long timescale on which CE objects evolve during this phase (of order 1000s of years) still prevents us from following it in its entirety in 3 dimensions, and the majority of our knowledge of this phase is still dependent on 1-d simulations.

As such, there remains significant uncertainty about how the shared envelope can be removed in this phase, in which we generally cannot rely on an ejection being driven by rapid dynamical evolution of the embedded binary (with the notable exception of systems which are liable to undergo delayed dynamical ejection). It is likely, therefore, that envelope ejection in this phase may rely at least as much on the dynamical properties of the envelope as it does on those of the binary.

Ivanova et al. (2015) (hereafter I15) carried out 1-dimensional simulations of a Red Giant envelope during the slow spiral-in phase of a synthetic CE event. Using these simulations, the authors of I15 were able to examine the effects on a giant envelope of the heat released by an embedded binary. In many of their models the combination of this

heating and the recombination energy released as the envelope expanded was able to render the envelope dynamically unstable. This result, in addition to demonstrating the importance of recombination energy to the overall CE energy budget, indicates that the energy input rates expected to be seen in the slow spiral-in phase are widely sufficient to destabilise the envelope and render it liable to undergo a later dynamical ejection. However, because the simulations reported by I15 were based only on hydrostatic stellar models they were unable to accurately describe the evolution of the envelope after it became dynamically unstable.

In this chapter, we describe a method for the hydrodynamical simulation in 1-dimension of a Red Giant primary undergoing a similar synthetic CE event (see below) using the stellar evolution code `MESA` and discuss the results of these simulations and their implications for dynamical mass loss during CE evolution. The simulations reported here have been designed to match closely the initial parameters used by I15 within this independently-developed code, in order to facilitate comparison. As both that study and this work make use of fully-featured stellar evolution codes, both sets of simulations contain all the relevant stellar physics which is accessible in 1 dimension, up to the fact that I15 carried out their simulations under the assumption that their giant envelope was hydrostatic. In this chapter, we first reproduce the hydrostatic results of I15, then we relax the hydrostatic assumption and make use of the hydrodynamics treatment in `MESA`. This distinction allows us to follow the dynamical behaviour that arises as these models destabilise.

In Section 3.1, we describe the model star we use and the parameters of our simulations; in Section 3.2, we report on the results of those simulations; in Section 3.3, we discuss the interpretation and implications of these results for the study of CE evolution and the study of giant envelopes in general; and in Section 3.4, we will provide a brief summary.

3.1 Simulation methodology

The most common method used for 1-dimensional studies of CE events is to co-opt a stellar evolution code to model a lone giant star, and then to simulate the presence of an embedded companion by adding quantities such as heat, gravitational mass, and angular momentum at appropriate locations within the envelope. In this manner, it is possible to study the response of the envelope to an ongoing CE event in a slow phase without being forced

to make the compromises in included physics and resolution that would be necessary in a 3-d simulation. In this chapter, we apply this method using the hydrodynamical treatment available in MESA and make use of the dynamical mass removal mechanism described in Chapter 2. All calculations reported in this chapter were performed using version 7624 of MESA.

It is important to consider the usefulness of a simple 1-dimensional model of a CE object within the context of the wider project of seeking to understand CE evolution. 3-d simulations represent the state-of-the-art in modelling the dynamical evolution of CE objects, with recent examples including Ricker & Taam (2012); Nandez et al. (2015); Ohlmann et al. (2016); Staff et al. (2016); and Iaconi et al. (2017). In contrast, while 1-d models clearly lack the ability to follow the full CE dynamics, they do provide the opportunity to examine some of the physics relevant to the CE process within a more conceptually simple framework and to explore much larger regions of the parameter space due to their vastly computationally cheaper nature. For these reasons, 1-dimensional CE models remain important tools to study the CE process (see Ivanova & Nandez 2016 for discussion of the ongoing 1-d simulation effort). In particular, 1-d models have proven useful for studying the slow spiral-in phase, in which the timescales of the CE process are sufficiently long that 3-d simulations of the entire phase remain prohibitively expensive (Ivanova et al. 2013b; Ivanova 2017), and the deviations from spherical symmetry in the envelope are less extreme than during the plunge-in phase within which most 3-d simulations have been performed (Ivanova & Nandez, 2016). A 3-d simulation of the start of the slow spiral-in was reported in Ivanova & Nandez (2016), and has been used by those authors to provide a valuable analysis of the prescriptions of 3-dimensional physics of the CE process in use in 1-d codes.

No 1-d model can be considered to accurately describe CE evolution. Rather, 1-d models should be thought of as one of a set of available tools which can be applied to gain insight into the physics of the CE process. As a pertinent example, the nature of the physics underlying the “shell-driven ejection” process which appears in the 3-dimensional simulations reported by Ivanova & Nandez (2016) was only understood when compared to the simulation results presented in this chapter, and is now believed to represent the same ejection-driving process as is described below (Ivanova, private communication 2016). At

the conclusion of this chapter in Section 3.3.4, and in Chapter 6, we make recommendations for ways in which the results presented here can be synthesised with information extracted from 3-d simulations to produce a more nuanced version of the commonly used alpha formalism for CE evolution.

The initial model we use in this chapter has been chosen to match as closely as possible the one used in I15 – a $1.6 M_{\odot}$ RG. As that work used a different stellar evolution code (see Ivanova & Taam 2004; Podsiadlowski et al. 2002 for a detailed description), small differences in the model cannot be avoided, but we endeavoured to reproduce it as closely as possible. We chose to match our giant’s core mass to the value used in I15, and tuned the mixing length parameter α_{MLT} arising in the Cox & Giuli formulation of mixing length theory (Cox & Giuli, 1968) to obtain approximately the $100 R_{\odot}$ star used in that paper. We evolved a metallicity of $Z = 0.02^1$ star of $1.6 M_{\odot}$ from the zero-age main sequence (ZAMS) through to the Red Giant branch, stopping its star’s evolution when its helium core, which we define as that region within which the fractional mass density of hydrogen X drops below 10^{-10} , reached $0.422 M_{\odot}$. Selecting a mixing length parameter, α_{MLT} , of 1.95 gave this star a radius of $100.7 R_{\odot}$. We adopted a simple Eddington grey atmosphere model and neglected wind mass loss. We shall refer to this giant model as the “initial model”, which serves as the starting point for all simulations reported in this chapter.

In order to study the response of this initial model to a simulated CE event, we must emulate the presence of an embedded binary companion within the envelope. This is accomplished by injecting additional heat into the envelope to represent the frictional dissipation of the binary’s orbital energy during a self-regulated spiral-in. The distribution of such heating during CE evolution is not well understood (see Taam & Sandquist, 2000), with multiple processes involved which are expected to dump heat in different regions of the envelope. The relevant heating mechanisms are likely to be

- **Hydrodynamical drag** – Energy can be dissipated directly from the binary by drag forces, which would tend to release heat near the radius of the companion’s orbit (as in Taam et al., 1978). The exact nature of this drag is unclear: viscous drag in a boundary layer around the orbital elements may be dominant, but its strength is highly

¹The metallicity Z is defined as the fraction of a star’s mass which is made up of elements that are not hydrogen or helium.

dependent on the effective viscosity, and therefore whether the region of the envelope the binary moves through is convective or radiative (Ivanova, 2002); aerodynamic drag due to the pressure differential caused by the presence of a turbulent wake may be more effective, especially in radiative regions, and it is unclear whether or not the motion of the secondary relative to the surrounding matter will be supersonic – in which case dissipation would occur primarily in a detached bow shock.

- **Tidal drag** – It has been claimed that tidal forces may be highly efficient in dissipating energy from the orbit (e.g. Taam et al., 1994), which would lead to a situation where the regions interior to the binary are locked in corotation with its orbit (as in Meyer & Meyer-Hofmeister, 1979). This would minimize viscous drag due to the low relative velocity between the secondary and the surrounding material, and instead the main source of heating would be viscous shear drag between differentially rotating layers of the envelope in an extended region outside the binary orbit. It has been claimed, however, that this model leads to an underestimate of the drag on the binary because the true effect of gravity in this scenario is not one of tidal locking but of direct gravitational drag on the orbital components (Ricker & Taam, 2012).
- **Accretion** – A simple Bondi-Hoyle-Lyttleton estimate for the accretion rate onto the secondary during the CE event suggests that it may accrete at the Eddington limit and release energy at a comparable rate to orbital drag and tidal interactions (Ivanova, 2002; Voss & Tauris, 2003), but hydrodynamical simulations of this process have suggested that these simple estimates may be vastly too large (Ricker & Taam, 2012; MacLeod & Ramirez-Ruiz, 2015), and it remains uncertain whether this process contributes significantly to the energy budget. It is also possible that accretion processes may drive jets which may enhance mass loss or act to heat the envelope in “bubbles” perpendicular to the orbital axis (Soker, 2004; Akashi & Soker, 2016; Shiber et al., 2017).
- **Nuclear burning** – Of course, the ongoing nuclear luminosity of both the giant core and the secondary (if a star) may be considerable, but it is also important to consider the effect of burning which may occur in accretion flows of, for example, hydrogen-rich material from a secondary onto the core of the primary, or onto the

secondary if it is itself a stripped star or a compact object. These effects, if they are important, are however only likely to occur in cases when either the secondary or the core of the primary overflow their Roche lobes in the final stages of CE evolution and the beginning of a slow merger (Ivanova, 2002; Ivanova & Podsiadlowski, 2003; Podsiadlowski et al., 2010).

We follow I15 in adopting two alternate prescriptions for the heating. In one case, which we will refer to as “base heating”, heat is added to a thin layer at the base of the envelope (but still within the convective zone), where an embedded companion is expected to orbit during the self-regulated spiral in. In contrast, we also have the “uniform heating” case, in which heat is added throughout the entire convective envelope of the star. These two prescriptions therefore represent the opposite extremes of heat distribution during a CE event. In both cases, heat is added at a constant total rate per unit time.

In both of our heating cases, we define the heating region by Lagrangian mass coordinate within the star (where we define the mass coordinate of a given spherical shell within the star as being equal to the total amount of mass contained within that shell). These heating regions are fixed at specific mass coordinates at the beginning of our simulations and move with the Lagrangian motion of the stellar material. Heat is then added within the specified region at a uniform rate per unit mass. The base heating case injects heat into a layer $0.1 M_{\odot}$ thick, starting at the base of the initial convective envelope at a mass coordinate of $0.4257 M_{\odot}$, and ending at $0.5257 M_{\odot}$, whereas the uniform heating case has the same inner boundary but has its outer edge at $1.5995 M_{\odot}$, at the top of the convective zone. Rather than heating beginning sharply at the edges of these regions, in both cases we smooth the rate of heat injection linearly over a region $0.01 M_{\odot}$ in mass on the inside of the boundaries.

The rates at which we have added heat to the envelope replicate those used in I15, which are intended to simulate the slow spiral-in of a $0.3 M_{\odot}$ companion. We can estimate the total energy released from orbital decay as

$$E_{\text{tot}} = GM_2 \left(\frac{M_{\text{core}}}{r_f} - \frac{M_1}{r_i} \right), \quad (3.1)$$

where G is Newton’s gravitational constant, M_1 and M_2 are the masses of the primary and the secondary, M_{core} is the mass of the primary core alone, and r_i and r_f are the initial and

final orbital radii. If the slow spiral-in phase begins when the secondary is located at the base of the giant’s convective envelope at approximately $12 R_{\odot}$, and the final orbital radius is taken to be the Roche limit (Roche, 1849) of the secondary with respect to the giant core, given to a reasonable approximation by

$$R_{\text{Roche}} \approx R_2 \left(\frac{M_1}{M_2} \right)^{1/3}, \quad (3.2)$$

where R_2 is the radius of the secondary, which for a main sequence star of $0.3 M_{\odot}$ is approximately $0.3 R_{\odot}$, then we recover a total available orbital energy of approximately 10^{48} erg before the secondary is disrupted. A reasonable range of timescales for the deposition of this energy into the envelope is 10–1000 years (see Ivanova et al., 2013b), which gives us energy deposition rates of between 10^{45} and 10^{47} erg yr⁻¹. This added heat and the star’s existing nuclear heating (which is almost completely constant on the timescales we’re dealing with) are the only energy sources present in our model.

In addition to the two heating cases, the simulations we present are run either with or without the use of a hydrodynamics scheme. In cases without a hydrodynamic treatment, the inertial terms in the stellar pressure equation are ignored, which amounts to an assumption that the star is in hydrostatic equilibrium. In all our hydrodynamic simulations, we apply artificial viscosity in order to allow MESA to resolve hydrodynamic shocks; we adopt a shock width of 5% of the local radius for this purpose. Variations of the numerical parameters of our simulations were found not to affect the qualitative nature of our results.

In our hydrodynamic simulations, unless otherwise stated, we make use of the custom mass-loss routine described in Chapter 2 to remove unbound material which is expanding at above the local escape velocity. If part of the mass within the heating zone is removed by our mass-loss routine, the specific heating rate increases to keep the total heating rate constant.

3.2 Results

3.2.1 Hydrostatic simulations

We carried out hydrostatic simulations in both uniform heating and base heating cases with heating rates of 10^{45} , 2×10^{45} , 5×10^{45} , 10^{46} , and 10^{47} erg yr⁻¹. The evolution of the outer




Figure 3.1: The radii of hydrostatic models as a function of the total heating energy deposited where the artificial heating is uniformly spread throughout the convective envelope. All simulations begin from the same $100 R_{\odot}$ model at $E_{\text{heat}} = 0$. The blue, pink and green cases terminate due to convergence failures, the red case is terminated when the surface velocity reaches twice the escape velocity.




Figure 3.2: As Fig. 3.1 with artificial heating restricted to the base of the convective envelope. Simulations are terminated when their surface velocities reach twice the escape velocity.

radii of these models can be seen in Figs 3.1 and 3.2. These simulations agree closely with the results presented in I15, despite the independent evolution code and implementation we use. As in I15, we see that in the base heating case only the lowest energy deposition rate leads to a stable hydrostatic model, with all higher heating rates causing an accelerating expansion of the star that quickly reaches and exceeds escape velocity (these simulations were all terminated when the surface reached double the escape velocity). Of course, this outcome is only meaningful in a hydrostatic simulation insofar as it shows us that the hydrostatic assumption is inappropriate in this regime due to instability of the envelope, and we must switch to a hydrodynamic treatment. The lowest heating rate of 10^{45} ergs yr⁻¹, although it does not lead to the uncontrolled and clearly hydrodynamical expansion seen at higher heating rates, does nonetheless cause the star to expand by a factor of 4 in radius, and causes the hydrogen and helium partial-ionization zones to expand sufficiently in both radius and mass coordinates to render the majority of the star’s envelope formally dynamically unstable, with $\langle \Gamma_1(m) \rangle < 4/3$, strongly implying that the true evolution of this model may be out of hydrodynamical equilibrium also.

The uniform heating cases also tend to match the results of I15, but with one notable difference: in I15, the models with the second and fourth highest heating rates, 10^{46} and 2×10^{45} ergs yr⁻¹, expanded into converged hydrostatic models; however, the model in between these two, 5×10^{45} ergs yr⁻¹, also expanded to its equilibrium luminosity but then destabilised and failed to converge – an outcome expected when there exists no unique hydrostatic and thermal equilibrium solution. As this numerical instability is thought to reflect a real, physical instability in the envelope, the greater numerical stability of the model with the faster heating rate is unexpected. We do not recover this result; instead, in our simulations the 10^{46} and 5×10^{46} ergs yr⁻¹ models both destabilise, with the larger heating rates destabilising faster (although the 5×10^{46} ergs yr⁻¹ case appears to destabilise first in Fig. 3.1, as the x -axis of that figure is in energy, the higher heating rate does destabilise first in time). This outcome is more aligned with the natural expectation that greater energy deposition rates are more destabilising to the model. The lines plotted in Fig. 3.1 representing these three models all end where the simulations terminate due to convergence failures; by lengthening the timesteps used in our simulations it is possible to artificially delay these failures somewhat, or in the 5×10^{46} ergs yr⁻¹ case, to suppress

the failure entirely, in which case the models remain for a while in a state of unstable equilibrium. These equilibrium states also display extremely large partial-ionization zones and associated low values of $\langle \Gamma_1(m) \rangle$, indicating that much of the envelope is dynamically unstable. The simulation with the very largest rate of uniform heating experiences the same rapid expansion as seen in the equivalent base heating case and quickly exceeds double the escape velocity, whereupon that simulation is terminated.

The runaway expansion seen in both uniform and base heating cases and the destabilising of the higher uniform heating rates all clearly indicate that the models in question cannot be hydrostatically stable, and that we must use a hydrodynamic treatment instead to perform meaningful simulations of the evolution of these stars after they destabilise. In fact, as we shall see shortly, the use of a hydrodynamical code leads to significant differences in the behaviour of the majority of these models, even before the point where their behaviour becomes obviously dynamical.

3.2.2 Hydrodynamic simulations

The hydrostatic simulations reported above make it clear that a hydrodynamic treatment is necessary to understand the evolution of these unstable models. We carried out a large number of hydrodynamical simulations in both the base and uniform heating cases with different heating rates. A selection of these simulations showing the regimes of behaviour that arise are presented in Figs 3.3 & 3.4, which show the evolution of the radius and mass of our simulated giant over time for different heating rates. The very lowest and highest heating rates shown in our hydrostatic simulations do not appear in these figures, as they behave the same as the hydrostatic cases: an expanded equilibrium state for a heating rate of 10^{45} ergs yr^{-1} , and a direct ejection in the case of 10^{47} ergs yr^{-1} . At intermediate heating rates, however, occupying almost the entire parameter space between these extremes, we see the emergence of large scale pulsations which come to dominate the evolution of these models.

The behaviour of these pulsations varies strongly with heating rate, and the cases shown in Figs 3.3 & 3.4 are a sample chosen to demonstrate the different regimes that emerge. In all cases in which pulsations develop (which is all heating rates other than some very low values just above our minimum and some rapidly ejecting models just below the maximum), we see




Figure 3.3: The surface radii (blue) and masses (red) of a selection of hydrodynamical models in the uniform heating case, chosen to be representative of the behaviours observed in these simulations. Each panel describes a model heated at a different rate. The relevant heating rate appears above each panel. Sudden mass changes accompany dynamical ejection events, where the outer layers of the model have exceeded escape velocity and have been removed (see Section 3.2.2). The grey lines are the evolution of radius for the other heating rates, overplotted for comparison.




Figure 3.4: The surface radii (blue) and masses (red) of a selection of hydrodynamical models in the base heating case, chosen to be representative of the behaviours observed in these simulations. Each panel describes a model heated at a different rate. The relevant heating rate appears above each panel. Sudden mass changes accompany dynamical ejection events, where the outer layers of the model have exceeded escape velocity and have been removed (see Section 3.2.2). The grey lines are the evolution of radius for the other heating rates, overplotted for comparison.

the rapid, exponential growth of those pulsations, quickly reaching very large amplitudes and becoming supersonic. These pulsations continue to grow until damped by nonlinear effects (discussed in section 3.2.3 below). The growth rate of the pulsations increases with heating rate and can reach extremely high values, with the growth timescale reaching the order of the pulsation period for higher heating rates and the amplitudes attaining limiting size within one pulsation cycle.

The behaviour of the pulsations seen in each of our models falls into one of three regimes:

1. **Self-limiting** – At low heating rates, such as 1.3×10^{45} ergs yr^{-1} in the uniform heating case (Fig. 3.3) and 1.2×10^{45} ergs yr^{-1} in the base heating case (Fig. 3.4), the pulsations grow exponentially over multiple pulsation cycles until nonlinear effects act to limit their amplitude (usually at the point at which these pulsations become supersonic and shocks begin to develop during the compression phases of the pulsations). These pulsations then proceed to grow again in cycles of alternating growth and modulation.
2. **Ejecting** – At intermediate heating rates, such as 1.7×10^{45} ergs yr^{-1} for uniform heating and 1.5×10^{45} ergs yr^{-1} for base heating, the pulsation amplitudes grow sufficiently large and the velocities sufficiently supersonic that the compressions and accompanying shocks which occur are strong enough to dynamically eject shells of matter from the surface of the models at greater than the star’s escape velocity. These ejected shells, which can be up to $\sim 0.1 M_{\odot}$, are completely unbound and represent a form of dynamical mass loss from the star. The matter in these ejected shells is excised from the simulation after it exceeds escape velocity, and we follow the evolution of the matter which remains bound. The points at which mass ejections occur can therefore be seen by the rapid changes in remaining model mass (the red line) in Figs 3.3 & 3.4, where a shell of material has been removed. The ejection phenomenon is discussed in more detail in Section 3.2.4 below.
3. **Non-ejecting** – The third regime, seen in cases such as 2.5×10^{45} ergs yr^{-1} and 10^{46} ergs yr^{-1} for uniform heating, and 2×10^{45} ergs yr^{-1} for base heating, pulsations reach amplitudes large enough to launch ejections, yet no ejections occur. Instead,

the shocks which develop in these cases dissipate the energy of the pulsation and excite higher-order pulsations (that is, higher harmonics) which then decay away. The primary pulsation proceeds to grow once again, resulting in stable repeating cycles. These non-ejecting cycles can be seen at the largest heating rates, but also in cases such as 2.5×10^{45} ergs yr⁻¹ with uniform heating, when both lower and higher heating rates are seen to exhibit ejections. It is also possible for a simulation to exhibit one or more ejections before settling into a repeating non-ejecting cycle. This phenomenon, and why the switch between ejecting and non-ejecting behaviours is not a simple bifurcation, is discussed in Section 3.2.6 below. The pulsation periods of these very large amplitude pulsations can be as long as the freefall time of the material near the stellar surface, as the trajectories of this material are approximately ballistic.

In order to understand the physics behind the pulsations we observe, it is instructive to examine the shape of the pulsations in the HR diagram, examples of which are shown in Fig. 3.5. Three simulations with uniform heating at different rates can be seen in this figure (all of these simulations also appear in Fig. 3.3). On the left, we have relatively slow pulsation growth for the low heating rate of 1.3×10^{45} ergs yr⁻¹ up to the point where the pulsation reaches its maximum amplitude. These pulsations produce “circular” spirals in the HR diagram initially, but as the amplitude increases we can see the emergence of the characteristic shape seen in the other two HR diagrams as the pulsation enters the nonlinear regime. A discussion of the nonlinear effects which appear in large amplitude pulsations can be found in Section 3.2.3.

The central plot of Fig. 3.5 shows the 2.5×10^{45} ergs yr⁻¹ case, also for uniform heating, which does *not* exhibit mass ejections despite the large amplitudes attained and the emergence of strong shocks in the envelope. One complete sequence of the growth and dissipation of the primary pulsation is shown. The pulsation growth of this model is extremely rapid, reaching limiting amplitude in less than two complete pulsation cycles. The initial expansion of the model from its unperturbed state, points 1–2 in the plot, is followed by pulsation around the “equilibrium” point labelled 2. The expansion phase at maximum amplitude coincides with point 3; this expansion phase is very long due to the increase in pulsation period that occurs with increasing radius. The giant envelope then

suffers a cooling catastrophe (see below), leading to the low temperatures seen at point 4. When contraction occurs at this amplitude, the outer layers contract at supersonic speeds, and a shock is formed as the pulsation approaches minimum radius. This shock moves outward through the envelope and breaks out from the surface, producing the dramatic but short-lived spike in temperature and luminosity at point 5. Due to the radiative losses sustained during the expansion phases, the model left after the shock is not in thermal equilibrium, so it then proceeds back up its original expansion trajectory from points 6–2, after which the pulsation begins to grow again. The higher-order pulsations excited by the shock can be seen around point 6, decaying as the model expands once again.

The right hand plot in Fig. 3.5 shows the 1.7×10^{45} ergs yr⁻¹ case for uniform heating, which *does* dynamically eject material. To generate this plot, our mass-loss routine was deactivated, so the ejected layer is not excised from the model, and continues to be simulated as it expands away. The same features can be seen as in the non-ejecting case next to it, but with the luminosity peak at shock breakout followed by an ejection event. As the HR diagram tracks only properties of the model’s photosphere, cases when that surface is ejected to infinity trail off into the cold region of the diagram. The details of the properties of this surface (the “wiggles” seen in the plot) depend on the model’s outer boundary conditions, which become increasingly less applicable in the unbound regime, so are not expected to be meaningful. When our custom mass-removal routine is active, this regime is avoided.

3.2.3 Nonlinear physics in the pulsation cycle

As the pulsation amplitudes in our simulations grow, they become dominated by nonlinear effects. To investigate how important these effects are, we will examine the repeating cycle appearing in the central simulation plotted in Fig. 3.5 in more detail. This simulation (2.5×10^{45} ergs yr⁻¹ uniform heating) is shown in Fig. 3.6, which plots the evolution of the model’s surface properties during this cycle, as well as the values of the internal, gravitational, ionization and kinetic energy of the envelope. In this figure, the pulsation cycle has been divided into 5 phases:

1. **Phase I – Initial compression** – The amplitude in this phase is too low for the model to experience significant nonlinear effects; the amplitude grows exponentially in time.




Figure 3.5: Three example Hertzsprung-Russell tracks for uniform heating cases, all of which appear in Fig. 3.3. Left: a case with no ejections, plotted until maximum amplitude is attained, Centre: A case that exhibits a repeated cycle of strong envelope shocks but no ejections, plotted until after the first complete cycle. This track begins at point 1 and progresses in numerical order to point 6, whereupon it moves back to point 2 and repeats. Right: A case that undergoes a large ejection, with our mass-loss routine deactivated so the evolution of the ejecting shell’s outer surface can be seen.

2. **Phase II – Expansion** – In this phase the amplitude of the pulsation continues to increase and nears its maximum. The pulsation period lengthens as the model’s radius grows, because the dynamical (freefall) timescale of the envelope increases with radius. The increase in surface area leads to a reduction of the envelope’s radiative cooling timescale at the same time. As the star’s radius approaches its maximum, the radiative cooling timescale of the envelope actually drops below its dynamical timescale – for example, at the end of phase II in Fig. 3.6, the dynamical timescale is 646 days, but the radiative cooling timescale is only 161 days. The envelope can therefore approach thermal equilibrium (cool exponentially towards an equilibrium temperature set by its radius) faster than it can approach hydrostatic equilibrium (collapse to an equilibrium radius set by its thermal properties).

The increase in pulsation period with radius also causes the pulsation to begin to decohere, as the outer layers of the envelope, being at a larger radius, pulsate with a longer period than layers deeper within the star. This causes the inner layers of the envelope to pulsate out of phase with the outer layers, with the inner layers reaching their maximum radii and beginning to contract before the layers outside them do.

3. **Phase III – Cooling catastrophe** – During this phase, the results of the envelope’s short cooling timescale become manifest. The envelope effectively cools faster than it contracts, leading to an exponential decrease in temperature down to extremely low values, and as a result of this cooling, there is a near-complete loss of pressure support to the outer layers. This leads to an extremely fast collapse of matter which is practically in freefall, whilst having lost most of its internal energy. During this phase, almost the entire envelope ($\gtrsim 1.1 M_{\odot}$) is fully neutral, having radiated away both its thermal energy and its ionization energy (this can be seen in the bottom plot of Fig. 3.6, where the ionization energy (green) drops almost to zero in this phase). This process constitutes a cooling catastrophe, with energy transport from the inner layers of the star unable to prevent the cooling of the outer envelope to very low temperatures.

During this contraction phase, the pulsation continues to decohere, with internal layers reaching minimum radius whilst material further out is still collapsing. The phase lag between different layers of the star leads to the development of large relative velocities and to the formation of shocks.

4. **Phase IV – Shock breakout and ringdown** – As the envelope’s collapse approaches its minimum radius, its outer layers are highly supersonic, and a strong compression shock is formed, which moves outwards towards the surface as rapidly infalling layers of the envelope collide with the material interior to them which has already decelerated. The shock is effective at converting the kinetic energy of the infalling material into internal energy, helping to reheat and reionize material in the envelope which lost almost all of its internal energy during the cooling catastrophe. The shock travels towards the surface over approximately a year, as the decoherence of the different layers within the envelope leads to the inner layers reaching minimum radius long before the outer layers (in this case, approximately one year before). As the shock reaches the surface, it produces a dramatic but short-lived spike in the luminosity and temperature of the model, and excites higher-order, shorter-period pulsations: the presence of the shock near the envelope’s surface leads to the deposition of the pulsation’s kinetic energy less deep within the envelope than would have occurred if

the contraction had been coherent (that is, if all layers of the envelope had been in phase and reached their minimum radii at the same time). This deposition of energy excites higher-order pulsation modes, which are not effectively driven and therefore decay as the envelope rings down.

5. **Phase V – Quasistatic re-expansion** – The large radiative energy lost during phases II & III cause the star to regain hydrostatic equilibrium at a lower energy state than the unstable state about which the observed pulsations can grow. The star therefore expands quasistatically towards this unstable state before pulsations can begin to grow again. This expansion is analogous to the initial expansion of the model from its starting state.

The simulation used for this analysis is an example of a non-ejecting case. In the cases which do display ejections, the expansion phase typically reaches a lower maximum radius and lasts for less time. This has two effects: it prevents the internal layers of the star from decohering to as great an extent, which leads to the compression shock being stronger near the surface; and it gives the star less time to radiate away energy during the cooling catastrophe. This leads to the star retaining sufficient energy to rebound quickly and launch an ejection, rather than having all pulsation energy used up in reheating cool envelope material. Fig. 3.7 shows an equivalent analysis of pulsation phases for a simulation undergoing a mass ejection (specifically, it shows the first ejection experienced by the 1.7×10^{45} ergs yr⁻¹ uniform heating model), in which the shorter expansion phase can be seen to be accompanied with much higher total and ionization energies in the envelope during that phase.

3.2.4 Dynamical ejections

To gain greater insight into the physics of the ejection-launching process, we will examine a specific example in more detail. A shell ejection experienced by the 1.7×10^{45} ergs yr⁻¹ uniform heating model is shown in Figs 3.8, 3.9 & 3.10. This simulation was performed with our mass-loss scheme turned off so that the ejected material appears in the plots. These plots record the values of 12 output variables throughout the giant envelope during the ejection process. In these plots, we show several variables whose range of interest spans




Figure 3.6: The pulsation phases dominated by nonlinear effects in the $2.5 \times 10^{45} \text{ ergs yr}^{-1}$ uniform heating model. The first pulsation cycle experienced by this model is shown on a Hertzsprung-Russell diagram (top); we show the evolution of the model's luminosity in variable colours and effective temperature in black (second from top); below this is shown the model's radius in black and surface velocity in red (third from top); and in the bottom plot we show the internal energy (thermal + recombination) in black, ionization energy in green, kinetic energy in blue, gravitational energy in red, and total energy in purple, summed over the giant envelope between a mass coordinate of $0.43 M_{\odot}$ and the surface.




Figure 3.7: The pulsation phases that appear in the $1.7 \times 10^{45} \text{ ergs yr}^{-1}$ uniform heating model. This ejecting model has our custom mass-loss routine turned off, so that after the ejection is launched the unbound mass remains part of the model, and the surface properties refer to the outer edge of this material. The first dynamical ejection experienced by this model is shown on a Hertzsprung-Russell diagram (top); we show the evolution of the model's luminosity in variable colours and effective temperature in black (second from top); below this is shown the model's radius in black and surface velocity in red (third from top); and in the bottom plot we show the internal energy (thermal + recombination) in black, ionization energy in green, kinetic energy in blue, gravitational energy in red, and total energy in purple, summed over the giant envelope between a mass coordinate of $0.43 M_{\odot}$ and the surface. Note that as this diagram covers two pulsation periods, phases I and II appear twice.

several orders of magnitude and can be both positive and negative, so to allow them to be clearly followed we have made use of the log modulus transformation as described by John & Draper (1980), which we will refer to as the logmod function:

$$\text{logmod}(x) = \text{sgn}(x) \log_{10}(|x| + 1). \quad (3.3)$$

This function is linear for small values and logarithmic for large values of x , whilst maintaining sign and symmetry about the zero point.

The ejection seen in Figs 3.8–3.10 is typical of our results in that it is launched upon the surface breakout of a compression shock, which can be seen at approximately 64.6 years and a logarithmic radius of 2.2. The expansion following this compression is sufficiently fast to raise a shell of $\sim 0.09 M_{\odot}$ onto an escape trajectory. Material below the ejected layer expands on bound ballistic trajectories and eventually returns to the star with a supersonic fallback shock. This material becomes optically thin as it expands, a measure of which is given in Fig. 3.9 as the logarithm of one tenth of the product of local radius, opacity, and density ($\frac{1}{10}r\kappa\rho$). The loss of optical depth of ejected and almost-ejected matter is not well accommodated by our stellar model, as the treatment of heat transport we employ is only strictly applicable in optically thick regions, but as this occurs after the ejection has been launched, and because ejected material follows approximately ballistic trajectories after this point, we do not expect it to affect our results significantly. Our models retain a high degree of dynamical instability throughout the pulsation cycle as measured through the formal stability criterion (see Section 1.4), with regions above the hydrogen and first helium ionisation zones keeping $\langle \Gamma_1(m) \rangle < 4/3$ in all phases except in the immediate aftermath of a cooling catastrophe.

We can also gain insight into the energetics of the ejection from these figures: the compression shock – which can be seen most clearly in the plot of velocity divergence – leaves the envelope’s hydrogen completely ionized, and produces a region of singly ionized helium near the surface. In the shock’s aftermath, there is heavy deposition of recombination energy into the expanding material – the helium’s energy is deposited very quickly and then the hydrogen’s follows as the expansion continues. As the compression shock moves towards the surface, the kinetic energy of the pulsation’s infalling material is temporarily stored in ionization of the envelope, and then released to accelerate the expansion of the model’s




Figure 3.8: The first mass ejection displayed by a model heated uniformly throughout the convective envelope at 1.7×10^{45} ergs yr $^{-1}$ showing non-dimensionalised entropy per particle; the ratio of velocity and local escape velocity for regions where this ratio is above 1; the ratio of velocity and local sound speed for supersonic regions; and the relative proportions of ionized hydrogen and singly and doubly ionized helium. Also shown are contours containing 100% in black, and 99, 98, 95, 90, 85, 80, 75, 70, 60, 50, 40, and 30% in dashed grey, of the total mass of the model.

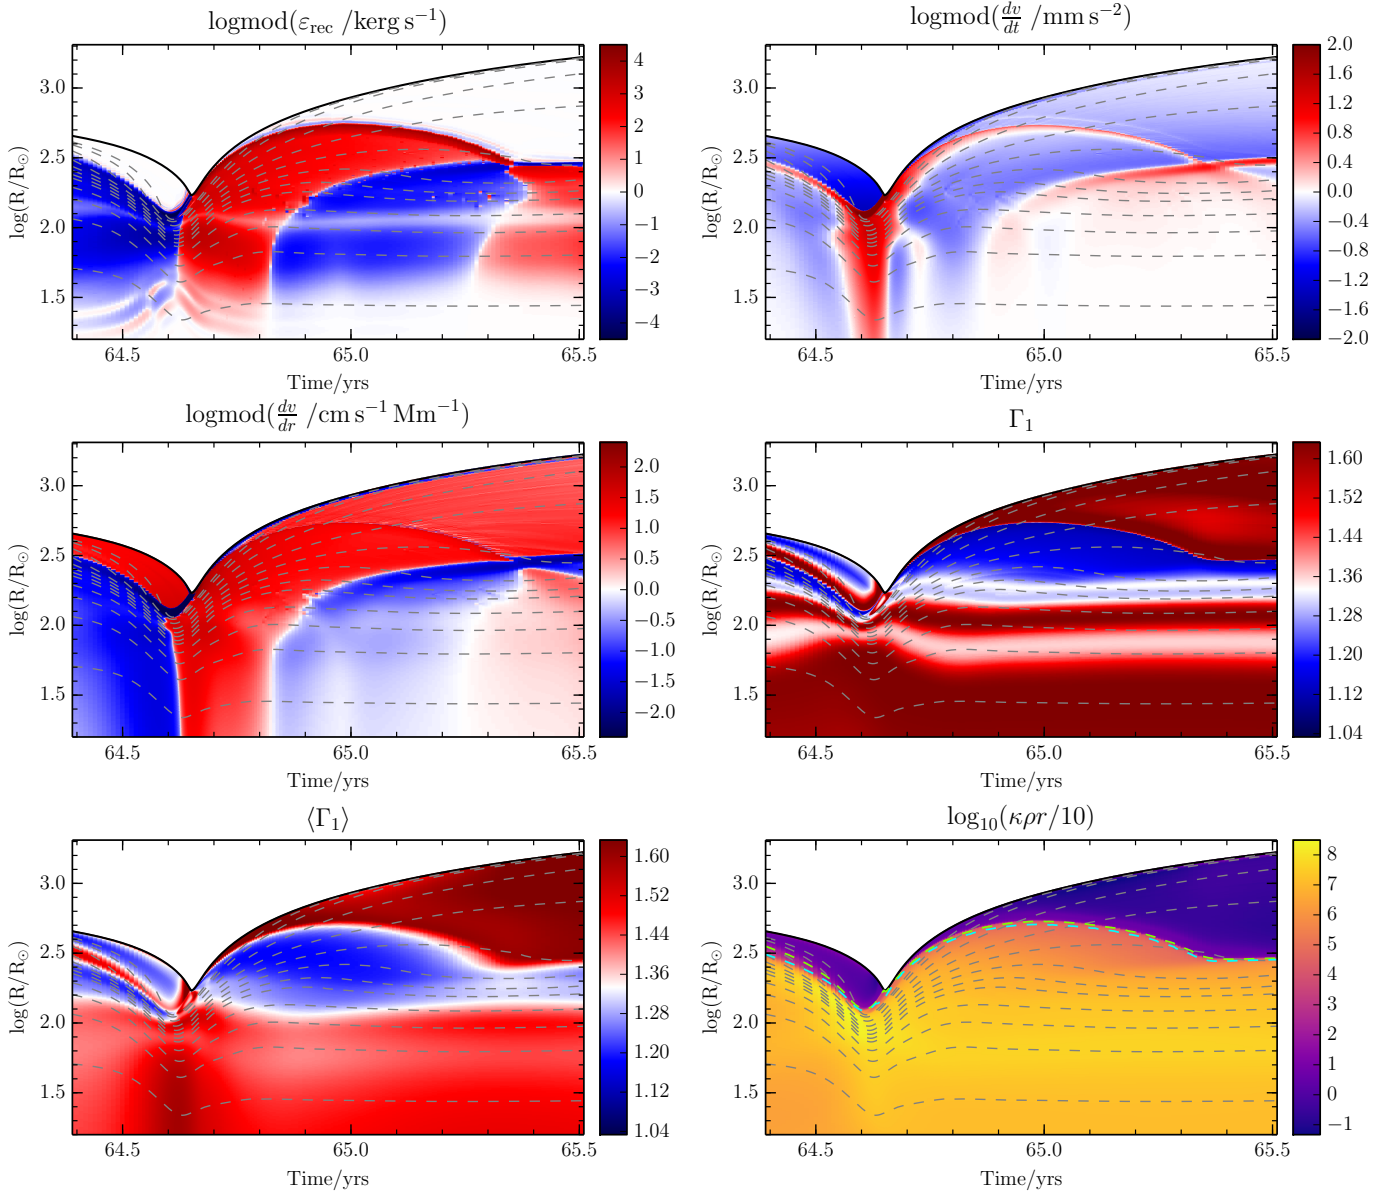


Figure 3.9: For the same event shown in Fig. 3.8 (the first mass ejection displayed by the model heated uniformly throughout the convective envelope at 1.7×10^{45} ergs yr $^{-1}$), we show the specific release rate of recombination energy; acceleration; velocity divergence; the first adiabatic exponent with white corresponding to the critical value of $4/3$; the pressure-weighted, volume averaged value of the first adiabatic exponent with white corresponding to the critical value of $4/3$; and one tenth of the product of opacity, density and local radius, a dimensionless quantity representative of the local optical thickness of the stellar material. Panel 6 also shows the radii at which the optical depth of the star is 10^2 in dashed green, and 10^4 in dashed cyan. Also shown in all panels are contours containing 100%, and 99, 98, 95, 90, 85, 80, 75, 70, 60, 50, 40, and 30% in dashed grey, of the total mass of the model.




Figure 3.10: For the same event shown in Fig. 3.8 (the first mass ejection displayed by the model heated uniformly throughout the convective envelope at 1.7×10^{45} ergs yr $^{-1}$), we show the temperature; density; velocity; sound speed; opacity; and pressure. Also shown in all panels are contours containing 100% in black, and 99, 98, 95, 90, 85, 80, 75, 70, 60, 50, 40, and 30% in dashed grey, of the total mass of the model.

outer layers in a recombination-powered bounce. The total amount of recombination energy released during the acceleration of this ejection is approximately 10^{46} ergs in half a year, the majority of which occurs at optical depths greater than 10^4 , as can be seen by comparing panels 1 and 6 of Fig. 3.9. This recombination-powered bounce is analogous to the process of "shell-triggered" ejection seen in the 3-d hydrodynamical simulations of Ivanova & Nandez (2016).

3.2.5 Non-ejecting simulations

For comparison, similar plots for the first supersonic compression of the 2.5×10^{45} ergs yr⁻¹ uniform heating model are shown in Figs 3.11, 3.12 & 3.13.

This compression exhibits a strong shock, but does not result in a shell ejection. The shock in this compression is long-lived, and can be seen between approximately 41.7 and 43.1 years. In contrast to the ejecting case shown in Figs 3.8–3.10, this model exhibits a large degree of decoherence between the inner and outer layers of the envelope, with the inner regions collapsing significantly earlier than those near the surface. This decoherence leads the shock to form very deep within the envelope and become choked by the infalling matter of the outer layers. The shock then has a much weaker effect when it finally reaches the surface, and there is no coordinated rebound of matter from all but the outermost layers of the envelope.

Another major difference that can be seen in Fig. 3.11 is that the infalling matter is almost completely neutral before it hits the shock, down to a mass coordinate of approximately $0.5 M_{\odot}$ – for comparison, the neutral layer seen in Fig. 3.8 reaches down only to a mass coordinate of $\sim 1.45 M_{\odot}$. The very thick neutral layer seen in the non-ejecting case is a result of the large radiative energy losses sustained by that model when its radius was large during the cooling catastrophe. The kinetic and gravitational energy thermalized by the shock is used to reheat and reionize this neutral material, which has the effect of damping out the primary pulsation, as the energy stored in that pulsation mode is used for this purpose; in effect, the gravitational and kinetic energy released by the envelope's collapse is used to return the cool, neutral material to a near-equilibrium state, after which additional energy is not available to drive a re-expansion. A comparison of the total envelope energies and ionization energies in our example non-ejecting and ejecting simulations can be seen in




Figure 3.11: The first catastrophic damping and non-ejecting collapse episode experienced by a model heated uniformly throughout the convective envelope at 2.5×10^{45} ergs yr $^{-1}$, showing non-dimensionalised entropy per particle; the ratio of velocity and local escape velocity for regions where this ratio is above 1 (which occurs nowhere in this plot); the ratio of velocity and local sound speed for supersonic regions; and the relative proportions of ionized hydrogen and singly and doubly ionized helium. Also shown are contours containing 100% in black, and 99, 98, 95, 90, 85, 80, 75, 70, 60, 50, 40, and 30%, of the total mass of the model.




Figure 3.12: For the same event shown in Fig. 3.11 (the first catastrophic damping and non-ejecting collapse episode experienced by a model heated uniformly throughout the convective envelope at 2.5×10^{45} ergs yr $^{-1}$), we show the specific release rate of recombination energy; acceleration; velocity divergence; the first adiabatic exponent with white corresponding to the critical value of $4/3$; the pressure-weighted, volume averaged value of the first adiabatic exponent with white corresponding to the critical value of $4/3$; and one tenth of the product of opacity, density and local radius, a dimensionless quantity representative of the local optical thickness of the stellar material. Panel 6 also shows the radii at which the optical depth of the star is 10^2 in dashed green, and 10^4 in dashed cyan. Also shown in all panels are contours containing 100% in black, and 99, 98, 95, 90, 85, 80, 75, 70, 60, 50, 40, and 30% in dashed grey, of the total mass of the model.




Figure 3.13: For the same event shown in Fig. 3.11 (the first catastrophic damping and non-ejecting collapse episode experienced by a model heated uniformly throughout the convective envelope at 2.5×10^{45} ergs yr^{-1}), we show the temperature; density; velocity; sound speed; opacity; and pressure. Also shown in all panels are contours containing 100% in black, and 99, 98, 95, 90, 85, 80, 75, 70, 60, 50, 40, and 30% in dashed grey, of the total mass of the model.

the bottom plots of Figs 3.6 & 3.7, in which the total energy of the envelope (in purple), and the ionization energy of the envelope (in green) drop almost to zero during the cooling catastrophe in the non-ejecting case (Fig. 3.6), but remain much higher in the ejecting case (Fig. 3.7).

When the compression shock in this simulation reaches the surface, it does not display the helium ionization feature seen in Fig. 3.8, is not associated with a coherent rebound by the envelope's inner layers, and results in only a small re-expansion of the upper layers, which excites a short-lived higher-frequency pulsation that then decays.

Loss of coherence between the pulsations of different layers within the envelope and the loss of large amounts of energy during a cooling catastrophe are the main physical differences seen in pulsations which lack the strong rebound required to eject a mass shell. Our interpretation is that these effects act to damp the primary pulsation and prevent mass ejections in cases when the pulsation amplitude is high but ejections are not seen. Both effects occur whilst the star is expanded to very large radii (phases II and III), and are a result of the increase in pulsation period with radius – the star's outer layers feel this effect more strongly than layers at lower radii, leading to decoherence, and a longer period means more time spent near maximum radius where radiative losses are strongest and the envelope can undergo catastrophic cooling.

It is worth noting that ejections also appear to be suppressed by the presence of incoherent pulsation modes within the envelope excited by shocks generated at previous pulsation minima. Thus the question of whether a given compression will launch a shell ejection depends heavily on the simulation's recent history. This is one way in which our simulations exhibit elements of chaos, with the future behaviour of a particular model highly sensitive to changes in initial conditions.

3.2.6 To eject or not to eject

As we can see from Fig. 3.3, whether or not dynamical ejections are exhibited by a particular model is not a simple function of heating rate, and there is no single contiguous set of heating rates which encompasses all ejecting models (note for example that the 2.5×10^{45} ergs yr⁻¹ case does not exhibit ejections, despite being surrounded by models that do). The primary effect on a simulation from varying the heating rate is a change in the growth rate of




Figure 3.14: A schematic figure showing a simplified model of phenomena seen in the large-amplitude pulsations arising in our simulations. Four invented examples of the evolution of radius with time are plotted, representing four different heating rates, and therefore pulsation growth rates, which increase to the right. Upper and lower radius thresholds are shown which represent radii that must be achieved for a model to undergo catastrophic cooling and mass ejections, respectively. In order for a star to eject mass, it must reach the lower threshold before the upper one. This gives rise to a complex structure of ejecting and non-ejecting regions in the heating rate parameter space. The major features of the simplified pulsation histories appear also on a simplified HR diagram showing the major features seen in HR diagrams of our simulation results, such as those shown in Fig. 3.5.




Figure 3.15: The surface radius (blue) and mass (red) of an extended simulation with uniform heating at 1.69×10^{45} ergs yr⁻¹, showing 18 successive ejections that together remove $\sim 51\%$ of the initial mass of the envelope.

pulsations; however, this growth is monotonically increasing with heating rate, so a more complex dependence of ejection on growth rate is required to explain these phenomena. In addition, some heating rates lead to multiple sequential ejections, as in the case of uniform heating at 1.7×10^{45} ergs yr⁻¹, whereas some lead to only one ejection followed by an apparently stable cycle of pulsation growth and shock dissipation. This behaviour would be expected to arise if the set of heating rates which leads to ejections changes with envelope mass, as seems intuitive, as these models change mass without changing heating rate.

Mass shell ejections occur when models rebound after compressions of sufficiently high amplitude (that is, compressions which are sufficiently deep), but are seen to be suppressed when these compressions display a high degree of decoherence between layers within the envelope, and when the internal energy of the envelope is very low. These two suppressing phenomena both arise as a result of nonlinear effects that emerge at very large pulsation amplitudes: very large amplitudes during the expanded phase of a pulsation lead to the decoherence of layers within the envelope, and to the radiative loss of large amounts of energy during a cooling catastrophe, and both of these effects act to prevent a strong rebound and to damp the primary pulsation. For a mass-ejection to occur, we require that the pulsation avoid being damped by these effects, but also that it have a sufficiently large amplitude at the minimum radius point in the pulsation, where the ejection is actually launched. In short, we require a pulsation whose amplitude is large, but not too large.

It is also important which phase the pulsation is in when it achieves maximum amplitude. The nonlinear damping mechanisms described above both depend on the pulsation amplitude during the *expanded* phase of the pulsation (phases II and III), whereas the ejection itself is launched in the trough of the compression (between phases III and IV). In order for ejections to occur, therefore, we require the pulsation amplitude to not be too high during the expansion phase, but also that it *is* sufficiently high during the following compression phase. Because pulsations grow until they are damped by the nonlinear effects discussed above, whether or not a given model will display mass ejections therefore depends on the interaction between the pulsation growth timescale and the period of pulsation, as it is the interplay between these two timescales that determines which phase the pulsation will be in when it reaches its amplitude peak. This means that the ranges of heating rates that will lead to ejections for a given model do not necessarily occupy a single contiguous region,

but instead form a more complex shape, in a manner which can be thought of as forming resonances between the pulsation period and the amplitude growth timescale.

A simplified “toy” model of how such a shape can emerge is described in Fig. 3.14. In this model, there are two radius thresholds, one high (outer) and one low (inner).

A star begins at an “equilibrium” radius between these two thresholds and pulsates around this initial point. The amplitude of this pulsation grows, and the star’s behaviour depends on which thresholds it reaches, and at what times: a model which passes above the outer radius threshold will undergo sufficient catastrophic cooling and decoherence of its internal layers to damp out the primary pulsation, whereas a model which passes below the inner radius threshold has a sufficiently strong compression and associated shock to launch a shell-ejecting rebound, unless the pulsation has already been damped out. Therefore, in order for a shell to be ejected, a model must pass below the inner threshold **without first passing above the outer one**. As can be seen in this figure, such a model naturally gives rise to a complex structure of ejecting and non-ejecting regions in the heating rate parameter space.

Models with certain heating rates exhibit multiple successive ejections, whilst others eject only once. A change in the mass of the model due to an ejection will likely cause a corresponding change in pulsation period, and possibly also the amplitude growth timescale. Since the behaviour of the model depends on the interplay between these two timescales, such changes are likely to cause the ranges of heating rates which produce ejections to shift, so the suppression of subsequent ejections in some models is unsurprising; the heating rate is held constant, but the regions of the parameter space may have moved such that that heating rate is now in a non-ejecting region.

An extended simulation of a heating case which exhibits repeated mass ejections is shown in Fig. 3.15, which was run for 1,500 years of star time. This simulation experiences two extended periods with no ejection events, lasting first 265 and second 692 years. Although the short-term behaviour of the pulsations we report is chaotic, and therefore unpredictable, the presence of long periods with no ejection events, followed by the later reappearance of repeated mass ejections, suggests the possibility of an underlying structural relaxation occurring within the envelope during these quiescent periods, operating on a much longer timescale than the pulsations of the model. However, a simple “by eye” analysis was unable

to discern any significant changes in internal structure (e.g. in entropy, density, ionization profiles) on top of the variation inherent in the pulsation. It should also be noted that the presence of long gaps between ejection events makes it difficult to say with certainty whether a given simulation will experience ejections in the future without calculating its evolution for an extended period of time.

3.3 Discussion

3.3.1 Numerical considerations

We have not attempted to define the boundaries of the ranges of heating rates which will cause dynamical mass ejections. This is because the exact value of the amplitude growth rate of pulsations appears somewhat sensitive to changes in the numerical parameters of our simulations. We have performed convergence testing with regard to the temporal and spatial resolution of our models, the artificial-viscosity-induced shock thickness parameter (l_η), the model’s outer boundary conditions, and the timescale of mass-removal for ejected shells; although such changes do not alter the qualitative nature of our results, our simulations exhibit chaotic behaviour, and whether or not a given pulsation will exhibit mass ejections is sensitive to the initial conditions and simulation parameters. Thus the regions of heating rate parameter space which lead to ejections, and the specific locations of the boundaries of these regions, can shift as a result of changes to the simulation parameters, and attempts to define the locations of these regions are not likely to be reliable without a large number of simulations over which “average” behaviour can be investigated (and this is before considering uncertainties introduced by our choice of model, see Section 3.3.4).

3.3.2 Dynamical instability and mass loss

Although our hydrostatic results match those of I15 closely, our hydrodynamic models evolve very differently in most cases. We see the same behaviour at the lowest heating rate of 10^{45} ergs yr⁻¹: expansion followed by pulsational stability, and at the highest, 10^{47} ergs yr⁻¹: direct dynamical ejection of the entire envelope. However, our models at the intermediate heating rates see neither the direct ejection of the base-heated hydrostatic case, nor the temporary stability of the uniformly-heated hydrostatic case. Instead, we find

that there is a large region of parameter space with insufficient heating to drive a prompt ejection that produces models which are dynamically unstable to large-amplitude pulsations.

The pulsations we recover are similar to those obtained by authors such as Wood (1974) and Tuchman et al. (1978, 1979) in the context of dynamically unstable single giants, which likewise display repeated mass-loss events in the form of dynamical shell ejections. The more recent work published by Heger et al. (1997) and Yoon & Cantiello (2010) reported the appearance of similar pulsations growing in models of RSGs, but those authors terminated their simulations before they could be affected by shocks, so could not study the shock-dominated regimes reported in this work.

The shells ejected by our models have masses of up to $\sim 0.1 M_{\odot}$. In some cases these ejections are repeated within a few decades, leading to an effective time-averaged mass-loss rate of order $10^{-3} M_{\odot} \text{ yr}^{-1}$. This is sufficiently high to clear the entire envelope of our model star within approximately 1000 years, making the shell ejections seen here a possible mechanism for the delayed ejection of common envelopes during the slow spiral-in phase. Although some of our models exhibit single, rather than repeated, ejections, we believe the suppression of further ejections in these cases to be an effect of changing the envelope’s mass, and therefore its pulsation characteristics, putting the heating rates of those models into non-ejecting regions of the parameter space. In reality, we would not expect heating rates to remain constant when such ejections occur, as in our models, so the problem of stars becoming “stuck” at non-ejecting heating rates is not likely to be relevant. We will be able to examine this possibility when we study other giant types in the following chapters, particularly in the case of AGB stars discussed in Chapter 4, as the nuclear luminosity of those stars can vary considerably even on the timescales accessible to our hydrodynamical simulations during the fast phases of the thermal pulse cycle.

3.3.3 Ejection efficiency

The situation we model in this chapter represents a highly non-adiabatic phase of the CE process, and is therefore not the situation for which the α efficiency parameter was originally conceived. However, it is still possible to define an equivalent efficiency value for the mass-ejection process that emerges from our calculations.

One common definition of the efficiency parameter is

$$\alpha \Delta E_{\text{orb}} = E_{\text{bind}} \approx \frac{GM_{\star}M_{\text{env}}}{\lambda R_{\star}}, \quad (3.4)$$

where ΔE_{orb} is the energy lost from the orbit during the CE process, E_{bind} is the gravitational binding energy of the envelope (for the purposes of this calculation, this does not include contributions from thermal energy), G is the gravitational constant, M_{\star} and R_{\star} are the mass and radius of the giant, M_{env} is the mass of the giant's envelope, and λ is a parameter commonly used to encode the dependence of the envelope's gravitational binding energy on the internal structure of the giant (see de Kool, 1990), equivalent to λ_g in Dewi & Tauris (2000).

In this way, a proportion α of the energy lost from the binary orbit is equal to the energy required to eject the envelope. Problems in calibrating values for α and λ , and in determining the correct value of E_{bind} (which depends on, for example, the precise location of the envelope/core transition, the amount of kinetic energy left in the ejected material, the amount of energy that is radiated away from the CE object and the internal energy sources available within the envelope), render this formalism somewhat inaccurate for making predictions, but it remains useful for its simplicity and ease of application.

A value of the efficiency parameter for the mass-ejection process seen in our simulations can be defined as the ratio of the binding energy that was possessed by the ejected material prior to an ejection event and the heating energy put into the system in the run-up to a that event.

An approximate expression for this might be

$$\alpha_{\text{eject}} = -\frac{m_{\text{eject}}\epsilon_{\text{grav}}}{P_{\text{heating}}\tau_{\text{eject}}}, \quad (3.5)$$

where m_{eject} is the mass of the shell that becomes unbound in an ejection event, ϵ_{grav} is the mass-averaged specific gravitational energy of the envelope, P_{heating} is the rate at which the envelope receives synthetic orbital heating, and τ_{eject} is the time between ejection events. In this case, if the entire envelope is removed by repeated shell ejections, we would have

$$\alpha \approx \frac{1}{M_{\text{env}}} \sum \alpha_{\text{eject}} m_{\text{eject}}. \quad (3.6)$$

It is not meaningful to take the average of the specific gravitational energy of the envelope

at any one point during the pulsation, as it varies a lot over any one pulsation cycle, so we will calculate it first for the initial RG model used as the starting point of our simulations, and then for an “equilibrium” model, after the envelope has expanded to approximately the radius about which pulsations oscillate.

Using an extended simulation with a heating rate of 1.7×10^{45} ergs yr⁻¹ and uniform heating, covering 500 years and averaging over 14 consecutive ejections which together removed $\approx 0.6 M_{\odot}$ of envelope material, we calculated the average efficiency of the ejection process. We find an α value of 0.25 and a λ of 0.55 when gravitational binding energy is calculated relative to the initial RG model, and an α of 0.046 and λ of 0.11 when calculated relative to an expanded radius of $518 R_{\odot}$ at 40 years. For purposes of comparison to existing literature, the first of these two definitions (comparison to the initial RG model) is closer to the one most often used by other authors.

To get an idea of how these values can vary, we performed the same calculation using a 500 year simulation with a heating rate of 3×10^{45} ergs yr⁻¹ and uniform heating, the greatest heating rate for which successive ejections were seen. This simulation yields an average α value of 0.046 and a λ of 0.55 when calculating energies relative to the initial RG model, and an α of 0.013 and λ of 0.075 when calculated relative to an expanded radius of $740 R_{\odot}$ at 20 years. This simulation ejected a total of mass of $0.27 M$ over 6 successive ejections. It should be noted that, in addition to the efficiency values varying between different simulations, the values calculated for each specific ejection *within* a given simulation also vary considerably due to variation in the time (and therefore energy input) between ejections, and the masses of the ejected shells.

The lower energy efficiency in this second simulation is a result of there being a larger average time between ejections, as well as a greater energy input rate, both of which increase the total energy input per ejection. A decrease in both average mass-loss rate and ejection efficiency should be a result of *increasing* the rate at which energy is injected into the model, which does not at first appear intuitive. The reduced ejection efficiency seen here is indicative of increased loss of energy to radiation, which is accelerated at higher heating rates by greater expansion of the envelope and higher average surface luminosities.

Both of these estimates for the ejection efficiency are considerably lower than the values commonly used to model the CE process in BPS codes, which tend to assume that the

product $\alpha\lambda$ is approximately equal to 1 (e.g. Kinugawa et al., 2014; Hurley et al., 2002; Belczynski et al., 2008). However, the values described in this section are specific to the slow spiral-in phase and therefore relate only to CE systems that enter this phase. The different ejection mechanisms that are likely to apply in these two phases make it probable that different ejection efficiencies are to be expected. A simple CE prescription for BPS codes which takes this difference into account can be imagined: one value of the efficiency parameter could be applied for CE systems that are expected to undergo envelope ejection during the fast plunge-in phase, while a second value of the efficiency parameter could be applied instead to systems expected to enter the slow spiral-in phase, utilising some (as yet unknown) criterion for determining which category a given system falls into.

3.3.4 Limitations of our simulations

The model we have employed in this work depends on several major simplifications of the CE process. Primarily, our simulations are confined to one dimension. The main problems caused by this simplification of geometry affect the modelling of the secondary's orbit, which we make no attempt to follow in this work, but we would also expect the presence of the secondary at one particular position in the envelope to have additional effects on the envelope in the forms of rotation, gravitation and localised heating, none of which we include in our simulations. Ivanova & Nandez (2016) found that the envelopes of 3-d CE models which entered the slow spiral-in phase exhibited non-negligible differences in density between polar and equatorial directions, but that these differences were considerably smaller than the up to order of magnitude differences seen during the plunge-in phase, which we are used to seeing in 3-d simulations. Despite the inability of our model to treat these asymmetries, it is likely to exhibit the correct qualitative behaviour, with quantitative differences from reality due to geometric effects. It is possible that the complete 3-d system may actually be more susceptible to the kind of mass ejections we find, as its rotation will reduce the energy requirements for mass loss in the equatorial plane. In addition, it is important to remember that the frictional heating, which is averaged across spherical mass shells in our models, is expected to be more localised in the true 3-d system; our results have proven largely insensitive to the radial location of heating within the envelope, but it is possible that localisation of this heating in 3-d has a stronger effect on the pulsational properties of

the envelope.

An important feature missing from our models is that frictional heating during the slow spiral-in phase is not expected to be constant, but to depend on the density of material at the radius of the secondary's orbit. The large-amplitude pulsations we recover are therefore likely to lead to considerable variation in heating rate. However, as the true heating rate is expected to increase with density, it should be higher during phases of envelope compression; this effect therefore represents a powerful potential *driving* mechanism for pulsations, and is likely to enhance the effects we see in our simulations. It is also worth noting that our simulations are based on heating zones defined at constant Lagrangian mass coordinates within the envelope, whereas in reality the region in which heating occurs should be better defined with radial coordinates, as it is dependent on the position of the secondary's orbit.

Another major consideration is the implementation of convection in our models using the mixing length theory, which is not properly time dependent (convective energy transport responds instantly to changes in stability under MLT, so it is time-dependent but in an unrealistic way). The lack of a true time dependent implementation of convection is a major source of uncertainty in our results, as differences in energy transport are liable to have a significant effect on the growth rates of pulsations; however, convective turnover times in our models are generally a fraction of a year, on timescales approximately two orders of magnitude below the pulsations we observe, so it is likely that the instantly-adapting convective fluxes in our simulations, while clearly not fully realistic, do form a reasonably close approximation to the correct energy transport behaviour in a way in which they would not for pulsations with significantly shorter periods. We will discuss the likely effect of different convective energy transport prescriptions on our results in more detail in Chapter 6.

3.3.5 Observational signatures

It is possible to make some predictions about what the objects modelled in this chapter would look like if observed. Most notably they would exhibit strong variability over periods in the range 3–20 years. It is possible that this variability would be complex and quasi-periodic, especially at larger amplitudes. Due to this complexity, and the large variation seen between pulsations at different heating rates, it is not possible to make detailed predictions about the light-curve shapes of these pulsations. However, the pulsations we report *are* highly

asymmetric at high amplitudes; in particular, the high-luminosity sections of the pulsations persist for much longer than the low-luminosity sections.

Our code does not have an explicit treatment of the stellar atmosphere and cannot follow the radiation-hydrodynamics of optically thin outflows, so we cannot make detailed predictions of the observable properties of these pulsations, but we can provide estimates of their general qualities using the surface properties of our models; we can estimate that the central points of these objects' pulsations would occupy a region of the HR diagram with $\log(T_{\text{eff}})$ between 3.4 and 3.5 and $\log(L/L_{\odot})$ between 4.0 and 4.4. The models which display the largest pulsation amplitudes (associated with large heating rates) have values of $\log(L/L_{\odot})$ that vary between approximately 2.5 and 5.2, with lower amplitude, non-shock-dominated models with lower heating rates saturating at a variation between approximately 3.3 and 4.6. The values of $\log(T_{\text{eff}})$ of the highest amplitude pulsators vary between approximately 3.2 and 3.8, with non-shock dominated models occupying the region between 3.5 and 3.7.

These predictions, however, are based solely on the surface properties of the envelope. As our models eject shells of matter containing a significant fraction of their total mass, it is possible that the envelope's surface will be heavily obscured by this ejected material, which may have a large effect on the observable properties of these objects. In particular, as the ejected mass shells continue to expand and cool, we would expect them to form large amounts of dust as their temperature drops below about 1500 K (Ferguson et al., 2005), which would lead to rapid increases in opacity and may obscure the central CE object entirely in the visible spectrum, rendering it observationally similar to an OH/IR star. The beginning of this cooling process can be seen in Fig. 3.7, and an increase in opacity within the ejected shell (in this case due to molecule formation) can be seen within one year of the launch of the ejection in Fig. 3.10.

It is possible that strongly episodic mass loss like we see in our simulations would leave observable relics in the material around post-CE objects. Recent analyses of planetary nebulae have noted the existence of “rings” or “arcs”, which are indicative of several episodes of strong, highly periodic mass loss from the progenitors of PNe at the ends of their lives, leading to discrete shells of matter in the nebulae (see, for example Corradi et al., 2004; Phillips et al., 2009). These rings are seen surrounding not only spherical PNe,

but also highly bipolar nebulae (e.g. Kwok et al., 1998) which are the most likely to be post-CE objects (De Marco, 2009). It is possible that these rings are indicative of strong episodic variation in the mass loss of the pre-PNe system, possibly during the CE phase. It has also been found (Corradi et al., 2014) that multiple structures in the bipolar lobes of a nebula surrounding a tight binary central star can be indicative of multiple mass-loss episodes during a CE event.

It is also possible that a similar process of episodic mass loss has occurred in the enigmatic luminous blue variable and possibly post-CE system η Carinae. The multiple epochs of mass ejection reported by Kiminki et al. (2016) require that the mechanism causing this system’s outbursts reoccur multiple times over at least 600 years. Our results, in particular that an episodic series of mass ejections can be generated by a constant rate of heating, suggest that it may be possible to explain the behaviour of this system using a similar model to the one presented here.

3.4 Summary

We have carried out stellar evolutionary calculations of a $1.6 M_{\odot}$ RG undergoing a synthetic CE event. By applying additional heating to the star’s envelope, we simulate the presence of a $0.3 M_{\odot}$ companion embedded within it during the slow spiral-in phase. We applied this heating at rates ranging from 10^{45} – 10^{47} ergs yr⁻¹, representing spiral-in times of 10–1000 years, and distributed it either throughout the convective envelope or in a thin shell at the envelope’s base. The response of the giant was modelled using the stellar evolution code MESA in both the hydrostatic regime and then using a shock hydrodynamics treatment.

Our hydrostatic results match closely those obtained using an independent code by Ivanova et al. (2015), whose initial model we reproduced. In these simulations, models heated at the base of the envelope tended to expand promptly to escape velocity, whilst those heated throughout the envelope expanded by a factor of ~ 10 in radius before destabilising. The highest and lowest heating rates examined behaved very similarly in both cases, expanding promptly to escape velocity or expanding into a non-pulsating equilibrium respectively. These results indicate the necessity of a hydrodynamical treatment for this scenario, as dynamical effects quickly become dominant in the models’ evolution.

Our hydrodynamical simulations revealed that almost the entire range of heating rates we investigated causes the expanded giant envelope to become dynamically unstable to pulsations with periods in the range 3–20 years and extremely high growth rates, whose timescales range from of order 10 pulsation periods to less than one period. These pulsations grow into a supersonic regime and develop strong compression shocks that pass outward through the outer layers of the envelope, and are eventually damped by the nonlinear effects of catastrophic cooling and internal decoherence that emerge at large pulsation amplitudes. In certain cases, the rebound following a high-amplitude compression can be strong enough to accelerate a layer of matter at the star’s surface to above escape velocity, dynamically ejecting a shell of matter from the star. This rebound is partially powered by the recombination of material ionized during the compression. The shells ejected in this way can be up to $0.1 M_{\odot}$ in mass, and ejections can repeat within a few decades, leading to a time-averaged mass-loss rate of order $10^{-3} M_{\odot} \text{ yr}^{-1}$. This mass-loss rate is sufficiently high to represent a candidate mechanism for removing the entire envelope over the duration of the slow-spiral in phase, and represents an $\alpha\lambda$ ejection efficiency in the range 0.025–0.14, with α in the range 0.046–0.25.

Chapter 4

Asymptotic giant branch stars

AGB stars have long been considered one of the most likely classes of stars to undergo dynamical mass loss. In this chapter, we report on a series of simulations of AGB stars as they evolve along the AGB and encounter dynamical instability. By coupling hydrostatic evolutionary calculations and hydrodynamical simulations of the entire star using a sampling approach, we are able to follow individual stars in detail for the entirety of their time on the AGB, including during the thermal pulses, and to study how their dynamical behaviour develops as they approach the ends of their lives, covering in detail vastly longer time periods than would be accessible to hydrodynamical simulations alone. We will report on dynamical mass-loss events that occur in these simulations, and consider the mechanism of pulsation-driven dynamical mass loss as a possible significant contributor to the mass-loss histories of AGB stars.

In Section 4.1, we describe the simulation methodology and parameters we adopted. In Section 4.2, we report on the results of those simulations and construct a parameterised model of the processes seen therein. In Section 4.3, we discuss the physics of dynamical ejection events in AGB stars and how they relate to those seen in Chapter 3. In Section 4.4, we test our model by comparing the results of evolutionary calculations incorporating it with the observationally inferred mass distribution of WDs, and assess how successful it is at reproducing this distribution from first principles. In Section 4.5, we discuss the interpretation of dynamical instability in AGBs in light of our results, the implications for the observation of stars undergoing similar dynamical events, the theoretical and observational uncertainties present in our simulations and analysis, the ability of the dynamical events we report to explain structures seen surrounding AGB stars and PNe, and the likely importance

of dynamical mass loss on the AGB.

4.1 Simulation methodology

Our investigation of dynamical mass loss on the AGB is based on models of 6 stars with initial masses of 1, 1.3, 1.6, 2, 4, and 6 M_{\odot} . We carried out hydrostatic evolutionary simulations of these stars from the zero-age main sequence to the end of the TP-AGB using the MESA stellar evolution code. To study the dynamical behaviour of these stars, we also carried out a series of hydrodynamical simulations using the hydrodynamics treatment available in MESA. Intermediate models taken from the hydrostatic evolutionary sequences were used as starting models for our hydrodynamical simulations. Because the timescales on which a hydrodynamical simulation must proceed in order to resolve the dynamical behaviour of these stars is so much shorter than their lifetime on the AGB, performing hydrodynamical simulations of the entire AGB is not feasible. We therefore ran multiple hydrodynamical simulations for each star, starting each simulation from a different point on the AGB. In this way, we are able to sample the dynamical behaviour of these stars at several points in time and to follow how their behaviour changes as they evolve along the AGB, as well as at different points during the thermal pulse cycle. A similar “hydrodynamical sampling” approach has been applied recently by, for example, Fadeyev (2017a,b) in studying the pulsation properties of Mira pulsators and by Yoon & Cantiello (2010) in the context of RSG pulsations.

Our hydrostatic simulations were performed using MESA version 8845, and began with zero-age main sequence models with an initial metallicity of $Z = 0.02$ and an initial helium mass fraction of 0.28. As discussed in Chapter 2, we adopted an Eddington grey atmosphere outer boundary condition and used the wind-loss prescription derived by Schröder & Cuntz (2005) (the Schröder wind, see Chapter 2), which we activated once the stars had left the main sequence. Mass loss on the main sequence itself was neglected. We employed the formulation of the mixing length theory of convection according to Cox & Giuli (1968) and adopted a mixing length parameter α_{MLT} of 1.9, the value which we found generated the most accurate solar model given our other simulation parameters. We employed the Schwarzschild condition for convective stability, along with a step overshooting model

extending for one quarter mixing length beyond convective boundaries; however after central helium exhaustion we move to a custom single-zone overshoot model which mixes material over one grid cell beyond the convective boundary (which causes the convective zone to adjust in size appropriately over time, see Iben, 1986). We found that this approach avoided unphysical mixing during the thermal pulses, whereas generic overshoot models were found to mix material across strong stabilizing chemical gradients which would in reality arrest the motion of overshooting convective elements.

It is well known that simulating stars on the AGB, particularly at higher masses, is numerically difficult due to the marginal stability of these stars and the presence of thin burning shells and requires careful selection of simulation parameters to ensure that calculations converge. The parameters of our hydrostatic simulations were chosen for the dual purpose of assisting in this convergence and enabling the process of using models from our hydrostatic sequences as starting points for our hydrodynamical simulations. In particular, we found that, if there are significant changes in the layout of mass zones or the choice of boundary conditions between the two simulation types, artificial oscillations emerge at the beginning of the hydrodynamical simulations as the result of this adjustment, which makes studying the emergence of pulsations difficult for obvious reasons. The compromise thus forced on us has led to difficulties in ensuring the convergence of the more massive of our hydrostatic models as these stars reach the more extreme end of the AGB, which we discuss more below.

We carried out hydrodynamical simulations using the backward-differenced hydrodynamics treatment available in MESA version 7624, which we found to be the most numerically stable version for our purposes. For each one of our hydrostatic evolutionary sequences, we ran multiple hydrodynamical simulations starting from different points along the star's evolutionary track. Individual models from the hydrostatic simulations were used as the initial models for these hydrodynamical simulations. We distributed these starting points over the stars' lifetimes on the AGB, so that we were sampling their dynamical behaviour as they aged. We also chose starting models from different phases of the thermal pulse cycle: the long hydrogen burning phase, the short helium burning flash, and the luminosity trough following the helium burning phase. This ensured that we sampled the widest possible variety of stellar parameters exhibited by these stars.

In our hydrodynamical simulations, we applied a maximum timestep between models of 0.01 years, which we found through experimental exploration to be sufficiently small to ensure that the dynamical behaviour of our models was resolved and did not change significantly with changes in timestep. Because our hydrodynamical simulations begin from models produced by purely hydrostatic calculations, our simulations start near to hydrostatic equilibrium. However, we found that there is some initial relaxation when switching to a hydrodynamical simulation, even when keeping simulation parameters as similar as possible. To remove the purely artificial oscillations caused by this relaxation, we lengthen the first 30 timesteps, which has the effect of damping out these oscillations and equilibrating the model. We adopt an artificial viscosity treatment to allow the effective modelling of hydrodynamic shocks and chose a shock spread width of 5% of local radius. We found that variation of this parameter did not affect the qualitative nature of our results. Unless otherwise stated, we ran our hydrodynamical simulations for a total of 200 years or until they encountered unrecoverable numerical difficulties. We neglected wind losses during our hydrodynamical simulations and in order to ensure that our simulations were able to survive dynamical mass ejection events we applied the unbound-mass-removal routine described in Chapter 2 and Appendix A.

4.2 Results

4.2.1 Hydrostatic models on the AGB

Our initial hydrostatic simulations were calculated from zero-age main-sequence models of masses 1, 1.3, 1.6, 2, 4, and 6 M_{\odot} . These stars' evolution along the TP-AGB is shown in Figs 4.1–4.6, beginning with the onset of the first thermal pulse and ending with the evolution of the stars off the AGB and towards the WD track. This evolution is driven by the growth of the core and by the loss of mass from the envelope due to the stellar wind. For our lowest-mass star, the wind is sufficiently strong that the entire envelope is removed within the first 6 thermal pulses, whereas with higher-mass stars more pulses occur, up to 80 in the case of our 4 M_{\odot} model. This difference can be partially attributed to increased envelope mass and partially to the shorter inter-pulse period experienced by thermally-pulsing stars with higher core masses. For the two most massive of our stars, at 4 and 6 M_{\odot} , our simulations

encounter fatal convergence failures before their envelopes have lost all their mass (see Section 4.1). However, both of these stars do reach the point of dynamical instability before their convergence failures, so for the purposes of this work, continuing these simulations further would be unnecessary and likely unphysical – it has been argued that several of the cataclysmic instabilities that arise when performing evolutionary calculations on the AGB are in fact physical instabilities rather than merely numerical ones (see, for example, Lau et al., 2012; Wood & Faulkner, 1986), and real stars may not be expected to survive past this point.

In order to examine the dynamical properties of the stars at various points during these simulations, we performed 96 hydrodynamical simulations to sample how their dynamical behaviour changed over the course of the AGB. We began these simulations using intermediate models taken from different points during our hydrostatic simulations. We chose these starting points to cover a variety of different phases during the thermal pulse cycle: during the hydrogen burning phase where the star spends most of its time, during the helium shell flash when luminosity and radius reach their maximum values, and during the luminosity trough that follows this flash, when the star is at its most compact and least luminous. We will discuss the behaviour of the hydrodynamical simulations and the physics behind the pulsations which appear in them in some detail.

4.2.2 Hydrodynamic simulations

In the majority of our hydrodynamic simulations, the evolution of our models was dominated by high amplitude fundamental mode pulsations of the envelope, which grew out of baseline numerical noise inherent in the simulations – no artificial driving or seed oscillations were applied. Each of our 6 stars began the TP-AGB in a dynamically stable state, but as they evolved pulsations eventually emerged in all 6 of them, although the time required on the TP-AGB before this occurred varied wildly, between 1,500,000 years in the case of our $2 M_{\odot}$ star and almost immediately for our $6 M_{\odot}$ star. A small subset of the results of our 96 hydrodynamical simulations is shown in Fig. 4.7. Specifically, those simulations which begin during the helium burning shell flash luminosity peaks of our $1.6 M_{\odot}$ stellar model during different thermal pulses are shown. This subset has been chosen as it contains representative examples of the three regimes of behaviour we see in our results:




Figure 4.1: The evolution of a $1 M_{\odot}$ hydrostatic stellar model on the TP-AGB. The first panel shows the radius at the photosphere in blue and the stellar mass in red; the second panel shows the specific binding energy of the envelope in blue (see text), along with a moving average of the specific binding energy averaged across a time window of approximately 1 thermal pulse period in dashed green; the third panel shows the luminosity at the surface; the fourth shows the effective temperature; and the fifth shows the helium core mass defined as the mass coordinate within which the hydrogen mass fraction drops below 10^{-7} in blue, and the carbon/oxygen core mass defined as the mass coordinate within which the helium mass fraction drops below 10^{-3} in red. Also marked in the top two panels are the models which were used as the initial models of hydrodynamical simulations. Models marked with red circles did not experience dynamical pulsations, those marked with yellow squares experienced dynamical pulsations but did not eject mass, and those marked with green triangles experienced episodes of pulsation-driven mass ejection.




Figure 4.2: The evolution of a $1.3 M_{\odot}$ hydrostatic stellar model on the TP-AGB. This figure follows the same display scheme as Fig. 4.1.




Figure 4.3: The evolution of a $1.6 M_{\odot}$ hydrostatic stellar model on the TP-AGB. This figure follows the same display scheme as Fig. 4.1.




Figure 4.4: The evolution of a $2 M_{\odot}$ hydrostatic stellar model on the TP-AGB. This figure follows the same display scheme as Fig. 4.1.

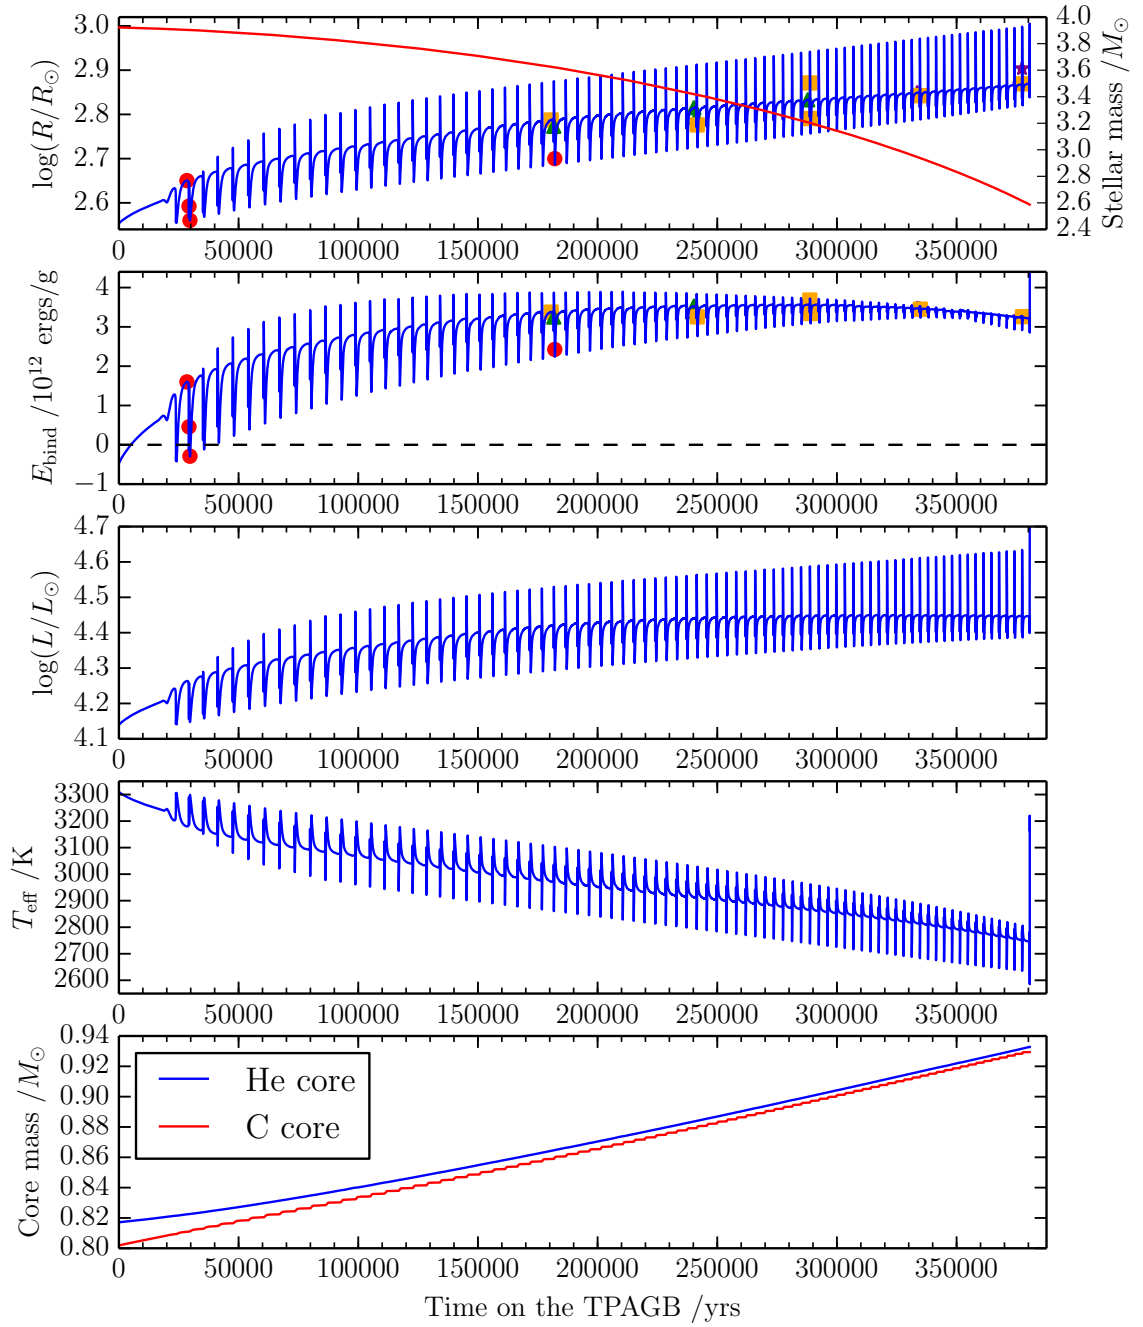


Figure 4.5: The evolution of a $4 M_{\odot}$ hydrostatic stellar model on the TP-AGB. This figure follows the same display scheme as Fig. 4.1, apart from the absence of the moving average of specific envelope binding energy from panel 2 (see Section 4.2.1). The colour scheme for hydrodynamical simulations is also expanded to include purple stars for those simulations which proved to be too numerically unstable for their properties to be reliably determined.




Figure 4.6: The evolution of a $6 M_{\odot}$ hydrostatic stellar model on the TP-AGB. This figure follows the same display scheme as Fig. 4.5.




Figure 4.7: Each panel shows the evolution of photospheric radius (in blue) and total mass (in red) of a $1.6 M_{\odot}$ star during hydrodynamical simulations started from a different point in the star’s life on the AGB – during the 7th, 10th, 12th, 14th, and 17th thermal pulses. The sudden changes in mass occur when unbound shells of matter are excised from the simulation by our mass removal algorithm (see Section 2.3) and indicate dynamical mass ejection events. Each simulation is begun from a point near to the luminosity peak during the rapid helium burning phase of the thermal pulse. These simulations were selected because they display the three regimes of behaviour we observe – no pulsation, self-limiting pulsation without mass ejection, and pulsation with mass ejection.




Figure 4.8: A Hertzsprung-Russell diagram showing the evolution of surface luminosity and effective temperature for a hydrodynamical simulation of our $1.6 M_{\odot}$ stellar model, beginning from a point near the helium burning luminosity peak of the 12th thermal pulse. The inset shows a magnified view of the centre of the spiral, where the pulsation first emerges.

1. **Non-pulsating** – In this regime, displayed by stellar models which are the most stable, no pulsations are seen. This regime is displayed in the simulation begun from pulse number 7 (panel 1) in Fig. 4.7. It is unlikely that these stars do not pulsate at all in reality, as stars in this regime are expected to be long-period variables, rather they do not have sufficiently strongly-driven pulsations to appear in our simulations, which are likely to experience a certain amount of numerical damping due to their backwards-differenced nature (see, e.g. Appenzeller, 1970).
2. **Self-limiting** – In this second regime, which can be seen in the simulations starting from pulses 10 and 12 (panels 2 and 3) in Fig. 4.7, our models experience high-amplitude dynamical pulsations. However, these pulsations are self-limiting and do not cause mass ejections. The pulsations in this regime grow exponentially until they reach their limiting amplitude, whereupon they undergo nonlinear damping and their amplitude is reduced. In some cases, these simulations exhibit a repeating cycle of simulation growth and damping, although they can also enter a more stable pattern, with a smaller amplitude that only varies by a small amount – a process which appears similar in nature to that observed by Ya’ari & Tuchman (1996). The pulsations we observe exhibit strong nonlinear damping at high amplitudes, associated with the emergence of strong shocks in the envelope, that are not significant at lower amplitudes, and it is these effects which limit the pulsation amplitude. We shall discuss the physics behind these nonlinear damping effects in detail in Section 4.3.
3. **Ejecting** – The third regime, in which we are the most interested, is that in which dynamical mass ejections occur. Examples of this behaviour can be seen in the simulations started from models during pulses 14 and 17 (panels 4 and 5) in Fig. 4.7. Dynamical pulsations in these models are driven sufficiently strongly that, when they have attained the necessary amplitude, they accelerate shells of material at the surface of the model to above the local escape velocity. These ejected shells are completely unbound and represent a form of episodic dynamical mass loss from the star. In order to allow our simulations to continue after such an ejection event occurs, we excise these ejected layers once they exceed escape velocity. It is the removal of these layers which causes the sudden decreases of the star’s mass that can be seen in Fig. 4.7.

These dynamical mass ejections can each remove a mass of up to approximately $0.1 M_{\odot}$, and generally repeat on timescales of between 10 and 100 years. The time-averaged mass-loss rate due to these repeated ejections is between approximately 10^{-4} and $1.5 \times 10^{-3} M_{\odot} \text{ yr}^{-1}$. The process by which mass shells are accelerated to ejection velocities is associated with the effect of recombination following strong shocks in the envelope, and is discussed in Section 4.3.

The time between successive mass ejections is dependent on the growth rate of the pulsations, as it is necessary for there to be enough time for the pulsation to regain sufficient energy and violence to accelerate another set of material to escape speed, a time interval determined by the pulsation's growth timescale. The cycle of amplitude growth and modulation seen in the non-ejecting regime also follows this pattern – a large amplitude pulsation does not necessarily result in an ejection, but can instead experience strong nonlinear damping resulting in the loss of significant energy from the pulsation mode, necessitating a further period of amplitude growth. This leads to considerable variability in the period between successive ejection events; the simulation shown in panel 4 of Fig. 4.7 displays this behaviour for three growth and damping cycles at an age of approximately 150 years. It is also possible for simulations which display repeated mass ejections to settle down into steady repeating patterns after a number of ejections have occurred, behaviour which can be seen in the last 50 years of the simulation shown in panel 5 of Fig. 4.7. Such stable periods may later be followed by the resumption of ejecting behaviour after multiple centuries and dozens of pulsation growth/modulation cycles.

The evolution of our simulations is highly chaotic (as expected for pulsators of this type, see Icke et al. 1992), due to the history dependence of the pulsation process and the nonlinear physics at work, so the exact course of evolution of a given hydrodynamical simulation cannot be considered reliable. Rather, each calculation should be considered to be one of a set of qualitatively similar possible outcomes. Small variations in initial parameters, while they will change the exact evolution of the simulation, do not affect its qualitative nature (for example: whether or not a given model displays mass ejections).

The overdriven regime seen in our simulations of CE objects in Chapter 3, in which pulsations grow extremely quickly to high amplitudes but no mass ejections occur, is not seen in these stars. It is possible this is because that regime only occurs at luminosities so

high our AGB models do not encounter them. The lowest CE heating value which led to this behaviour corresponds to a value of $\log(L/L_{\odot})$ of approximately 4.2, which is greater than the luminosity of any of our AGB models with initial mass less than $2 M_{\odot}$ in any of their hydrodynamical models, and the L/M ratio (which may be the more physically-relevant quantity) of those CE models exceeds that seen in any of our hydrodynamical simulations on the AGB. Although our more massive stars do exceed this luminosity, it seems unlikely that the point at which this regime emerges would be independent of stellar mass. It is also possible that we do not encounter this instability because, as suggested in Chapter 3, the luminosity of our AGB models change during our hydrodynamical simulation time periods, sometimes highly significantly. The overdriven non-ejecting regime displays a high level of regularity, and it is possible that variations in luminosity will prevent such a steady pattern of repeating behaviour from emerging, especially in the light of the chaotic nature of our simulations.

In Fig. 4.8, we present an HR diagram showing the dynamical pulsation that develops in a hydrodynamical simulation of our $1.6 M_{\odot}$ model beginning from during the helium shell flash luminosity peak of the 12th thermal pulse. The growth of the pulsation is shown until it reaches maximum amplitude. Features visible once the pulsation attains high amplitude include a rapid nonlinear cooling excursion, and the sharp luminosity and temperature spike revealing the breakout of a shock from the envelope, processes which are both discussed in Section 4.3.

The starting points of our hydrodynamic simulations, which are marked on Figs 4.1–4.6, are coloured according to which behavioural regime they fall into. It is clear that there are patterns in which regimes emerge at what times: simulations are more likely to exhibit mass ejections later along the AGB, at larger luminosities, at higher radii, and at lower temperatures. In particular, it is evident that ejections are far more likely to occur during the helium shell flashes than at other phases of the thermal pulse cycle, and it is during this phase that dynamical pulsations and mass ejections emerge first. For stars sufficiently far along the AGB, mass ejections can also be observed during the hydrogen burning phase, and in only one simulation out of the 96 that make up this dataset can mass ejections be seen during the luminosity trough that follows the helium burning phase.

Although the time-averaged mass-loss rate of the dynamical ejection process is very

high (of order $10^{-3} M_{\odot} \text{ yr}^{-1}$), the helium burning peak is also very short (of the order of 100s of years). The mass loss experienced by our models is therefore sufficiently fast that the entire envelope could be removed during the shell flash phases of a handful of thermal pulses. In Section 4.2.3 below, we construct a parameterised model of the dynamical mass loss we observe and apply it to study the effect of dynamical mass loss on the evolutionary history of AGBs.

Han et al. (1994) hypothesised that dynamical mass loss begins when the total energy of the envelope becomes positive (i.e. it becomes energetically unbound). To examine this hypothesis, we have calculated the specific envelope binding energy in our models (summing gravitational potential energy and internal energy, including recombination energy) averaged over the entire envelope – the values can be seen in the second panels of Figs 4.1–4.6. The value one obtains for the binding energy of the envelope can depend strongly on the precise location chosen for the envelope’s lower boundary, and we have followed Han et al. in choosing the location of this boundary to be at the point where this dependence is most weak. In their paper, Han et al. used a specialised stellar-evolution code which did not resolve the thermal pulse cycle, but instead produced a kind of “averaged” evolution. In our simulations, the binding energy of the envelope varies dramatically during a thermal pulse, so to enable comparison with the results of Han et al. we have applied a moving average filter calculated over a window with width approximately equal to the thermal pulse period. It is difficult to do this precisely, as this period varies with the mass of the core, decreasing approximately exponentially as the core mass increases (Paczynski, 1975).

In our simulations, the time-averaged value of the envelope’s binding energy increases as stars progress along the AGB until the final thermal pulses, at which point it may again decrease. In general, the increase in total energy is associated with the emergence of pulsations and then dynamical mass ejections, but it is not clear that crossing the zero-point specifically has any particular relevance in our models, although it is difficult to conclude that it *does not*, due to the difficulty in identifying the absolute value of the binding energy. In addition, the two most massive of our simulations attain positive envelope energies almost immediately at the start of the TP-AGB, whereas we find that dynamical instabilities emerge much later than this. It is possible that stars of higher mass are less well described by the hypothesis of Han et al. and, as we shall discuss in Section 4.4, it is for these stars for

which comparisons to observation based on the calculation of WD mass distributions such as those performed by Han et al. are the least sensitive.

In order to test the validity of our “sampling” approach to determining how the dynamical behaviour of stars varies along the AGB, we performed a small number of extended hydrodynamical simulations which covered a period of several thousand years of star time, beginning during the hydrogen burning phase of the thermal pulse and running throughout the entire helium shell flash. The results of one such simulation can be seen in Fig. 4.9. As can be seen in that figure, the behaviour of the star is similar to what would be expected from our shorter simulations. Initially, the hydrogen burning phase displays self-limiting pulsations with no mass ejections, with alternating amplitude growth and damping periods. Once the helium burning phase begins, approximately 1,700 years after the start of the simulation, pulsations cease as the star contracts, then emerge again as the star’s luminosity begins to peak during the helium shell flash. Three mass ejections occur during this peak, before the model once again enters a non-ejecting regime. The star eventually ceases to pulsate altogether as it enters the luminosity trough that follows helium-burning extinction. This simulation and similar calculations we performed broadly support the validity of our sampling approach.

4.2.3 Modelling nonlinear pulsations

Of the 96 hydrodynamical simulations we performed, 53 display dynamical pulsations. We measured the growth rate of these pulsations by modelling their amplitudes as growing exponentially in time, as is predicted by theory (see Cox, 1980; Unno et al., 1989; Kippenhahn et al., 2012). This exponential growth is eventually arrested by the emergence of nonlinear damping effects (see Section 4.3), so we extracted the pulsation growth rate from the beginnings of the simulations only, before the amplitude growth rate deviates significantly from exponential. Because the pulsation period varies with amplitude, it was not feasible to fit exponentially modulated sine waves directly to the radius or surface velocity evolution of our models, and we instead fit the local maxima and minima of the pulsations simultaneously as two diverging exponentials with the same growth rate, effectively fitting the amplitude envelope of the pulsation. In addition, in some of our simulations the underlying stellar parameters vary due to the thermal pulse cycle on sufficiently short timescales that

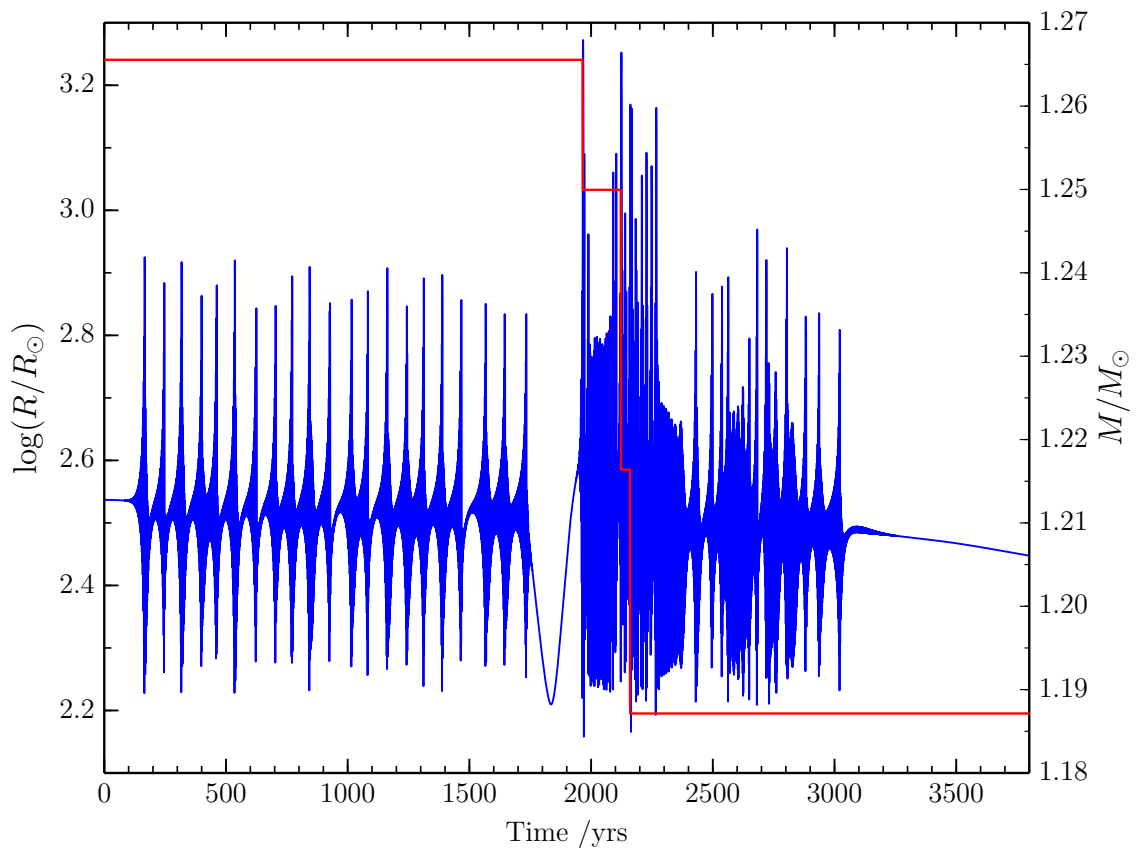


Figure 4.9: The evolution of radius (blue) and stellar mass (red) in a hydrodynamical simulation beginning approximately 2000 years before the helium luminosity peak of the 13th thermal pulse experienced by our $1.6 M_{\odot}$ model. The high-luminosity phase due to the helium shell flash occurs between times of approximately 1900 and 2300 years.

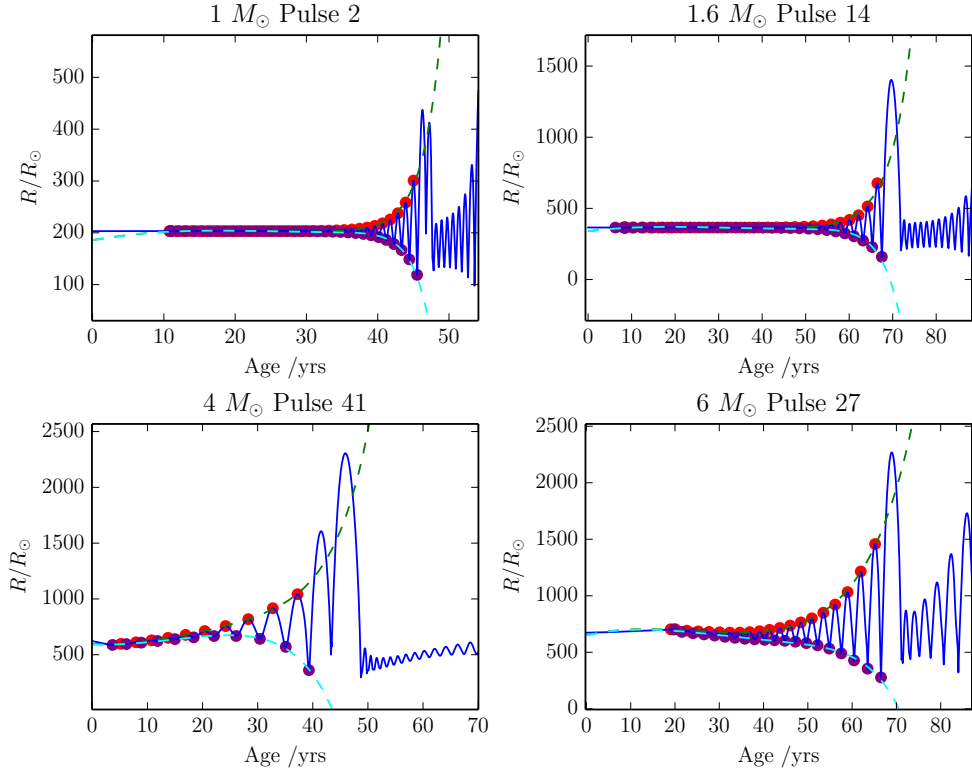


Figure 4.10: Four examples of diverging exponential fits to the initial pulsation growth phases of hydrodynamical simulations. The two top panels show simulations started during the hydrogen burning phase, and the bottom two show simulations started during the helium burning shell flash. The blue line shows the radius of the photosphere, the red and purple points show the maxima and minima used for the fitting process, and the green and cyan dashed lines show the diverging exponentials which are fitted to those extrema.

the “baseline” equilibrium value of the stellar radius cannot be treated as constant for this calculation, so we also fit this baseline value with a cubic polynomial. Our fitting function for stellar radius R as a function of time t was therefore

$$R = \alpha_1 + \alpha_2 t + \alpha_3 t^2 + \alpha_4 t^3 \pm \alpha_5 e^{\alpha_{GR} t}. \quad (4.1)$$

Four examples of the fits we obtain can be seen in Fig. 4.10. This method allowed us to extract the exponential pulsation growth rate α_{GR} from 50 of our simulations, with the remaining 3 pulsating models being too unstable to characterise in this way (they do not have enough local extrema during initial amplitude growth to allow us to perform a meaningful fit).

We find that the best-fitting relationship between pulsation growth rate and stellar mass

M , surface luminosity L and effective temperature T_{eff} is

$$\alpha_{\text{GR}} = 1.21 \times 10^{-6} / \text{yr}^{-1} \left(\frac{M}{M_{\odot}} \right)^{-0.55} \left(\frac{L}{L_{\odot}} \right)^{0.33} (T_{\text{eff}} / \text{K})^{1.4} - 1.11 / \text{yr}^{-1}. \quad (4.2)$$

We measure goodness-of-fit with the coefficient of determination R_{CD}^2 , which we define as one minus the ratio of the residual sum of squares to the total sum of squares,

$$R_{\text{CD}}^2 \equiv 1 - \frac{\sum (y_i - f_i)^2}{\sum (y_i - \bar{y})^2}, \quad (4.3)$$

where y_i are the data values with mean \bar{y} , and f_i are the model's predictions for each datapoint. Values of R_{CD}^2 for models fits are confined to the interval between 0 and 1, with higher values indicative of stronger fits.

The coefficient of determination for the above fit is 0.965. The result of this fit is shown in panel 1 of Fig. 4.11. Although it appears surprising that amplitude growth rate should have a positive dependence on effective temperature, as cooler stars are expected to be more expanded and therefore more unstable, this model takes into account the differences between stars of different masses, not just the evolution of individual stars along the TP-AGB, and that the stellar mass and luminosity are coupled in non-trivial ways; when the dependence of amplitude growth rate on effective temperature is looked at in isolation, we find that for any given *individual* star, as it progresses along the AGB, the pulsation growth rate does increase as temperature decreases.

We also constructed a parameterised model of the time-averaged mass-loss rate due to dynamical mass ejections. We found that our two most massive stars, $4 M_{\odot}$ and $6 M_{\odot}$, were not easily modelled in this way and have been excluded from this analysis. For comparison, a growth rate model fit using only the remaining 28 simulations of our 4 lowest-mass stars can be seen in panel 2 of Fig. 4.11. We extracted the mass-loss rate from our simulations via the rather crude method of dividing total mass loss by total simulation time. This ignores the fact that our simulations begin from models close to hydrostatic equilibrium, in which the pulsation amplitude is zero, so there is a period occupied only by pulsation growth, where in reality one would expect pulsations to already be present. This initial growth phase means that our mass-loss rate measurements are likely to be underestimates.

We found that the best-fitting model for mass-loss rate was

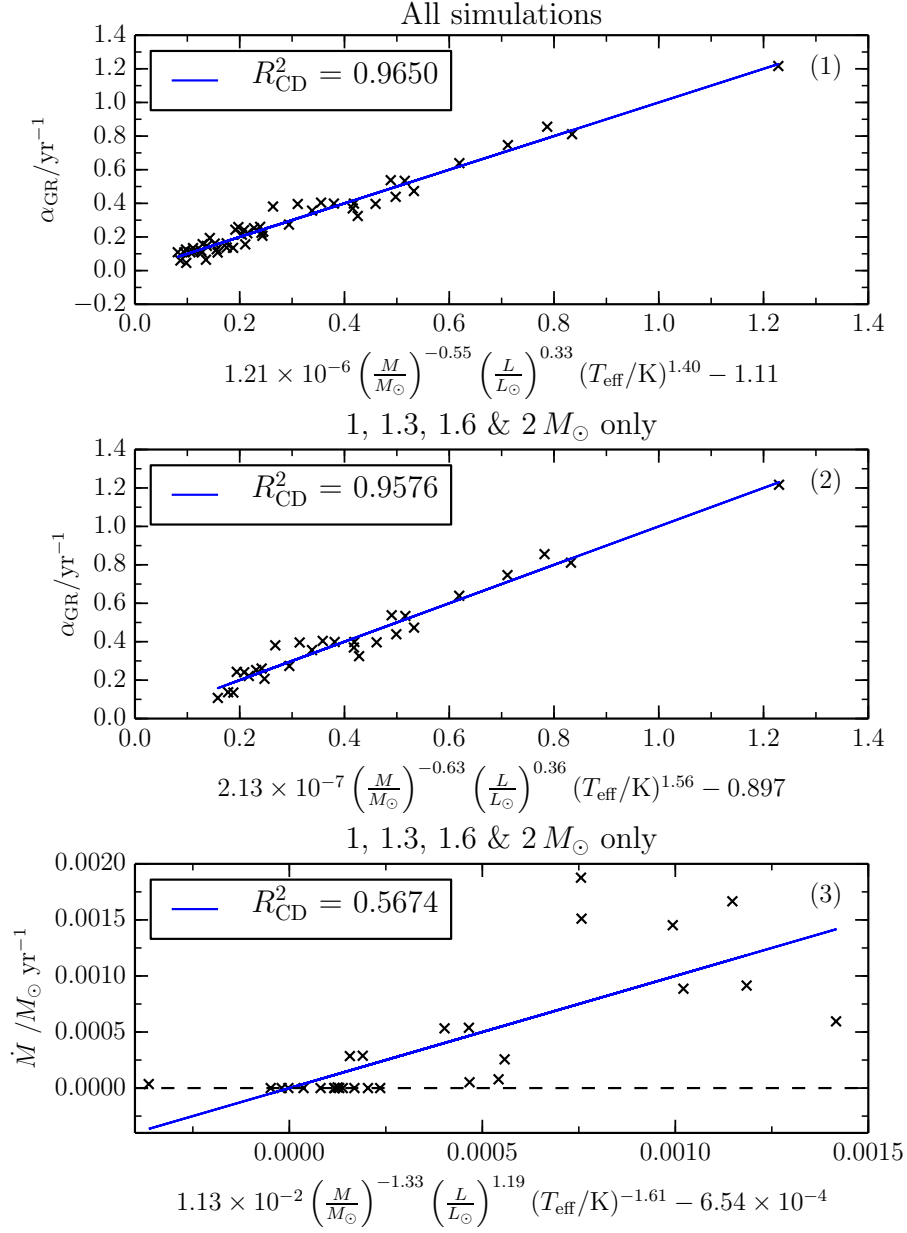


Figure 4.11: Panels 1 and 2 show our best-fitting models for initial pulsation growth rate α_{GR} as a function of stellar mass, surface luminosity, and effective temperature. In panel 1, all 50 pulsating simulations for which growth rates could be obtained are used for the fit, in panel 2, only the 28 simulations from stars with initial masses 1, 1.3, 1.6 & 2 M_{\odot} are used. Panel 3 shows our best fitting model for time-averaged mass-loss rate using the 28 simulations from stars of masses 1, 1.3, 1.6 & 2 M_{\odot} . In each case, the x-axis is the best-fitting model, and the y-axis is the true value. Coefficients of determination are shown inset in all panels.

$$\dot{M} = 1.13 \times 10^{-2} / M_{\odot} \text{ yr}^{-1} \left(\frac{M}{M_{\odot}} \right)^{-1.33} \left(\frac{L}{L_{\odot}} \right)^{1.19} (T_{\text{eff}} / \text{K})^{-1.61} - 6.54 \times 10^{-4} / M_{\odot} \text{ yr}^{-1}, \quad (4.4)$$

where mass-loss rate \dot{M} is in units of solar masses per year, with coefficient of determination 0.567. This model fit is shown in panel 3 of Fig. 4.11. The mass-loss measurements on which this fit is based carry significant uncertainty due to the crude nature of our measurement method and due to the chaotic nature of the system – how many ejections occur over the course of a given calculation can vary due to small changes in initial conditions. In contrast, our growth rate measurements, because they do not depend on the “messy” evolution of the models after limiting amplitude is reached, are largely unaffected by this chaos, and so have much lower scatter. Surprisingly, we found that models which treated mass-loss rate as a function of envelope mass rather than total stellar mass produced lower quality fits, and models which treated envelope and core mass separately, or envelope and total mass separately, did not increase the quality of fit sufficiently to justify the additional complexity of the fitting function. We do note, however, that luminosity is a good proxy for core mass (although the relationship between these two quantities does vary over the thermal pulse cycle), so an explicit inclusion of both core and envelope mass is not necessary to encode this information in the model.

It is important to note that we find a large number of simulations in which our best-fitting model suggests that small amounts of mass loss should be seen in the simulation but in fact none occurs. These are the models in which dynamical pulsations appear, but no mass ejections are seen. It is reasonable to conclude that there is a threshold value of “predicted” \dot{M} below which mass ejections do not occur, as pulsations are driven insufficiently strongly, and their amplitudes are limited by nonlinear damping effects before ejections can take place.

We considered the possibility that a source of scatter in our mass loss-model may be that the mechanism has some dependence on the mass-loss history of the star; the fact that our underlying hydrostatic AGB simulations do not incorporate this mass-loss mechanism may therefore represent an inconsistency in our approach. To test this, we ran additional hydrostatic evolutionary simulations using the model described above as a custom mass-loss prescription in addition to or in replacement of the Schröder wind. We then carried out new

sets of hydrodynamical simulations based on these new sequences of hydrostatic models. When performing these second-generation calculations, we applied mass loss according to the above formula only when it exceeded some threshold value, for which we found the most reasonable values to be between 1.2×10^{-4} and 2.4×10^{-4} . In practice, however, we did not find that this approach improved the quality of the mass-loss model we obtained, despite the fact that the mass loss treatment used in the hydrostatic models was more consistent with the hydrodynamic simulations.

This second set of hydrostatic simulations was useful, however, in studying the mass-loss histories that our dynamical mass-loss model predicts for stars on the AGB. We found that dynamical mass loss begins, in all cases, during the luminosity peaks of the helium shell flashes, with 1–3 thermal pulse cycles often being sufficient to remove the entire remaining envelope. Eventually, stars sufficiently advanced along the AGB become unstable enough that they exhibit dynamical mass ejections during the hydrogen burning phase also. We found that, in those cases where our model predicted dynamical mass loss should occur during the hydrogen burning phase, it was predicted to be of sufficient strength to quickly remove the remainder of the envelope. An example simulation is shown in Fig. 4.12, where the effect of our dynamical mass-loss model is visible during two helium shell flashes, followed by a final episode of mass loss during hydrogen burning which removes the remainder of the envelope.

4.3 Physical analysis

4.3.1 The pulsation cycle

To gain insight into the physical processes governing dynamical mass ejections in the AGB regime, we repeated the hydrodynamical simulation beginning at the luminosity peak of the 14th thermal pulse experienced by our $1.6 M_{\odot}$ star with the mass removal scheme deactivated so we can see the ejected material. The evolution of several quantities that are important during the first ejection event that this model undergoes is shown in Fig. 4.13. For the purpose of discussion, we have divided the evolution shown in that plot into 3 phases. These phases are similar to those reported in Chapter 3, and although the differences in stellar regime are sufficiently important that an independent analysis is appropriate, we have




Figure 4.12: The final 4 thermal pulses exhibited by a $1.3 M_\odot$ star undergoing mass loss according to the model described in Section 4.2.3 with a threshold value of $1.2 \times 10^{-4} M_\odot \text{ yr}^{-1}$, in addition to the Schröder wind. The top panel shows the time evolution of stellar radius in blue and mass in red whilst the helium core mass is plotted in dashed black; the central panel shows the time evolution of surface luminosity; and the bottom panel shows the evolution of effective temperature.




Figure 4.13: The first mass ejection event experienced by a hydrodynamical simulation of the helium luminosity peak of the 14th thermal pulse undergone by our $1.6 M_{\odot}$ model. We show an HR diagram (panel 1) coloured by phase (see text); the evolution of surface luminosity coloured by phase and effective temperature in black (panel 2); the surface radius in black and surface radial velocity in red (panel 3); the internal energy (thermal and recombination) of the envelope in black, recombination energy in green, kinetic energy in blue, gravitational energy in red and total energy in blue (panel 4); and the thermal timescale of the star in black and the dynamical timescale in red (panel 5). For the simulation shown in this plot, our custom mass-loss scheme was disabled, so that ejected material is still present in the simulation after the ejection event.

simplified the discourse in light of this fact. Note however that the phases we have selected in this analysis are not identical to those used in Chapter 3, as here we will divide the process into fewer phases overall.

- **Phase I – Exponential growth** – Exponential amplitude growth occupies most of this phase, beginning to diverge strongly from exponential towards the end as nonlinear effects emerge. These effects are due principally to the fact that the star’s dynamical and thermal timescales both vary with stellar radius. A higher radius increases the rate of radiative energy loss from the stellar surface, decreasing the thermal timescale, and at the same time the dynamical timescale increases due to the decrease in the gravitational force experienced by the outer layers of the star. When the pulsation amplitude becomes large enough, the radius is sufficiently high during the expanded phase of the pulsation cycle that, as can be seen in Fig. 4.13, the thermal and dynamical timescales become approximately equal, and in fact the thermal timescale drops below the dynamical timescale temporarily. The pulsation period also increases with the dynamical timescale. This means that, for sufficiently large amplitudes, the pulsation begins to internally decohere, with different mass layers of the star reaching maximum radius at different times. The contraction of the star therefore begins deep within the envelope, with the outer layers contracting last.
- **Phase II – Cooling catastrophe** – In this phase, the envelope cools and contracts, but because its thermal timescale is shorter than its dynamical timescale, the envelope is not able to contract fast enough to keep up with the rate of cooling – it equilibrates thermally faster than hydrodynamically – and the outer layers of the envelope experience strong thermal losses, dropping to very low temperatures. This event, with the star unable to adjust to the rate of radiative losses, constitutes a cooling catastrophe. The dramatic loss of heat leads to a loss of pressure support, which causes the envelope’s outer layers to contract extremely rapidly, effectively in freefall. As this occurs, the pulsation continues to decohere; layers deeper in the envelope reach their minimum radius and cease contracting before the layers above them. This causes a shock to develop, as the outer layers of the star fall onto approximately stationary inner layers, with relative velocity between layers reaching several times the local

sound speed.

- **Phase III – Shock breakout** – As the star approaches its minimum radius, the infalling of the envelope’s outer layers is highly supersonic, and the compression shock this causes moves outward towards the surface. This shock acts to convert the kinetic energy of the infalling material into internal energy, reheating and reionizing material which lost much of its energy during the cooling catastrophe. When the shock reaches the star’s surface, it causes a dramatic peak in luminosity and temperature, and the energy freed by the shock drives the rapid outward acceleration of the envelope’s outer layers as the star begins to expand again. It is this acceleration which pushes a region near the surface to above escape velocity. By this point the total energy of the envelope is significantly *lower* than it was in phase I, but the compression shock has the effect of concentrating more energy into the shocked layers near the surface, enabling that material to become unbound.

Once the unbound outer layer expands significantly, the surface boundary condition of our simulation begins to be inappropriate, which causes the strange fluctuations in surface temperature and luminosity seen towards the end of Fig. 4.13. By excising unbound material from our simulations, our custom mass scheme (deactivated in this example) prevents this problem from emerging and affecting our results.

4.3.2 Dynamical ejection

We shall briefly examine the mass ejection process in more detail. For the same simulation discussed above (and shown in Fig. 4.13), the period of time surrounding the mass ejection event is shown in detail in Figs 4.14, 4.15, & 4.16, showing the evolution of physical quantities throughout the envelope. Visible in those figures is the end of the cooling catastrophe, the emergence and breakout of the succeeding shock, and the acceleration of a layer $0.081 M_{\odot}$ in size to above escape velocity.

The compression shock is visible in these figures between approximately 36.7 and 36.8 years, between $\log(R)$ values of approximately 2.1 and 2.4. The source of the shock can be easily seen in panel 3 of Fig. 4.14, in which material from the outer layers of the envelope is seen entering the shock region at over 10 times the local sound speed. Shortly after this shock

reaches the surface, material near the surface of the star begins to exceed escape velocity moving outward. We draw the reader’s attention to the features seen in the ionization fraction plots in Fig. 4.14, which show the dramatic effect the shock has on the ionization of the envelope: causing the hydrogen partial-ionization zone to reach within a tenth of a solar radius of the surface and producing an isolated pocket of singly-ionized helium. By comparing these features to the recombination energy release rate, acceleration, and velocity divergence values plotted in Fig. 4.15, we can see that the recombination of the material which the shock had ionized plays an important role in accelerating the ejected material, producing a “recombination-powered bounce” by depositing energy heavily in those regions in which acceleration occurs. Within 1 year of shock breakout, over 7×10^{45} erg of energy is released by recombination of hydrogen and helium in the envelope. Material which is accelerated out to high velocities but less than the escape velocity follows approximately ballistic trajectories and falls back onto the star causing a weak subsidiary fallback shock at times too late to be seen in these figures.

As can be seen in Fig. 4.15, there is a significant region of the envelope in which $\langle \Gamma_1(m) \rangle$ is less than $4/3$, extending deep within the envelope. This unstable layer persists throughout the entire pulsation cycle, and at its greatest extent (approximately 1 year before the region shown in Figs 4.14–4.15), it extends right down to the core/envelope boundary, indicating that the entire envelope is dynamically unstable.

It is important to consider that as unbound material expands, it becomes optically thin, which our calculation is not equipped to deal with as it cannot model non-local radiative transfer of heat. In panel 6 of Fig. 4.16, we have plotted a measure of local optical thickness, the logarithm of the product of one tenth of local radius, opacity, and density ($\frac{1}{10}r\rho\kappa$), along with contours of optical depth. Importantly, the majority of the acceleration of the unbound layer occurs *before* that material becomes optically thin, so we do not expect our results to be significantly affected by this issue.

4.3.3 Self-limiting pulsations

Also of interest to us is why some of our hydrodynamical simulations display mass ejections while others pulsate strongly and yet do not eject mass. The nonlinear effects that become important when pulsations expand to extreme radii – catastrophic cooling and decoherence

– both have damping effects. The cooling catastrophe leads to the loss of a significant fraction of the total energy of the envelope being lost to radiation; as we can see from Fig. 4.13, the total energy of the envelope in our example simulation decreases by a factor of 2 over the course of the mass ejection event, similar losses are experienced when matter is not accelerated sufficiently to become unbound. A pulsation which accelerates the outer layers of the star to close to, but not exceeding, escape velocity therefore loses a significant amount of energy without losing mass. The decoherence of different layers within the envelope likewise removes energy from the pulsation. The result of this decoherence – the compression shock – thermalizes the pulsation mode’s energy in the outer layers of the envelope, and acts to excite higher-order pulsation modes. As can be seen in the non-ejecting pulsating models in Fig. 4.7, a peak in amplitude is followed by higher-frequency pulsations which decay away – the high amplitude compression shock is followed by ringdown. Higher-order pulsation modes are not efficiently driven and so their amplitudes decay with time, and their presence also appears to disrupt the growth of the fundamental mode pulsation. The energy that the shock moves out of the fundamental mode and into higher-order pulsation modes therefore acts as a damping mechanism for the fundamental mode.

Because the bifurcation between experiencing a mass ejection event and merely experiencing damping and amplitude modulation appears to depend sensitively on the maximum amplitude reached during a specific phase of the pulsation cycle – the high-radius phase where nonlinear effects emerge – and on the recent history of the star’s pulsation evolution, the exact evolution of any given model is chaotic, and it is very difficult to predict whether a given pulsation cycle will lead to a mass ejection. Instead, we can only make predictions about the mass-loss properties of a given model over many pulsation cycles.

From this analysis, it is clear that the similarities between the mass ejections seen in our simulations of CE objects and these that are found in our AGB models bear strong similarities in their underlying physics as well as their appearance and effects. We posit that this is indeed an identical mechanism at work in these two regimes despite the different energy sources and heating locations.




Figure 4.14: The first mass ejection experienced by a hydrodynamical simulation of a $1.6 M_{\odot}$ star beginning during the helium burning shell flash of the 14th thermal pulse. Our mass removal algorithm was deactivated for this figure. We show the non-dimensionalised entropy per particle; the ratio of fluid velocity to local escape velocity in those regions where this ratio exceeds 1; the ratio of fluid velocity to local sound speed for regions where the flow is supersonic; and the relative ionization fractions for hydrogen, and for singly and doubly ionized helium. Also shown are contours which contain 100% (black), and 99, 98, 95, 90, 85, 80, 75, 70, and 50% (dashed grey) of the total mass of the model.




Figure 4.15: The same event appearing in Fig. 4.14, showing the specific rate of release of recombination energy; the Eulerian acceleration; the velocity divergence; the value of the first adiabatic exponent with white representing the critical value $4/3$; the pressure-weighted, volume averaged value of the first adiabatic exponent with white representing the critical value $4/3$; and one tenth of the product of local radius, opacity and density, a dimensionless quantity representative of the local optical thickness. Panel 6 also shows the radii at which the total optical depth is equal to 10^2 (dashed green), and 10^4 (dashed cyan). Also shown are contours which contain 100% (black), and 99, 98, 95, 90, 85, 80, 75, 70, 60, and 50% (dashed grey) of the total mass of the model.




Figure 4.16: The same event appearing in Fig. 4.14, showing the temperature; the density; the velocity; the sound speed; the opacity; and the pressure. Also shown are contours which contain 100% (black), and 99, 98, 95, 90, 85, 80, 75, 70, 60, and 50% (dashed grey) of the total mass of the model.




Figure 4.17: The WD initial-final mass relations we obtain for, with points and solid lines: a mass-loss formula made up of the Schröder wind plus our dynamical mass-loss model’s estimate multiplied by either 10 (green), 1 (dark blue), or 0.1 (red); with points and dashed lines, a mass-loss formula made up of only our dynamical mass-loss model’s estimate multiplied by either 10 (light blue), or 1 (orange). Also shown, in black without points, are estimates by Catalán et al. (2008) (dashed), Weidemann (2000) (dash-dotted), and Han et al. (1994) (dotted).

4.4 The white dwarf mass distribution

Having recovered a formula for predicting dynamical mass-loss rates as a function of elementary stellar properties, we can compare how well evolutionary simulations implementing that formula reproduce observations. We performed a series of full-lifetime simulations of stars in the mass range $1\text{--}3 M_{\odot}$ from the ZAMS to the WD stage using the same parameters as those reported in Section 4.1, but rather than using the Schröder wind mass-loss model for all simulations, we applied either a mass-loss model implementing our new formula for dynamical mass loss, or a model in which mass loss follows the sum of our formula and the Schröder wind. To investigate the effect of the uncertainty in our dynamical mass-loss rate estimates, we performed simulations in which the strength of our dynamical mass-loss model was multiplied by either 10, 1, or 0.1. In all cases, we applied our dynamical mass loss only when the predicted mass-loss rate due to that mechanism exceeded 1.6×10^{-4} , representing the threshold value we see in our hydrodynamical models. We used the results of these new evolutionary simulations to construct a series of initial-final mass relations, giving the mass of the WD remnant as a function of the ZAMS mass of the star. These relations are shown in Fig. 4.17, along with an observationally inferred estimate by Catalán et al. (2008), a semi-empirical estimate by Weidemann (2000), and a theoretical estimate by Han et al. (1994). We have followed Catalán et al. (2008) in fitting linear relationships to our results (see also Kalirai et al., 2008; Williams, 2007). We have excluded the results of simulations incorporating only our mass-loss model with strength multiplied by 0.1, as those simulations were generally terminated due to numerical instability before they reached the WD stage. This is not unexpected, as those models were becoming increasingly unstable due to an instability they were unable to follow (due to their hydrostatic nature) and moving into a region of parameter space which we do not expect stars to actually enter because they would already have lost their envelopes.

Unsurprisingly, weaker versions of our episodic mass-loss model lead to the production of more massive WDs, and the addition of the Schröder wind decreases the final WD mass.

Using our initial-final mass relations, we performed Monte-Carlo calculations to produce theoretical predictions for the mass distribution of WDs which should be present in the Milky Way at the present day. We constructed a crude model of the initial stellar population

under the assumption that star formation follows the Salpeter IMF (Salpeter, 1955), and has been proceeding at a constant rate since the birth of the Galaxy, which we assumed to be 1.1×10^{10} years ago. This calculation proceeds by generating a population of stars (we used 1,000,000) drawn from the Salpeter mass distribution with uniformly distributed ages, selecting only those stars which would have had time to evolve to the WD branch, and calculating their final masses. A Monte-Carlo method is preferred because performing direct calculations, while possible, are strongly affected by the values of the first differential of the initial-final mass relation, and are therefore highly-sensitive to the choice of interpolation scheme used. To ensure a self-consistent calculation, we interpolated stellar lifetime estimates from the same sets of simulations we used to generate the initial-final mass relations. The resulting WD mass distributions are shown in Fig. 4.18. We also plot in that figure an observationally inferred WD mass distribution calculated by Kepler et al. (2016), and extracted from spectral fits of 5,863 hydrogen-atmosphere WDs (DAs) with high signal-to-noise ratio spectra ($S/N \geq 15$) in the Sloan Digital Sky Survey Data Releases 7, 10, and 12, the data for which was kindly provided by those authors, and is described in more detail in Kepler et al. (2016, 2015) and Kleinman et al. (2013).

The most sensitive measure of the length of the AGB that can be derived from the WD mass distribution is the location of the peak, which moves to lower masses when AGB mass loss is overestimated and higher masses when it is underestimated. The top panel of Fig. 4.18 therefore suggests that, when both episodic mass loss and the Schröder wind are included, the combined model overestimates mass loss on the AGB, and a reduced value for the episodic mass-loss rate provides the best peak mass. However, it should be noted that the peak locations produced by our model are remarkably similar, even when the mass loss rate is varied considerably. We suggest that this is because, although a change in the estimated dynamical mass loss rate has a strong effect on the amount of mass that is lost during each helium shell flash, once the star becomes sufficiently unstable that dynamical mass loss begins during the hydrogen burning phase, the star is likely to lose its entire envelope within a single thermal pulse cycle, even if the dynamical mass-loss rate is reduced by more than an order of magnitude. The age on the AGB when this instability emerges is therefore much more important than the exact value of the mass-loss rate adopted.

In addition, we note that the Schröder wind has not been tested at mass-loss rates as

high as those attained on the extreme AGB (see Schröder & Cuntz, 2007), and we suggest that wind prescriptions for cool giants are likely to overestimate mass loss in this regime *in general*, because they are calibrated and tested under the assumption that significant dynamical mass loss does not also occur. When we apply our mass-loss model alone with no wind loss prescription, as shown in panel 2 of Fig. 4.18, we are able to reproduce the location of the WD mass peak remarkably well with the mass-loss rate our model suggests. We emphasise that this distribution is a first-principle theoretical estimate, and is not the result of calibration against observational mass-loss estimates. This result demonstrates that our dynamical mass loss model predicts sufficiently strong mass loss to be consistent with the WD mass distribution’s constraint on the length of the AGB (see e.g. Catalán et al. 2008 for a comparison of other estimates of the WD mass peak’s location).

The right edge of the WD mass distribution estimates produced for different versions of our model vary considerably, as these depend on the mass of the highest-mass WD predicted for a given distribution of ZAMS masses; this variation leads to the different heights of the distributions as a normalisation effect. We have truncated our WD mass distributions at a ZAMS mass of $3 M_{\odot}$, because, as discussed in Section 4.2.3, our model does not provide a good fit to simulations of higher-mass AGB stars. Those stars, however, have very little effect on the location of the WD mass distribution’s peak, as they are formed in too low numbers.

It is clear that the observationally inferred mass distribution is much wider than our theoretical distributions, and the peak’s location is less clear. This is likely due in part to errors in the observation and fitting processes¹, in part to variations in the ZAMS chemistry of WD progenitors (see, e.g. Romero et al. 2015), and in part to alternative pathways for producing WDs. In particular, our calculations take no account of duplicity, which may have a significant effect on the mass-loss histories of many stars, particularly on the AGB, through for example wind Roche lobe overflow (i.e. Case D mass transfer, see Podsiadlowski & Mohamed 2007), and pulsational Roche lobe overflow. It has been estimated that as many as 70% of WDs with masses below $0.45 M_{\odot}$ are the result of binary evolution (Kilic et al.,

¹In particular, observational estimates of the masses of white dwarfs are sensitive to the atmosphere model used to fit the WD spectrum (see e.g. Tremblay et al. 2011, Tremblay et al. 2013) and depend on evolutionary models to calculate mass values from the quantities derived from the fit (see Kepler 2016 for details of how this is accomplished). These factors may introduce significant systematic errors into the location of the mass peak.

2011).

4.5 Discussion

4.5.1 Dynamical instability as a driver of mass loss

The pulsations appearing in our models are similar in character to those obtained by Wood (1974), Tuchman et al. (1978), and Tuchman et al. (1979), and are likely to represent the same dynamical instability reported in the work of those authors. By applying a treatment to numerically resolve shocks that those early works lacked, we have corroborated those authors' results with a more physically consistent approach. Despite arguments made by later authors (Fox & Wood, 1982; Bessell et al., 1996) that large-amplitude dynamical pulsations should not appear in cool giants, we find that they are in fact likely to have a significant effect on the mass-loss history of stars on the AGB. We note that the pulsations we see in our simulations resemble those found by Heger et al. (1997), and Yoon & Cantiello (2010) in the context of RSGs, and we shall discuss further the possibility of seeing similar dynamical mass-loss in those stars in Chapter 5. Our results are also similar to the dynamical pulsations and associated mass ejections seen in the CE objects discussed in Chapter 3, likely due to the similarities in the expanded envelope structures of AGB stars and our CE models, despite the differences in the location of energy generation.

In our simulations, we find that although envelopes which are more dynamically unstable as measured through $\langle \Gamma_1(m) \rangle$, as well as those with more positive binding energies, do exhibit stronger instability, neither condition seems to form a definitive boundary between dynamically stable and unstable stars. From a theoretical point of view, this is not unexpected, as both of those possible theoretical conditions are calculated for the entire envelope, whereas the processes we see leading to mass ejection are localised in the outer envelope specifically, and, in particular, shocks in the envelope act to concentrate energy in the outer layers of the star shortly before ejection events occur, allowing the outer layers to gain large amounts of kinetic energy during the recombination-powered bounce event. Our results support the idea that formal dynamical instability, as calculated using $\langle \Gamma_1(m) \rangle$, and energetic stability calculated through envelope binding energy, are useful as *indicators* of the likely extent of dynamical behaviour, but do not form clear bifurcation points between




Figure 4.18: Theoretical estimates of the mass distribution of WDs in the local environment, calculated using our initial-final mass relations and stellar lifetime estimates. In the top panel, we show the distributions calculated using a mass-loss prescription that is the sum of the Schröder wind and our dynamical mass-loss model, where the strength of our model has been multiplied by either 10 (red), 1 (blue), or 0.1 (green). We show also the DA WD distribution inferred from observation by Kepler et al. (2016), in black. In the bottom panel, we show the distributions calculated using our dynamical mass-loss model alone with no wind component, and the strength of our model multiplied by 10 (orange), or 1 (light blue). The observationally inferred result of Kepler et al. (2016) is plotted in this panel also.

stable and unstable stars.

4.5.2 Uncertainties in our models

There are several sources of theoretical uncertainty in our results that we can identify. One of the most important is likely to be our treatment of convection using mixing length theory, which is not properly time-dependent (convective energy transport responds instantly to changes in stability under MLT, so it is time-dependent but in an unrealistic way). In principle, changes in convective energy transport over time can interact with the pulsation cycle in non-trivial ways, and differences in energy transport are likely to have an effect on the growth rate of pulsations. Unfortunately, currently-available prescriptions for time-dependent convection, for example those described by Stellingwerf (1982) and Kuhfuss (1986), are highly dependent on multiple free parameters which are difficult to calibrate (see Gastine & Ditrans 2011 for a discussion). The difficulty in treating convection remains one of the biggest unsolved problems in stellar pulsation theory; however, we note that in the stars we examine in this work, the pulsations that emerge have sufficiently long periods that the convective turnover time remains significantly shorter than the pulsation period. This suggests that convection is likely to be able to adjust to changes in structure and stability caused by the pulsation sufficiently quickly that our “instantly adjusting” prescription should be approximately correct. Furthermore, in a study of similar long-period pulsations in RSG models, Heger et al. (1997) found that neither pulsation periods nor growth rates were strongly affected by the response rate of convective fluxes to stability changes, and the addition of artificial lag times between pulsation and convective flux changes has also been found not to cause strong changes to pulsation properties (Langer, 1971).

Also of concern is the 3-dimensional nature of AGB envelopes and outflows, which are likely to have significant deviations from spherical symmetry due to large-scale convective plumes (see, for example Ohnaka et al., 2016; Freytag & Höfner, 2008; Stewart et al., 2016). It is unclear what effect this non-sphericity will have on dynamical mass ejections, although it is noteworthy that large-amplitude radial pulsations are still seen to emerge in 3-dimensional hydrodynamical simulations of AGB envelopes, and display broadly similar properties to those calculated from 1-dimensional models (Freytag et al., 2017). The effects of 3-d structure and anisotropy on mass ejection episodes specifically require further study.

4.5.3 Turbulent viscous damping

Dynamical mass-loss in our simulations is dependent on pulsations being able to grow to large amplitudes and to do so quickly (i.e. within the duration of the helium shell flash). It has been suggested that models of nonlinear pulsations in cool giants exhibit pulsation growth that is unrealistically strong (Olivier & Wood, 2005), and that these pulsations are suppressed by the treatment in simulations of the damping effects of viscous pressure generated in turbulent convective eddies (the effects of turbulence on pressure and energy transport were also considered by those authors, but were found to have much smaller effects). It is therefore important that we consider the possible effects of these processes on our results. Attempts to incorporate terms representing turbulent viscous pressure into 1-d nonlinear simulations of stellar pulsations have been made by several groups such as Olivier & Wood (2005), Kolláth et al. (2002), and Keller & Wood (2006), following one of several similar prescriptions derived by, for example, Kuhfuss (1986) or Stellingwerf (1982). Olivier & Wood (2005) in particular applied such a prescription to nonlinear pulsations of red giant envelopes and found significant effects on pulsation amplitudes. Unfortunately, as with all of our parameterised models of convection in stars, the available prescriptions for turbulent viscous pressure depend on free parameters which are difficult to calibrate, and the true strength of this effect remains highly uncertain.

In order to examine how turbulent viscous pressure affects our results, we carried out a series of tests in which we performed hydrodynamical simulations with an added turbulent viscous pressure term with varying levels of strength. Although none of the prescriptions developed for this effect are convincingly well-founded, they are likely to provide reasonable estimates of its magnitude. In our simulations, we applied a viscous pressure prescription developed by Von Neumann & Richtmyer (1950), and Weaver et al. (1978), in which we apply a pressure of the form

$$P_v = \frac{1}{r^3 \rho} \frac{d}{dr} \left(\frac{3}{4} \mu r^4 \frac{d}{dr} \left(\frac{v}{r} \right) \right), \quad (4.5)$$

where ρ is density, r is local radius, v is local bulk velocity, and μ is a dynamical viscosity term whose value is equal to the product of some constant of order 1, a characterising length scale for convective eddies, and a characteristic velocity scale. We use $\mu = \beta r v_c$, where

β is the dimensionless constant of order 1, and v_c is the sound speed. These terms are likely to give correct order upper bounds for the characteristic scales of convective eddies in the envelope. The implementation of this damping term in MESA is designed to ensure simultaneous conservation of momentum and energy, and is given in detail in Paxton et al. (2015).

The results of one of these tests are shown in Fig. 4.19. We found in general that the addition of turbulent viscous pressure had no significant effect on our simulations until its strength reached very high values ($\beta \geq 10$). At these high values, a decrease in amplitude growth rate is visible and the rate of mass loss is suppressed. However, it is likely that a β of this magnitude is unphysically high; for comparison, Olivier & Wood (2005), Keller & Wood (2006), and Kolláth et al. (2002) found that, when calibrating similar expressions for the turbulent viscous pressure, they required a scale constant of between 0.1 and 0.2. Those authors also used estimates of turbulent eddy length and velocity scales derived from their convection models which are likely to be smaller than the local radius and sound speed we use in our model. Comparable strengths of effect are therefore likely to be found with β less than 0.1, for which we find no effect on the qualitative nature of our results.

We suggest that the reason we do not find a strong effect on our results from the addition of turbulent viscous pressure is because, in our simulations, the pulsation amplitude is limited primarily by nonlinear damping effects emerging at high amplitude which act as catastrophic damping methods for the pulsation once a threshold amplitude has been reached. The addition of another, comparatively weaker damping mechanism does not strongly affect the location of this threshold amplitude and therefore does not have a strong effect on our results. Our interpretation is that the pulsations seen in our models grow sufficiently strongly that, even with the addition of reasonable levels of turbulent viscous damping, the limiting threshold amplitudes are still reached. For very large strengths, the added damping *can* become dominant and have a strong effect, preventing pulsations from reaching the catastrophic damping thresholds, but only when applied at strengths which are likely to be unphysically high. Fig. 4.19 also provides a good demonstration of the chaotic nature of dynamical pulsations: even very small changes in the value of the turbulent viscous scaling parameter (for example, between the top two panels) can lead to significant changes in the *details* of the evolution of the system after only a small number of pulsation cycles,

whilst leaving their qualitative nature intact.

4.5.4 Observational signatures

4.5.4.1 Circumstellar relics

There is considerable evidence that mass loss from AGB stars exhibits short term variability towards the end of the stars' lives, derived from analyses of circumstellar material around AGBs in the form of detached dusty shells (see for example Olofsson 1996, Hashimoto et al. 1998, Schöier et al. 2005, Sanchez et al. 2015), and that this variability is at least in part associated with helium shell flashes (Olofsson et al., 1990). Calculations of the absolute values of mass loss involved in creating these structures are difficult, however, as they depend on distance estimates, and there remains much uncertainty as to how these structures are formed (see, for example, Schröder et al., 1998). If the episodic mass ejections reported in this chapter do occur in reality, they would be likely to produce detached shell structures of this kind, each containing between a tenth and a hundredth of a stellar mass of material. We make no attempt to model the evolution of these detached shells once they leave the star, but it may in principle be possible to differentiate between a shell produced by a variation in wind strength and one resulting from a dynamical ejection event due to the different kinematics involved in the two processes.

Concentric shells of material are also seen in the halos of planetary nebulae (called “rings” or “arcs”, although they are generally believed to be spherical), which suggest that the progenitors of those objects experienced episodic mass loss on timescales of hundreds of years – significantly longer than the pulsation timescale of Miras (up to ~ 1000 days) but shorter than the thermal pulse period ($\sim 10,000$ years or longer) (Corradi et al., 2004; Phillips et al., 2009). Some of these structures are spiral in form and likely the result of binary interactions (Kim & Taam, 2012; Maercker et al., 2012), but they also appear as complete, discrete shells, and in spherical PNe which lack evidence of shaping by a binary companion (see Corradi et al., 2007). The time between shell ejections in our simulations ranges between approximately 5 and 100 years, with separations of a few decades being common. Whilst this appears slightly too short to match the ring separation in PNe halos, it is important to point out that values given for the ring separation timescale depend on assuming a ring expansion velocity, which has never been measured directly and is basically




Figure 4.19: The evolution of radius (in blue) and stellar mass (in red) of a hydrodynamical simulation of our $1.6 M_{\odot}$ star beginning during the helium shell flash. Each panel shows an identical simulation performed with a different strength of turbulent viscous pressure. The numerical value appearing in the formula for μ is the quantity β .

unknown (see Ramos-Larios et al., 2016). Perhaps a more serious problem is the irregularity in the spacing between ejections seen in our simulations, whereas the spacing of PN halo shells appears to be highly regular. It is possible, however, that the radiation-hydrodynamics that determines the evolution of shells after their ejection may have a regularising effect, as discussed by Ramos-Larios et al. (2016), but studies of this evolution have not yet been published.

4.5.5 The possibility of direct observation

The stars in our simulations exhibit an extremely high degree of mass loss when in an “ejecting” mode, and are therefore likely to be shrouded by large amounts of dusty circumstellar material. It therefore seems reasonable to identify these stars with OH/IR stars – heavily dust-shrouded stars undergoing enhanced mass loss which are difficult to observe directly – which are believed to be the final phase AGB stars enter at the very ends of their lives. The models that exhibit self-limiting pulsations and do not experience mass ejections, however (for example in panel 2 of Fig. 4.7), display steady pulsations with periods of a few hundred days and velocity amplitudes of order 10 km s^{-1} , similar to the properties of Mira variables, which are expected to be seen on the TP-AGB. The models which exhibit mass ejection events (for example those shown in panels 4 and 5 of Fig. 4.7) show pulsations which are likely too violent and irregular to represent Miras.

Our results therefore suggest that a sliding scale in pulsation amplitude and growth rate connects Miras and OH/IR stars and that, if their pulsation properties change as a result of luminosity changes during the thermal pulse cycle, stars may migrate between the two regimes. All the pulsations we observe in our simulations are fundamental mode pulsations, which is believed to be the most likely operation mode for Miras (e.g. Bessell et al., 1996; Ireland et al., 2008). This identification of non-ejecting models with Miras and ejecting models with OH/IR stars suggests that the decreased stability of AGB stars as they evolve along the TP-AGB which we see in our simulations, leading eventually to episodes of dynamical mass loss, represents the evolution of AGBs from Miras into OH/IR stars. This evolution into a state of greatly-increased mass loss mirrors the well-known superwind theory (Renzini & Voli, 1981; Vassiliadis & Wood, 1993), commonly invoked to provide an explanation for the very strong mass-loss rates reported in OH/IR stars (Engels et al.,

1983; Habing et al., 1983), and filling the theoretical requirement for a rapid mass-loss mechanism to produce PNe.

In our simulations, we find that stars enter a mass-ejecting mode at first only during the helium shell flash and associated luminosity peak of the thermal pulse. We find also that, due to the short length of this phase, a star may well fail to lose its entire envelope during one shell flash, and may survive into the following, low-luminosity phase of the thermal pulse cycle, during which we see no dynamical mass loss occurring. This suggests that AGB stars may be expected to intermit between a high-mass-loss, OH/IR-like state during the helium shell flash, and a low-mass-loss (although still high by normal stellar standards), Mira-like state, when their mass loss would return to being dominated by (possibly pulsational) wind mechanisms.

4.6 Summary

We carried out coupled evolutionary and hydrodynamical simulations of 1-6 M_{\odot} stars to study their dynamical properties at different stages on the TP-AGB and at different phases of the thermal pulse cycle. We found that sufficiently evolved TP-AGB stars become unstable to large-amplitude dynamical pulsations with periods of between 100 days and 5 years. As a star progresses along the TP-AGB, the growth rate of these pulsations increases and they become increasingly powerful, eventually becoming powerful enough to eject shells of matter dynamically from the surface of the star in repeated mass-loss episodes, with individual ejections varying between approximately 0.01 M_{\odot} and a few times 0.1 M_{\odot} . These mass-loss episodes repeat on timescales between just under 10 years and hundreds of years, and represent a time-averaged mass-loss rate of up to approximately $10^{-3} M_{\odot} \text{ yr}^{-1}$. We found that both dynamical pulsations and mass ejections emerge earlier during the helium shell flash phase of the thermal pulse cycle than during other phases and that either all, or a significant part of the dynamical mass loss experienced by an AGB star occurs during these shell flashes.

We constructed a parameterised model describing dynamical mass loss as a function of elementary stellar parameters and carried out additional stellar evolution calculations using this model as a mass-loss prescription. We found that our episodic mass-loss prescription

acted to remove the envelope over the course of 1–3 thermal pulse cycles, with the majority of the mass loss occurring during the helium flashes, although stars that survived long enough did eventually become unstable enough to experience dynamical mass loss during the hydrogen burning phase.

We calculated theoretical estimates of the initial-final mass function for WDs using full-lifetime stellar evolution simulations incorporating our dynamical mass-loss model, and calculated theoretical estimates for the local WD mass distribution based on those results. Our model is successful in reproducing the location of the observationally inferred WD mass peak, and therefore at locating the termination of the AGB, to a good level of accuracy, despite relying on no observationally calibrated parameters. However, we found that when combined with a wind-loss prescription, the total mass loss is likely to be an overestimate. It is not clear at this point which mass-loss mechanism has been overestimated.

It is possible that dynamical mass ejections may be responsible for producing the observed discrete mass shells surrounding both AGB stars and the cores of PNe. Full radiative hydrodynamics simulations of ejected mass shells are necessary to understand how these shells behave and interact after they leave the star, and in particular whether the interaction between concentric shells has a regularising effect on their arrangement in space.

Our model for dynamical mass loss is not expected to be effective at describing the dynamics of higher-intermediate mass stars (on the so-called massive AGB), and more work is needed to explore the dynamical properties of these stars.

Chapter 5

Red supergiants

Mass loss from RSG stars is amongst the least well understood of all stellar regimes despite the fact that it is of primary importance in determining the upper mass and luminosity limits of the RSG phase and in determining type of supernova experienced by stars which pass through that phase. In this chapter, we will examine the possibility that dynamical mass loss plays an important role in the mass-loss histories of RSG stars and discuss whether this mechanism can help with our understanding of the Humphreys-Davidson limit, the initial mass limit for SNe II-P, and the origin of SNe IIn. The most advanced dynamical simulation work in this regime, by Yoon & Cantiello (2010), studied dynamical pulsations in RSGs but did not examine dynamical mass loss. Their simulations contained only the initial growth phase of RSG pulsations and they did not make an attempt to follow the dynamical evolution of their stars after their pulsations had become supersonic. We shall extend this work into the shock-dominated regime and examine how the growth of pulsations connects to dynamical mass loss.

5.1 Simulation methodology

In our examination of RSGs, we again employ the hydrodynamical sampling method introduced in Chapter 4, with a series of hydrodynamical simulations beginning from stellar models extracted from hydrostatic evolutionary simulations of RSG stars. We performed evolutionary simulations of stars with initial masses of 15, 17.5, 20, 25, 30, and 40 M_{\odot} and an initial metallicity of $Z = 0.02$ using MESA version 8845. We construct these stars as uniform composition pre-main-sequence models, as described in Chapter 2, and follow their

evolution until they approached core oxygen exhaustion, with a core silicon mass fraction exceeding 0.1. The silicon burning stage is omitted both because it is extremely numerically challenging to simulate and because its duration of order days – orders of magnitude shorter than the dynamical timescale of the envelope – is short enough that no significant further mass loss can be expected to occur (see e.g. Kippenhahn et al., 2012).

We adopted the Eddington grey atmosphere outer boundary condition and applied the “Dutch” mass-loss prescription (see Chapter 2). Chemical mixing in massive stars can be complex, with semiconvective layers developing outside the core at the end of the main sequence; we adopt the Ledoux criterion for convective stability and apply a semiconvection prescription due to Langer et al. (1985), with a semiconvection efficiency parameter α_{sc} of 0.1 (this parameter is ill-constrained; we adopt a moderate value following Choi et al. 2016). We apply convective step overshooting over a distance of 0.2 pressure scale heights outside central convection zones. We do not apply overshooting to convection zones which do not extend to the centre of the star in order to avoid unphysical mixing across strong abundance boundaries.

An important uncertainty in the modelling of RSGs is their temperature scale. The effective temperature of these stars is difficult to measure due to strong line blanketing and difficulty in modelling the molecular opacities that are dominant in their atmospheres. Recent estimates of RSG effective temperatures by Levesque et al. (2005) are approximately 400 K higher than the older generation of estimates given by for example Humphreys (1978), with current best estimates for the galactic RSG temperature scale occupying 3450–4100 K, with estimates for the Large and Small Magellanic Clouds cooler by 50 K and 150 K respectively (Levesque et al., 2006). The effective temperature scale of giant models is set by the efficiency of convective energy transport, specified in evolutionary simulations by the value of α_{MLT} . In general, stellar models are calculated under the assumption that this parameter is the same across different classes of stars, most importantly because this allows the use of values calibrated from the Sun to be used for simulations of other stellar regimes. However, there is no strong theoretical reason to believe this is the case and, in RSGs in particular where convective plumes can be a significant fraction of the star’s size and convective velocities may be supersonic, this comparison may be unreliable. To ensure that we cover the available temperature range well and to allow comparison to previous

work such as that by Yoon & Cantiello (2010), in which RSG temperatures were probably too cool, we use three values for α_{MLT} : 1.5, 2.0, and 2.5. Higher values correspond to more efficient convective energy transport and therefore to smaller giants and higher effective temperatures. We apply the Henyey formulation of the mixing length theory, an extension of the theory designed to give more accurate results in regions which are not strongly optically thick (Henyey et al., 1965).

We performed 3 sets of hydrostatic evolutionary simulations, one for each value of α_{MLT} , with 6 initial mass values per set. We then performed hydrodynamical simulations starting from models at different points along these evolutionary sequences. We selected starting points for these simulations to sample a large range of stellar parameters as the RSGs cross the Hertzsprung Gap and ascend the giant branch.

Our hydrodynamical simulations were performed using MESA version 7624, making use of its backward-difference hydrodynamics treatment. We employ an artificial viscosity shock spread parameter of 2% of local radius, a slightly smaller value than in Chapters 3 and 4, which we found to work effectively in this regime, possibly due to the generally higher envelope densities compared to AGB stars. The RSG is a numerically difficult regime to work in, and we found that to construct hydrodynamical simulations that remained reasonably numerically stable throughout high-amplitude pulsation cycles it was necessary to relax the position of the outer boundary condition inward at the start of the simulation from the optical depth of $2/3$ used in our hydrostatic simulations to an optical depth of 20. This effectively moves low-opacity material of low optical thickness near the surface of the star out of the directly simulated region and into the stellar atmosphere model. We found that selecting different positions for this new outer boundary condition did not have a strong effect on the evolution of the models other than in terms of their numerical stability. In particular, because the material moved to the atmosphere generally consists of neutral material, the majority of the partially-ionized matter which contributes to pulsation driving is not affected.

We adopted the same maximum timestep in our hydrodynamical models as that used in Chapter 4, with timesteps limited to not exceeding 0.01 years. This value remains orders of magnitude shorter than the pulsation period of the star. We also continued to apply our custom mass-removal routine (see Chapter 2, Appendix A) to excise shells of matter

that have been dynamically ejected from the star, with ejected material removed from the simulation at 100 times the mass of the ejected layer per year. Wind losses were neglected during our hydrodynamical simulations. All hydrodynamical simulations were run for 200 years of star time or until they encountered unrecoverable numerical difficulties, whichever was the earlier.

To prevent spurious oscillations created by the relaxation from hydrostatic to hydrodynamical models or by our changing of the outer boundary condition from appearing in our results, we apply strong artificial damping to the first 700 models of our hydrodynamical simulations (equivalent to the first 7 years of star time). We apply this damping by reducing all velocities by 4% per timestep. This damping method removes pulsations whilst still allowing the star to equilibrate in response to its new simulation parameters. Pulsation growth can then proceed after this damping period has ended.

5.2 Results

We performed 18 hydrostatic simulations with 6 starting masses and 131 hydrodynamical simulations, sampling the behaviour of these stars over their post-MS lifetimes. The evolutionary tracks of the hydrostatic simulations along with the positions of the starting models of the hydrodynamical simulations can be seen in Fig. 5.1. 70 of our simulations develop high-amplitude dynamical pulsations, all in the fundamental mode, which grow quickly and become nonlinear, with stellar radius varying by up to a factor of 10. Examples of such pulsating simulations are shown in Fig. 5.2. Some simulations, especially those begun at earlier points in the stars' lives, do not exhibit any pulsations, although it is possible that the stars they represent would in fact pulsate in reality, but their pulsations are not sufficiently strongly driven to appear in our simulations.

We measured the growth rate of pulsations using the same method applied in Chapter 4, fitting the local maxima and minima of the surface radius curve with two diverging exponentials. We performed a simultaneous least-squares fit to the two sets of radius extrema using the fitting function

$$R = \alpha_1 + \alpha_2 t + \alpha_3 t^2 + \alpha_4 t^3 \pm \alpha_5 e^{\alpha_{GR} t}, \quad (5.1)$$

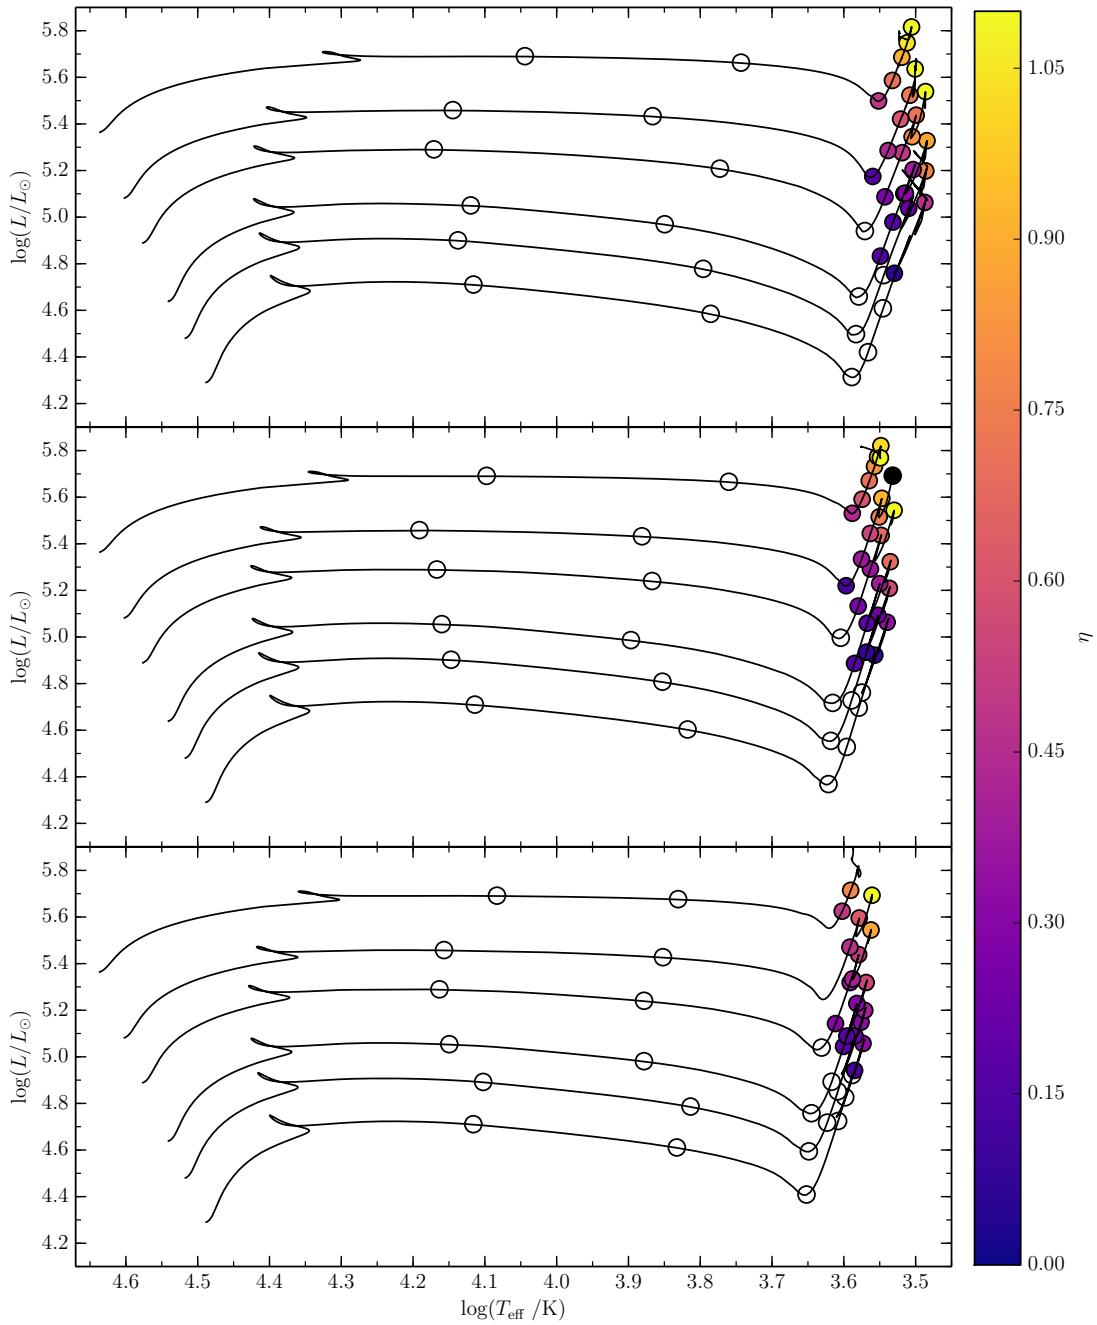


Figure 5.1: The starting positions of our hydrodynamical models on the HR diagram for mixing length parameter α_{MLT} of 1.5 (top), 2.0 (middle), and 2.5 (bottom). In each case, evolutionary tracks are shown for stars with initial masses of 15, 17.5, 20, 25, 30, and 40 M_{\odot} . Points are coloured according to the value of the non-dimensional growth rate η (see text). White points represent simulations which display no visible pulsations. The single black point represents a model which does pulsate, but in too violent a manner for its amplitude growth rate and pulsation period to be measured.




Figure 5.2: Four hydrodynamical simulations starting at different points in the life of a $25 M_{\odot}$ RSG from our $\alpha_{\text{MLT}} = 2.0$ simulation set. Photospheric radius is shown in blue, and total model mass is shown in red. For simulations 4 and 5 (top two panels) the evolution of the star’s radius over the same time period according to the *hydrostatic* model is shown in black dashed. In simulation 7 (bottom panel) the mass of the helium core (defined as the region within which the hydrogen mass abundance falls below 10^{-3}) is shown in green dashed.

giving radius R as a function of time t . The positive exponential corresponds to the series of local maxima and the negative exponential to the local minima. Although the “baseline” radius variation due to long-term nuclear and thermal evolution of the star is not as fast as in some of the AGB stars whose pulsations we fit in this way in Chapter 4, the use of a third-degree polynomial to allow for this variation remains helpful for attaining acceptable fits. We exclude all fits based on fewer than 6 extrema (which in practice is only 1 simulation), and all fits were visually checked to avoid overfitting and spurious results.

We defined growth rate in Chapter 4 as the value of the α_{GR} parameter, with units yr^{-1} ; however, in keeping with similar work in the cool supergiant regime (e.g. Yoon & Cantiello, 2010; Moriya & Langer, 2015), we will write our results here in terms of the dimensionless pulsation growth parameter η , which is defined as the product of growth rate and pulsation period P in years:

$$\eta = \alpha_{\text{GR}}P. \quad (5.2)$$

For an exponentially growing pulsation where the motion of an element at an initial radius r_0 is equal to $\delta r(r_0)e^{i\sigma t}$, where σ is in general a complex quantity. The parameter η is equal to minus the ratio of the imaginary part of σ to its real part. Note that while our definition of η is the same as that used by Moriya & Langer (2015), Yoon & Cantiello (2010) use a slightly different definition, so that η in this work is equal to $\log(\eta)$ in their analysis. Whether η or α_{GR} is the more basic physical property is debatable, but we found η to be modelled more easily than α_{GR} alone in this regime (see below). Pulsation periods were obtained by differencing the time positions of successive radius maxima.

Fig. 5.1 shows the values of η for each of the 69 simulations for which pulsation growth fits could be obtained. It is clear that stronger pulsation growth appears in the cooler simulations associated with lower values of α_{MLT} and that stars become more unstable as they ascend the RSG branch. All of these 69 simulations are also shown in Fig. 5.3 plotted on a L/M ratio versus T_{eff} plane.

Dynamical mass ejection events occur in 41 of our simulations. 15 of these simulations display only one mass ejection, as in the two middle panels of Fig. 5.2, while 26 simulations show repeated mass ejections, as in panel 4 of Fig.5.2. These mass ejection vary in size between of order $0.001 M_{\odot}$ to a few solar masses. Repeated mass ejections are separated




Figure 5.3: The positions of our hydrodynamical models' starting points in the L/M vs T_{eff} plane. Points are coloured according to the non-dimensional growth rate η (see text). Simulations in which no pulsations are seen or (in 1 case) are too violent to characterise, are not shown. Points are joined into sequences according to which hydrostatic model they begin from. Points joined in green have a mixing length α_{MLT} of 1.5, points connected in red have a value of 2.0, and points connected in blue have a value of 2.5. In general, simulations based on stars with a higher initial mass appear further to the top left of the diagram.

by intervals of between approximately 2 and 20 years. Some of our simulations experience quite strong pulsations but do not eject mass, settling into a steady pulsation period and amplitude; others eject mass on one or a small number of occasions, then settle into a steady pulsation pattern with a lower amplitude, and others eject mass many times until they have lost the majority of their envelope. Once the envelope is reduced to $1\text{--}2 M_{\odot}$, the model becomes increasingly numerically unstable and it becomes impossible to continue the simulation. Each of these regimes is visible in Fig. 5.2.

The RSG simulations which experience repeated mass ejections display far more regularity than in AGB or CE simulations undergoing similar repeated ejections and are capable of ejecting mass every pulsation cycle for many such cycles. This is possibly just an effect of the envelopes being more massive and therefore being more difficult to perturb, but, as can be seen in the example shown in the bottom panel of Fig. 5.2, this regularity can persist even as the envelope undergoes very strong changes in mass.

5.2.1 The Ya’ari effect

Our pulsating models commonly exhibit rapidly-growing pulsations which reach some maximum amplitude value and then settle into a steady pulsation with a lower amplitude. In many cases the period of the pulsation also changes noticeably when this happens (this is easily visible in Fig. 5.2). This phenomenon, first described by Ya’ari & Tuchman (1996), in which nonlinear pulsation models change amplitude and pulsation period significantly after some initial growth period, is believed to be due to a rearrangement of the thermal structure of the envelope as a response to the presence of the pulsation. This effect has been widely noted as possibly occurring in long-period variables, red giants, and type-II Cepheids undergoing observed or inferred period and amplitude changes (e.g. Lebzelter & Wood, 2005; Wood, 2006, 2007; Uttenthaler et al., 2011; Smolec, 2016), but has never been systematically investigated.

We have plotted the radius and entropy at various mass coordinates within an RSG model in Fig. 5.4; it is clear that the model plotted undergoes amplitude and period changes whilst at the same time both the shape of the pulsation and the baseline “equilibrium” values of radius and entropy change. The envelope as a whole contracts, and layers internal to the star lose entropy while layers near the surface increase in entropy. This figure can be




Figure 5.4: The evolution of a hydrodynamical model begun during the RSG stage of a star with an initial mass of $20 M_{\odot}$. In the top panel we show the radii of Lagrangian mass coordinates within the stellar envelope containing 100, 90, 80, 70, 60, 50, and 45% of the star's total mass. The bottom panel shows the evolution of the logarithm of dimensionless entropy per particle at Lagrangian mass coordinates containing 95, 90, 80, 70, 55, 47.5, and 42% of the star's total mass. We emphasise that these are not the same coordinates as in the top panel and have been chosen for ease of visibility. This simulation does *not* eject mass – the changes in structure and pulsation are not related to mass loss. In the bottom panel, lines appearing at higher entropy are deeper within the star.

compared to a similar one produced by Ya'ari & Tuchman (1996) which produces the same result down to a sign inversion¹. Although many of our simulations display period and amplitude changes due to this effect, we do not observe any mode-switching or the presence of overtones either before or after the star has undergone thermal adjustment.

This effect presents a difficulty for our analysis because it suggests that the properties of a pulsation during its initial growth are not necessarily strongly indicative of its properties once established. Our approach of using hydrodynamical simulations to sample the dynamical behaviour of RSGs requires us to start our hydrodynamical simulations from non-pulsating hydrostatic models, whereas in actuality such a star would already have an established pulsation. Of course all pulsations must start somewhere, but the growth process is likely to be far more gradual than in our simulations as the envelopes of RSGs become increasingly unstable due to the steady nuclear evolution of the stars. In particular, this issue is a bigger problem in this regime than in AGB stars, in which stellar luminosity can change significantly even on the timescales of dynamical pulsation growth, and pulsations *are* likely to grow suddenly to large amplitude (as seen in Fig. 4.9, where pulsations grow in a shorter time period than the duration of the ~ 100 year helium shell flash). RSG properties do change quickly relative to many other stars, especially at the start of the RSG phase (see for example the change in hydrostatic radius in panel 1 of Fig. 5.2, which occurs during the ascension of the giant branch when luminosity and radius double in approximately 2000 years), but not to nearly the same extent as AGBs undergoing thermal pulses (where luminosity and radius can change by a factor of 5 in less than a century). Whilst the initial pulsation growth period is not likely to be physical, the pulsation growth rate here *is* still valuable as a measure of the level of instability of the envelope.

Because of the likely-unphysical amplitude peak that follows the initial growth phase, many of the simulations in which we see only one mass ejection or a small number of mass ejections during this peak cannot be relied upon to represent stars which would actually

¹In their paper, Ya'ari and Tuchman claim that it is the layers within the star that gain entropy and those near the surface that lose it. It is likely that this is an error: it is clear from Fig. 4 of their paper that this arrangement would require an entropy profile that increases outward through the star, which would be convectively stable and therefore unrealistic for a giant, and it would also require that their pulsation produce much greater entropy perturbations at the base of the envelope than near the surface, which is not physically realistic. It is even possible to see the node of a first-overtone pulsation in their entropy plot, which is aligned in the opposite direction to the one claimed. It seems clear that they have mistaken the order of the entropy profiles in their analysis, in which case all these problems would be immediately resolved.

exhibit mass ejections in reality. The true behaviour of the star is likely to be closer to the steady-state pulsation that occurs after the amplitude has peaked and then fallen, which in many cases occurs without mass ejections. The pulsations shown in panels 2 and 3 of Fig. 5.2 fall into this regime. We do, however, see 10 simulations in which repeated mass ejections continue once the pulsation has settled into a steady state, as in panel 4 of Fig. 5.2. These stars are much more likely to actually eject mass, because their dynamical mass loss cannot be ascribed to the effects of initial pulsation growth before the envelope has thermally adjusted, and it continues after this process has finished.

This issue makes it difficult to properly differentiate between models which are expected to actually represent stars undergoing dynamical mass loss and stars which are not. A suitable comparison of pulsation properties after the envelope has adjusted and the pulsation has attained a steady state may be effective, but it would require a more advanced knowledge of the effects of dynamical mass loss on the envelope, as many of these simulations only reach a steady pulsation state after undergoing a possibly-spurious mass ejection. We also experienced great difficulty in ensuring that our models maintained numerical stability after experiencing initial mass ejections, and many simulations were terminated early for this reason. A complete understanding of this effect and its consequences may not be possible with this set of numerical tools.

5.2.2 Modelling pulsation growth

To better describe the regime in which dynamical pulsations become important, we can attempt to construct models that parameterise the growth rate of pulsations in our RSG simulations as a function of basic stellar properties. We performed least-squares nonlinear curve fits to our measurements of η using power-law models in a variety of physical properties of the star. The results of three of these fits can be seen in Fig. 5.5. In that figure we show a fit using the same parameters used to model pulsations in AGB envelopes in Chapter 4 (although it should be remembered that the measure of pulsation growth rate used in that chapter is not identical to that used here), and a fit using the same parameters as Yoon & Cantiello (2010) use in their similar model. We also found that a good level of accuracy could be attained using a model parameterised only in terms of the effective temperature and the dynamical timescale of the envelope, and this model is shown also in

Fig. 5.5.

We find that the pulsation growth rate in our RSG models scales similarly to our model of AGB stars, with the exception of its dependence on effective temperature. It is however difficult to disentangle whether this change represents an actual difference in the physics at work or merely an effect of the coupling of parameters: T_{eff} can be trivially reduced to a function of radius and luminosity using the Stefan-Boltzmann law (Eq. 2.12). Our results also scale similarly to those of Yoon & Cantiello, with higher L/M ratios and smaller values of τ_{KH} indicative of greater levels of instability. We found that our growth rates scale considerably less strongly than those of Yoon & Cantiello, and their expression tends to underestimate the growth rates seen in our simulations by a factor of two. We believe that these differences are likely to be due to differences in simulation parameters, although possibly also in the finite difference schemes being used, which can introduce possibly non-physical damping effects (see e.g. Appenzeller, 1970).

In order to test if the dependence of our η models on mass and therefore luminosity was robust, we carried out additional test simulations in which we artificially stripped increasing fractions of the envelope mass from our $15 M_{\odot}$ and $25 M_{\odot}$ models before starting the hydrodynamical simulations. We found that the changes in stellar properties and growth rates due to this artificial stripping matched our models reasonably well for envelopes of down to $3 M_{\odot}$ in mass, but were unable to investigate smaller envelopes due to numerical difficulties experienced by those models.

The fact that the temperature scaling in our M , L , T_{eff} fit is in the opposite direction from in our T_{eff} , τ_{dyn} fit demonstrates the difficulty in trying to interrogate these formulae to determine the dependence of pulsation growth on some particular variable – the coupling between stellar variables is highly confounding. Plotting the growth rate data against different parameters individually, as in Fig. 5.3, can be helpful in determining how each variable affects pulsation growth.

5.2.3 Dynamical mass loss

Although 41 of our hydrodynamical RSG simulations display dynamical mass loss, in many of these cases only a small number of mass ejection events occur, and these ejections are located around the amplitude peak that follows the initial pulsation growth phase. As this

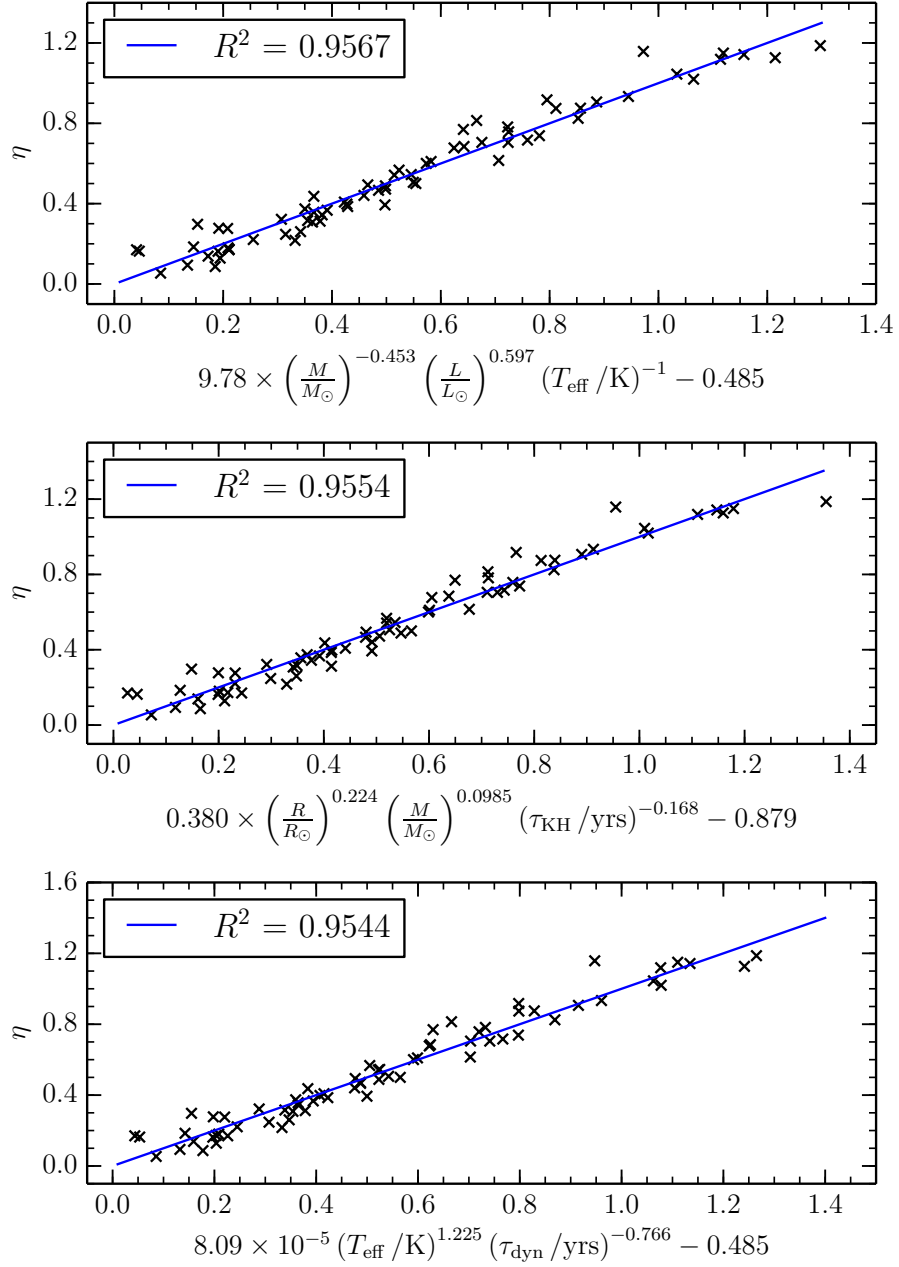


Figure 5.5: The value of the non-dimensional pulsation growth rate for all 69 of our well-characterised pulsating simulations plotted against 3 simple models which are parameterised in terms of basic stellar properties. Each model is derived via a least-squares fit. The coefficient of determination for each model is shown inset.

peak is not expected to be present in actual RSGs due to the expected slow (on the timescales of our simulations) onset of high-amplitude pulsations, many of these mass ejections should be treated as spurious. However, we do also see a regime of behaviour in which a large number of repeated mass ejection events occur which do *not* appear to be a result of the amplitude peak (for example, in panel 4 of Fig. 5.2) and as such are much more likely to represent realistic dynamical mass-loss behaviour.

The regime of repeated dynamical mass loss not associated with the initial pulsation growth or amplitude peak can be seen in 10 of our simulations, all based on our 25, 30, or 40 M_{\odot} stars. Several more simulations are terminated due to numerical stability problems too early to determine their long-term behaviour and so may also act in this manner. As such, it is difficult to determine the boundary of the region in which this behaviour occurs, but we can provide a basic estimate of that region. In Fig. 5.6, the locations of the 10 simulations which experience mass loss in this regime are shown, along with the time-averaged mass-loss rates that appear in these simulations (in calculating these values, this average is calculated only over the period of the simulation *after* the initial amplitude peak has decayed significantly). We also show the relationship between mass-loss rate and L/M ratio in Fig. 5.7. By simple linear regression we obtain the formula

$$\log(\dot{M} / M_{\odot} \text{ yr}^{-1}) = 5.93 \times \log\left(\frac{L/L_{\odot}}{M/M_{\odot}}\right) - 26.6. \quad (5.3)$$

It is important to remember, however, that this formula can only be used to describe the behaviour of models which *do* exhibit repeated mass ejections. There exists a cut-off below which no repeated mass-loss events are seen, and it is difficult to determine where this cut-off should best be placed. This is complicated by the fact that there are models in which this mass-loss regime is not seen which also occupy this region of parameter space, but, as can be seen in Fig. 5.6, most of these which appear above a $\log\left(\frac{L/L_{\odot}}{M/M_{\odot}}\right)$ value of 4.1 are terminated by convergence failures before their behaviour can be characterised and therefore could represent mass-losing models if those simulations could be continued. Because of this, our results are consistent with placing this cut-off at a $\log\left(\frac{L/L_{\odot}}{M/M_{\odot}}\right)$ value of between approximately 4.1 and 4.15.

In these simulations, a total mass of between 2 and 8 M_{\odot} is lost by the star due to repeated dynamical mass-ejection events, although in all cases this process continues




Figure 5.6: The distribution of the starting parameters of hydrodynamical simulations which display repeated dynamical mass ejections recurring over many pulsation cycles. Each point represents one such simulation, with the colour given by the logarithm of the time-averaged mass-loss rate calculated over the period containing repeated mass-loss events, excluding the initial amplitude peak. Simulations drawn from the same evolutionary sequence (in both cases, for a $40 M_{\odot}$ initial mass star) are joined by grey lines. Simulations which do not experience repeated mass ejections are shown as empty circles, while those simulations which terminate too early for their behaviour to be characterised are shown as empty squares.




Figure 5.7: The time-averaged mass-loss rates of the 10 simulations displaying repeated dynamical mass loss as a function of their initial luminosity to mass ratio. Time-averaged mass-loss rates are again calculated only over the period of the simulation in which repeated mass-loss events occur, excluding the initial pulsation amplitude peak. Plotted in blue is the best-fitting linear relationship (see text).

until the termination of the simulation (either because it reached 200 years or because of convergence failures), so more mass loss would certainly be expected if the simulations were extended. In one of our simulations, mass is ejected in every pulsation cycle for 54 successive pulsations, although each of these ejections consists of only a few hundredths of a solar mass. In others, such as the one shown in panel 4 of Fig. 5.2, mass ejections do not occur in every pulsation cycle, but are considerably larger in size, the largest exceeding $1 M_{\odot}$. It seems that our $40 M_{\odot}$ models tend to display more consistent but smaller mass ejections than models of lower mass.

This dynamical mass-loss regime is only seen in simulations based on our 25, 30, and $40 M_{\odot}$ models, and it occurs within a few thousand years of the expected supernova in the case of the 25 and $30 M_{\odot}$ models, and of order 100,000 years before this in the case of the $40 M_{\odot}$ star. This suggests that this mechanism, if it occurs in real stars, could produce strong pre-supernova mass loss in stars in the 25– $30 M_{\odot}$ range which would be likely to affect the supernova type by producing a massive shell of circumstellar material. For more massive stars, the onset of rapid dynamical mass loss appears to occur earlier in the star’s

nuclear evolution, after which we would expect a star, having lost its envelope, to evolve to the blue and enter a WR phase (see e.g. Meynet et al., 2015). The luminosity range in which dynamical mass loss seems to emerge begins at a $\log(L/L_{\odot})$ value of approximately 5.55, although the simulations which experience steady mass loss at the low end of this range seem to be the 25 and 30 M_{\odot} stars, with the 40 M_{\odot} stars entering this regime around a $\log(L/L_{\odot})$ value of 5.65. However, we must stress that these values are based on a small number of simulations and should therefore be considered uncertain. A better determination of the point of onset of this process could in theory be found using a denser set of hydrodynamical simulations, although the numerical stability problems are severe; our own original grid is roughly twice as dense in this area of parameter space, but roughly half of our simulations were terminated by unrecoverable convergence failures too soon to properly characterise their long-term behaviour.

In Figs 5.8, 5.9, and 5.10 we show in some detail a mass-ejection event experienced by a star with an initial mass of 25 M_{\odot} . This event is taken from a model with an α_{MLT} of 2.5, and is the third mass ejection experienced by that model, taking place after the initial amplitude peak has decayed. A notable difference from similar figures appearing in Chapters 3 and 4 is that these figures are plotted from a simulation in which the unbound-mass removal routine remains active – we found that the presence of unbound mass had too great an effect on the numerical stability of the model to deactivate this routine. In the mass ejection shown in these figures nearly 1.2 M_{\odot} is removed between ages of approximately 38.4 and 38.7 years. The onset of this mass removal produces the sudden decrease in surface radius at 38.4 years.

This mass ejection displays two compression shocks (most easily visible in the plot of dv/dr in Fig. 5.9), with one shock occurring early in the contraction phase of the pulsation preceding the mass ejection (it already exists at the start of these figures) and one occurring shortly before the point at which minimum radius is reached. The first shock is long-lived, moves inwards in both radius and mass, and does not appear to be directly involved in the mass-ejection process, but does have a dramatic effect on the helium partial-ionization zones. The second shock exists for a much smaller amount of time, approximately 0.1 years, and is confined to the outer 10% of the star’s mass, occurring just before the radius minimum. Interestingly, the hydrogen partial-ionization zone in this simulation is extremely

narrow in both radius and mass, and is very close to the surface (although recall that the surface of this simulation is at optical depth 20 not at the photosphere). This suggests that in these stars the helium partial-ionization zones may be responsible for driving the pulsation of these models rather than the hydrogen partial-ionization zone which is generally believed to be responsible for driving pulsations in AGB stars. This change may be the reason that we do not see dynamical pulsations in the early-RSG models of our $15 M_{\odot}$ stars despite these representing more extreme giants than some of our pulsating AGB models – these stars may be in a regime in which hydrogen has become inefficient at driving pulsations but helium has not yet become efficient. As can be seen from the plot of ϵ_{rec} in Fig. 5.9, significant ionization energy is released during the acceleration of the ejected shell, first from the recombination of doubly ionized helium, then at later times by the recombination of singly ionized helium. Hydrogen recombination does not seem to contribute significantly to accelerating material to escape velocity, but may be important in accelerating this material to still-higher velocities after it has achieved escape, however this cannot be seen in these figures because that material has been removed.

5.3 Discussion

Our results indicate that strong dynamical mass loss is expected to occur in RSGs as they approach the ends of their lives. Discounting the non-repeating mass ejections which are likely to be a result of the artificial pulsation-growth period at the start of our simulations, the regime of repeated mass ejections constitutes a very strong mass-loss mechanism, with time-averaged mass-loss rates as high as $0.1 M_{\odot} \text{ yr}^{-1}$. Whilst we are not in a position to simulate the switch between a steadily pulsating RSG and one undergoing dynamical mass ejections, it is possible that this may be relatively sudden, as a steady increase in the star’s pulsation amplitude may push it into the regime of dynamical mass loss, after which its L/M ratio will increase, which we find to be indicative of greater levels of instability against both dynamical pulsations and mass ejections. However, this may not necessarily translate to a sudden, dramatic increase in mass-loss rate, as there are several mass-loss mechanisms which may occur in stars that display high-amplitude pulsations but are not actually ejecting mass dynamically, for example pulsation-enhanced or pulsation-driven

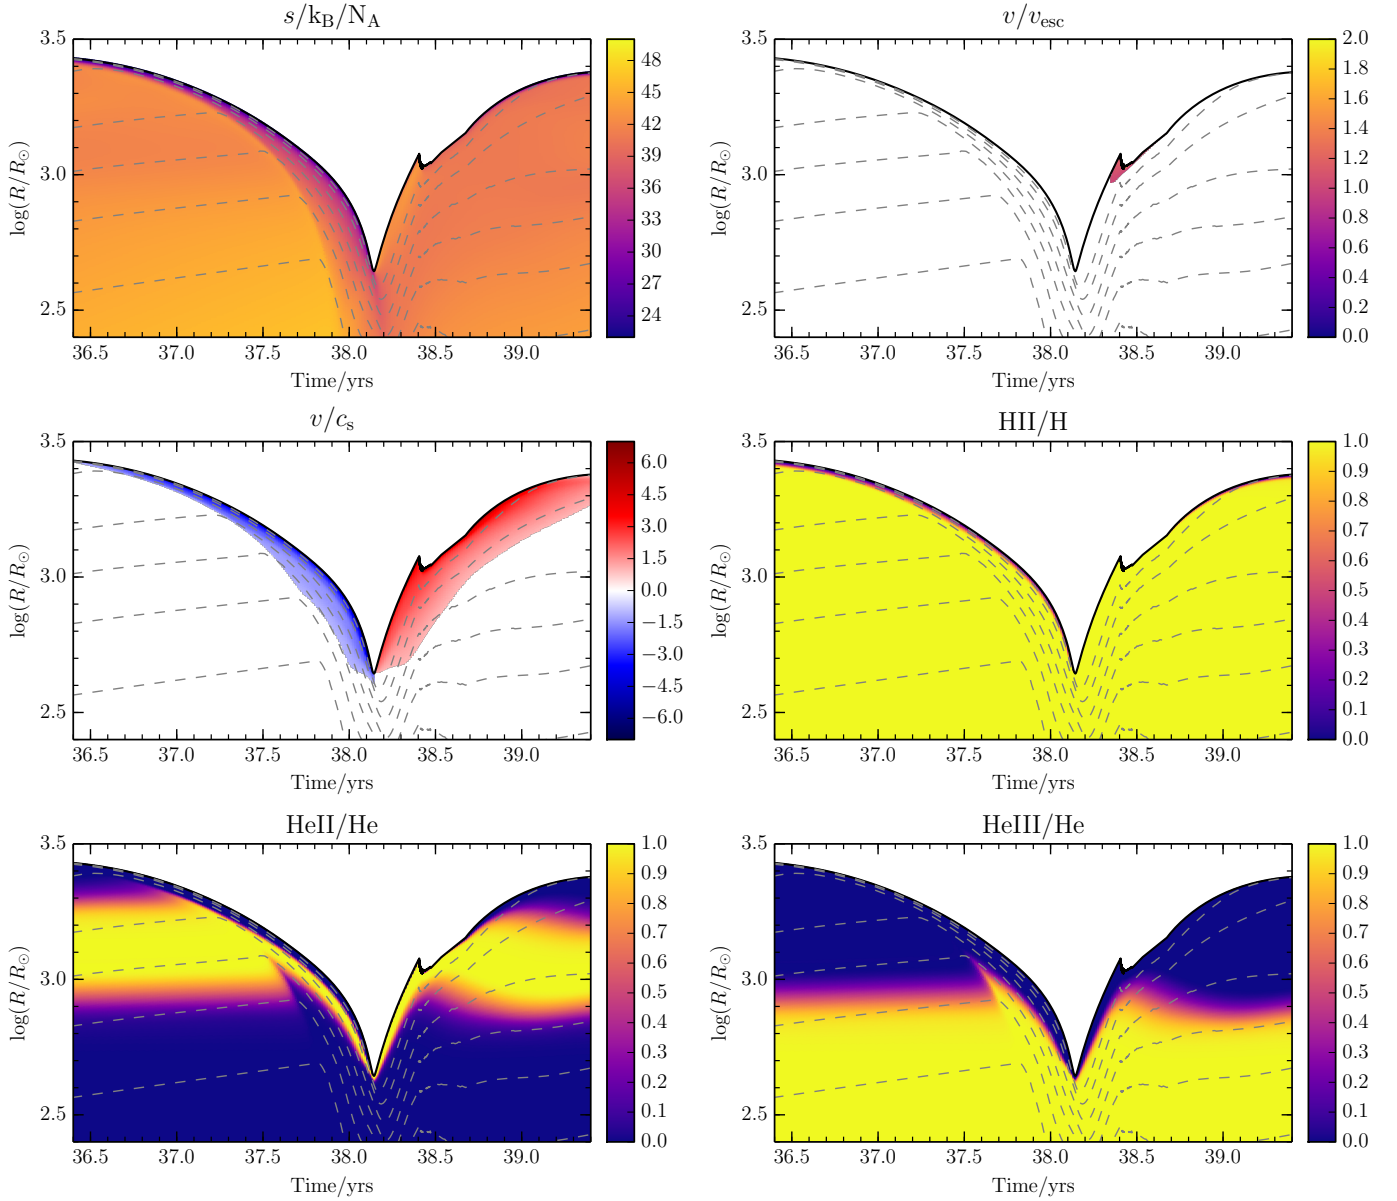


Figure 5.8: The third mass ejection experienced by a hydrodynamical simulation of a star with an initial mass of $25 M_{\odot}$ begun at age 7,320,463 years (the 8th and final simulation based on this star), with an α_{MLT} of 2.5. Our mass removal algorithm remained active for this figure. We show the non-dimensionalised entropy per particle; the ratio of fluid velocity to local escape velocity in those regions where this ratio exceeds 1; the ratio of fluid velocity to local sound speed for regions where the flow is supersonic; and the relative ionization fractions for hydrogen, and for singly and doubly ionized helium. Also shown are contours which contain 100% (black), and 98, 95, 90, 85, 80, and 75% (dashed grey) of the total mass of the model. Note that the total mass changes between 38.4 and 38.7 years as unbound matter is removed.

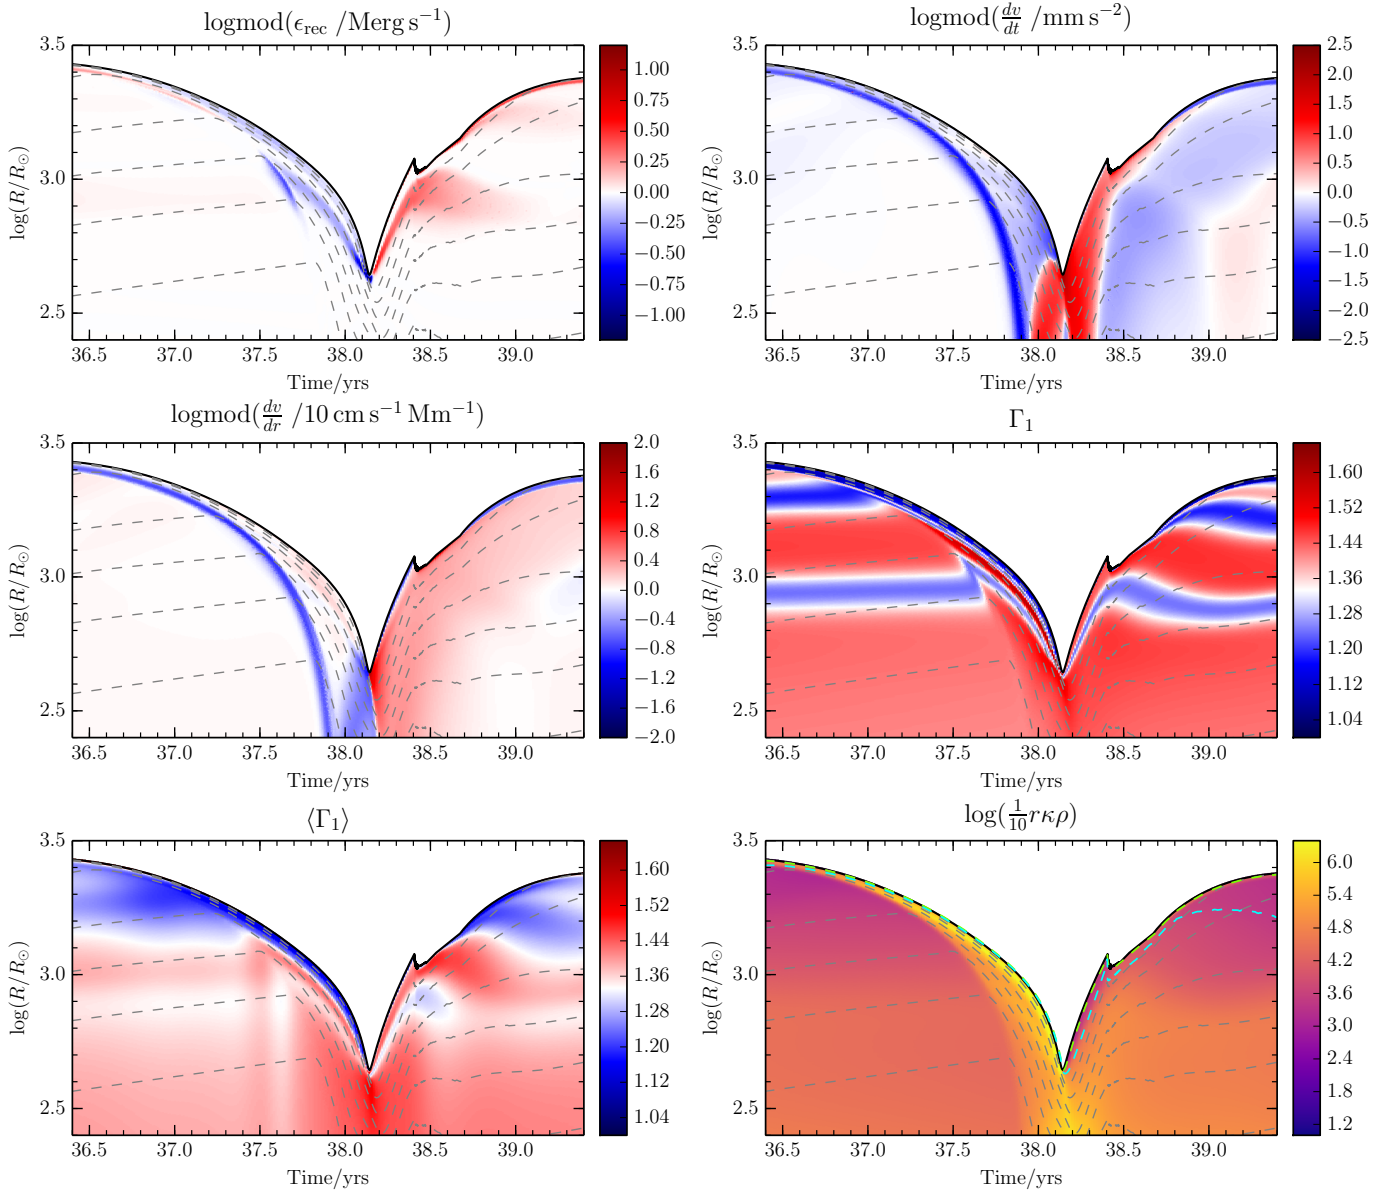


Figure 5.9: The same event that appears in Fig. 5.8. Our mass removal algorithm remained active for this figure. We show the specific rate of release of recombination energy; the Eulerian acceleration; the velocity divergence; the value of the first adiabatic exponent with white representing the critical value $4/3$; the pressure-weighted, volume averaged value of the first adiabatic exponent with white representing the critical value $4/3$; and one tenth of the product of local radius, opacity and density, a dimensionless quantity representative of the local optical thickness. Panel 6 also shows the radii at which the total optical depth is equal to 10^2 (dashed green), and 10^4 (dashed cyan). Also shown are contours which contain 100% (black), and 98, 95, 90, 85, 80, and 75% (dashed grey) of the total mass of the model. Note that the total mass changes between 38.4 and 38.7 years as unbound matter is removed.




Figure 5.10: The same event that appears in Fig. 5.8. Our mass removal algorithm remained active for this figure. We show the temperature; the density; the fraction of total pressure provided by radiation pressure; the sound speed; the opacity; and the pressure. Also shown are contours which contain 100% (black), and 98, 95, 90, 85, 80, and 75% (dashed grey) of the total mass of the model. Note that the total mass changes between 38.4 and 38.7 years as unbound matter is removed.

winds or pulsational Roche lobe overflow in binary systems – and we note that almost all massive stars are expected to appear in multiple systems (e.g. Kobulnicky & Fryer, 2007; Sana et al., 2012, 2013).

Our results suggest that dynamical mass loss may occur in RSGs whose luminosity exceeds a $\log(L/L_{\odot})$ of 5.5–5.7, with stars evolving to higher luminosities expected to undergo very strong mass loss that is likely to render them no longer visible as RSG stars. In the long term these stars are expected to lose all or almost all of their hydrogen envelopes and migrate to the blue region of the HR diagram (see e.g. Meynet et al., 2015). This is in good agreement with current estimates of the location of the Humphreys-Davidson limit, and dynamical mass loss may form a significant part of the mass-loss histories of RSGs in this luminosity range. This may be helpful in resolving the theoretical problem with providing sufficient mass loss to explain the recent lower luminosity estimates for the HD limit given by Davies & Beasor (2018).

The dynamical mass loss we see in our simulations is capable of ejecting very large amounts of material in short amounts of time, and for stars of initial masses 25 and 30 M_{\odot} this mass loss occurs within a few thousand years of core collapse. As such, dynamical mass loss seems to represent a candidate mechanism for the production of the massive CSM shells inferred to exist around the progenitors of IIIn supernovae (e.g. Schlegel, 1990; Chugai & Danziger, 1994; Filippenko, 1997). It will be necessary to perform more detailed simulation work to determine whether models incorporating this mass-loss mechanism can produce objects with the correct properties to be IIIn progenitors.

It is important that we stress the limitations of our simulations in the RSG regime. As in the case of CE and AGB stars, we are once again relying on parameterised prescriptions of convection which can only approximate the true nature of the process. We continue to utilise the mixing length theory, with its instantly-adjusting convective energy fluxes, although the pulsation periods of our models are again significantly longer than the convective turnover timescale, so the resulting behaviour of the system is likely to remain at least approximately correct. Experiments in RSG models performed by Heger et al. (1997) did not find that adding time delays to convective flux adjustments led to qualitatively different behaviour. In addition to issues with convective flux adjustment times, convective velocities in RSG envelopes may approach or exceed the speed of sound, whereas MLT assumes that at least

the relative velocities between different layers within a convective zone are small compared to the sound speed. We experimented with artificially limiting the convective velocity in our models but did not find that this led to any clear changes in our results. There is also evidence that the envelopes of RSGs are likely to deviate significantly from spherical symmetry due to the presence of large convective plumes (e.g. Freytag et al., 2002; Kervella et al., 2009; Chiavassa et al., 2010) which are likely to stochastically excite (nonradial) pulsations, and it is not clear what effect this will have on the radial pulsation properties of the star.

5.4 Summary

We performed hydrostatic and hydrodynamical simulations of RSG stars with initial masses of 15, 17.5, 20, 25, 30, and 40 M_{\odot} using 3 different values for α_{MLT} in order to ensure our models provided good coverage of the RSG temperature range. Of our 131 hydrodynamical simulations, 70 developed large-amplitude pulsations in the fundamental mode. We measured the pulsation growth rates of 69 of these simulations and examined the dependence of this quantity, a proxy for the level of instability in the envelope, on basic stellar properties.

In 41 of our hydrodynamical simulations, we see dynamical mass ejection events, although many of these events are associated with a likely-artificial pulsation amplitude peak and are therefore not considered reliable indicators of dynamical mass loss. Nevertheless, in 10 simulations we see repeated mass-loss events which continue after the star has thermally adjusted to the presence of the pulsation and do not appear to be related to the non-physical amplitude peak. These simulations lose between 2 and 8 M_{\odot} of material in less than 200 years, representing an extremely fast mass-loss mechanism with mass-loss rates as high as $0.1 M_{\odot} \text{ yr}^{-1}$. This strong dynamical mass loss occurs within a few thousand years of the expected supernova time for models with initial masses of 25 and 30 M_{\odot} , and of order 100,000 years before this for models with an initial mass of 40 M_{\odot} . These simulations occupy a region of the HR diagram with $\log(L/L_{\odot})$ of between 5.55 and 5.85. This is indicative of strong dynamical mass loss with onset late in the RSG phase that is consistent with current estimates of the location of the Humphreys-Davidson limit and able to produce massive shells of circumstellar material similar to those that are inferred to surround SN

In progenitors.

Chapter 6

Conclusion

We have examined the role which dynamical mass loss may play in the lives of three types of stellar objects: common-envelope objects, asymptotic giant branch stars, and red supergiants. Using hydrodynamical simulations of objects in these regimes, we have determined that this mass-loss mechanism is likely to appear in all of these classes of stars.

In our analysis of CE models we found that, for certain orbital energy dissipation rates, the expanded shared envelope becomes dynamically unstable and develops high-amplitude pulsations, which lead to repeated dynamical mass-loss events. These events eject shells of up to $0.1 M_{\odot}$, and repeat within a few decades, giving a time-averaged mass-loss rate of order $10^{-3} M_{\odot} \text{ yr}^{-1}$. For the $1.6 M_{\odot}$ Red Giant model we used, this is sufficient to remove the entire envelope within a reasonable duration of the slow spiral-in phase of 1,000 years. As such, dynamical mass loss in this regime should be considered a candidate mechanism for envelope ejection, and the termination of CE evolution, during this phase.

By relating the rate of orbital heating to the mass-loss rate, we were able to calculate a value of the ejection efficiency parameter α for this mechanism in the range 0.046–0.25, with the value of $\alpha\lambda$ for our giant model in the range 0.025–0.14. With estimates of these values, it may be possible to incorporate this envelope-removal mechanism into existing models of the CE phase used in binary population synthesis codes, although a few necessary elements are still missing (see below).

In the AGB regime, we found that for all of the stars we modelled, over a mass range of 1–6 M_{\odot} , their evolution along the TP-AGB brought them into a state of dynamical instability, with dynamical pulsations and then dynamical mass ejections appearing as their envelopes became larger and cooler and their cores more luminous. We found that the

dynamical properties of the envelope were heavily dependent on the thermal pulse cycle, with dynamical mass loss appearing earlier, and at higher rates, during the helium shell flashes than at other times. Simulated stars undergoing this process were seen to eject shells of matter in the $0.01\text{--}0.1 M_{\odot}$ range, on repeating intervals of between less than ten, and several hundred years, again giving time-averaged mass-loss rates of up to approximately $10^{-3} M_{\odot} \text{ yr}^{-1}$, high enough to remove the entire envelope within a handful of helium shell flashes, or during only one thermal pulse if experienced during the hydrogen burning phase. A detailed analysis of the physics underlying the pulsation cycle of these stars and the process of dynamical ejection revealed that this mechanism appears identical to that seen in our CE models.

We constructed parameterised models of the growth rates of dynamical pulsations and the rates of dynamical mass loss seen in our simulations of AGBs. By implementing this mass-loss model in MESA, we were able to perform full-lifetime evolutionary calculations that incorporated this mass-loss mechanism and compare the resulting models to the observed population of WD remnants, allowing us to test the predictions our model provides for the termination point of the AGB and therefore the masses of the WDs produced. We found a good level of agreement between our results and the observationally inferred location of the WD mass peak, a sensitive measure of the end of the AGB for lower-intermediate mass stars, despite the fact that our dynamical mass-loss model is a first-principles theoretical estimate that relies on no empirically calibrated parameters.

The presence of high levels of dynamical mass loss during the helium shell flash phase of the thermal pulse cycle suggests that stars in this phase may become heavily shrouded by ejected material and evolve into OH/IR stars, possibly intermitting between appearing as OH/IR stars and Miras as their pulsation and mass-loss properties change throughout the thermal pulse cycle. Furthermore, the fact that dynamical mass loss in our simulations occurs as a series of episodic events separated by a timescale of order decades to centuries suggests that it is a candidate mechanism for the formation of PN halo rings (Corradi et al., 2004; Phillips et al., 2009; Ramos-Larios et al., 2016) and for the discrete shells of material seen around certain AGB stars (Olofsson, 1996; Hashimoto et al., 1998; Schöier et al., 2005; Sanchez et al., 2015).

Our analysis of the RSG regime, based on models of stars with initial masses between 15

and $40 M_{\odot}$, also revealed large-amplitude pulsations and associated mass-ejection events. Although the mass ejections present in a number of our models were associated with unrealistic amplitude maxima, there did exist a regime of repeated mass-ejection events which was not associated with these phenomena. This regime appeared late in the lives of $25\text{--}40 M_{\odot}$ stars, producing time-averaged mass-loss rates of $10^{-3}\text{--}10^{-1} M_{\odot} \text{ yr}^{-1}$. This mass loss emerged a few thousand years before core collapse for 25 and $30 M_{\odot}$ stars, and is therefore expected to produce circumstellar shells of material of significant mass around these stars which may be present during the subsequent supernova, fulfilling the requirements for such structures inherent in SN IIn models (Schlegel, 1990; Chugai & Danziger, 1994; Filippenko, 1997). For stars with initial masses of $40 M_{\odot}$, we see the emergence of dynamical mass loss roughly 100,000 years before core collapse; the extremely high mass-loss rates provided by this mechanism are likely to strip the entire envelopes of these stars and cause them to evolve out of the RSG regime before their SNe, a result consistent with estimates of the Humphreys-Davidson limit (Davies & Beasor, 2018) and the lack of higher-mass RSGs seen as SN IIp progenitors (Smartt et al., 2009). We provide a simple model for estimating the rate of this mass loss based on a star’s L/M ratio that can in principle be implemented within evolutionary calculations, although more simulations will be required before this model can be considered reliable.

We found that dynamical instability was a feature common to all three of these regimes, and that it was in all cases associated with much of the same physics – pulsation-driven mass ejections associated with strong shocks in the envelopes. Although there are differences in the nature of the dominant mass ejection patterns across these regimes, we are able to identify general patterns, such as lower effective temperatures and higher L/M ratios leading to more unstable envelopes and stronger pulsation growth, and, in all cases, we found that dynamical mass loss is likely to contribute significantly to the mass-loss histories of all of these stars. The extremely high mass-loss rates that result from dynamical mass-ejection events are such that the lengths of time during which giants are actually undergoing this mass loss are expected to be very short, and the large amount of CSM produced may lead to stars not being visible as definitively AGB stars or RSG stars during this phase, but instead as OH/IR stars.

Despite the clearly dynamical nature of the instability seen in our simulations, we did not

find that the concept of formal dynamical instability, as determined using $\langle \Gamma_1(m) \rangle$, provides a very good determination of the boundary between stable and unstable models, although we did find that envelopes which were more unstable by this criterion were generally more unstable in our simulations. As such, although $\langle \Gamma_1(m) \rangle$ remains a useful quantity, it should not be considered the sole criterion for deciding whether or not an envelope is likely to undergo dynamical processes – the true physics governing the evolution of the envelope are more complex than the linear adiabatic theory from which this condition is derived. Nonlinear effects such as shocks and highly non-adiabatic cooling catastrophes play an important role in this evolution and can control the appearance of dynamical mass loss in unintuitive ways (such as by preventing mass loss when pulsation growth is too fast in CE models). Because of this, we cannot rely on simple dynamical stability conditions and must rely instead on empirical models fitted to grids of simulations.

The results we have discussed in this thesis suggest a number of avenues for future research that may be particularly informative. It is important that the properties of dynamically ejected mass shells after the ejection event are investigated so that their effects on the observable properties of their source giants can be understood. Simulations of these structures incorporating radiative-hydrodynamics treatments may be able to tell us how effectively they obscure their central giants and whether they would therefore appear as OH/IR stars, whether they can produce the halo ring structures seen in PNe (and in particular, whether their interaction with each-other produced a regularising effect as suggested by Ramos-Larios et al. 2016), and whether their interaction with supernova outflows can reproduce the spectra of IIIn supernovae. Simulations of these structures would also be useful in determining whether they could be observationally separated from shells potentially produced by the modulation of stellar winds.

More work studying the process of dynamical mass loss and its relationship to other processes affecting unstable giants is also required. In the case of common envelopes specifically, we can propose a plan for a more accurate version of the widely used α formalism that incorporates dynamical mass loss during the slow spiral-in phase: if we can find a criterion that determines whether a system will undergo envelope ejection in the fast plunge-in phase or whether it will instead enter the slow spiral-in phase, we can apply different values of α depending on which process is responsible for removing the

envelope. We have provided estimates of this value in this thesis for pulsation-driven mass loss during the slow spiral-in phase, while estimates for ejection during the plunge-in phase are provided by 3-d hydrodynamical simulations of that phase (e.g. De Marco et al., 2011; Nandez & Ivanova, 2016). The criterion for whether or not a CE system will enter the slow spiral-in phase may be determined by simulations of systems which undergo the transition into this phase (Ivanova & Nandez, 2016). Constructing such a system could significantly improve the quality of the CE treatments used in binary evolution and population synthesis calculations.

For the study of dynamical, pulsation-driven mass loss in general, it is important that we understand how radial pulsations interact with the multidimensional processes occurring within giant stars. In particular, the effects of large-scale convective motions that may stochastically excite nonradial pulsations and significantly deform the shape of the envelope should be investigated. In the case of CE models, the effects of rotation on the envelope and the effects of the differences in energy deposition due to the presence of the binary orbit within a specific plane should be studied, and it is possible that 2-dimensional simulations may be helpful in understanding this process. The recent interest in 2-dimensional stellar hydrodynamics and evolution codes (e.g. Goffrey et al., 2017; Halabi et al., 2017) is promising in this regard. There is also considerable scope for mass loss which, while not directly dynamical in nature, is driven by the presence of large-amplitude pulsations in the form of pulsation-driven and -enhanced winds and mass transfer via pulsational Roche lobe overflow (see Mohamed & Podsiadlowski, 2007; Dermine et al., 2009). The extent to which such mass loss may occur in stars undergoing large-scale pulsations should be considered in models of pulsation-driven dynamical mass loss.

Extensions of the techniques applied in this thesis can be made, both to improve the fidelity of the simulations and to extend them to other regimes. In terms of improved physics, an improvement that can be made in the CE case is to model the interaction between the binary orbit and the rate of energy release, rather than treating it as constant. By explicitly calculating the radius of the secondary's orbit and its evolution, the effects of changes of density in the orbital path can be taken into account in a self-consistent manner, although care must be taken that the chosen heat-dissipation model is appropriate, for there are many processes involved and it is not clear which are the most important in many cases

(see Chapter 3 and Ivanova 2002). In general, the simulation of dynamical pulsations and mass loss may be improved by implementing a time-dependent convection model and an explicit atmosphere simulation to better model the outer boundary of the star throughout the pulsation cycle. Existing nonlinear pulsation codes could be applied for this purpose, taking the place of MESA as the hydrodynamical arm of the dual evolutionary/hydrodynamical sampling process we have employed.

Finally, although we have studied the three regimes in which the possibility of dynamical mass loss seems the most promising, there remain other classes of stars to which the techniques used in this thesis might be applied. Although their properties are generally less extreme than in the stars we have studied, it is possible that dynamical instabilities may emerge in stars close to the tip of the RG branch also, where they may serve to enhance mass loss or lead directly to dynamical mass-ejection events. There is also a gap in mass between the AGB and RSG models we have studied in this work, and these stars may also experience dynamical mass loss during their super-AGB phase (Lau et al., 2012). The possibility of dynamical mass loss in stars with different metallicities should also be considered, and very massive cool giants in particular appear promising in this regard (see Moriya & Langer, 2015).

In summary, we must conclude that dynamical, pulsation-driven mass loss appears to be a vital part of the mass-loss histories of a wide range of unstable cool giant stars, and may be able to resolve a number of outstanding theoretical problems in this area. Rather than a marginal and dubious mechanism confined to simplistic simulations, as some authors had concluded in previous decades, we find that dynamical mass loss is likely to be a universal property of expanded, unstable giant envelopes; its appearance in simulations is robust against significant changes of both the included physics prescriptions and the stellar regime under study, and its implications for stellar evolution are consistent with modern theoretical and observational requirements. This mass-loss mechanism has important implications for stellar evolution theory in multiple classes of stars and has the potential to be a fertile and rewarding research area.

Appendix A

Mass-removal routine

The custom mass-removal scheme used in this work to excise dynamically ejected material is implemented in MESA using the `other_wind` hook routine, which is called at every timestep to find the appropriate amount of mass to remove from the star during that step.

The algorithm used to calculate this first calculates the local escape velocity for each simulation cell, and then finds the largest contiguous layer which includes the outermost cell in which all cells exceed their local escape velocity (if such a layer exists). The mass of this layer is the quantity m that appears in Equation 2.22, and mass loss is applied at a rate of $100m$ per year.

The algorithm is reproduced below using Fortran syntax.

```
! m(x) is the mass coordinate of cell x
! r(x) is the radius coordinate of cell x
! v(x) is the velocity of cell x
! G is Newton's gravitational constant
! msol is one solar mass
! cellnum is the total number of cells in the model
! mdot is the mass-loss rate to be calculated

limit_cell = 0

! Cells are numbered from 1 at the star's outer edge,
! to cellnum at the centre
do x = 1, cellnum

    v_esc = sqrt( 2 * G * m(x) / r(x) )

    if ( v(x) > v_esc ) then
        limit_cell = x
```

```
        else
            exit
        end if
    end do

    if ( limit_cell == 0 ) then
        mdot = 0
    else
        mdot = ( m(1) - m(limit+1) ) / msol * 100
        ! in solar masses per year
    endif
```

Bibliography

Akashi, M., & Soker, N. 2016, *MNRAS*, 462, 206

Alcock, C., Axelrod, T. S., Bennett, D. P., et al. 1992, in *Astronomical Society of the Pacific Conference Series*, Vol. 34, *Robotic Telescopes in the 1990s*, ed. A. V. Filippenko, 193–202

Appenzeller, I. 1970, *A&A*, 5, 355

Army, T. 1990, *Vistas in Astronomy*, 33, 211

Asida, S. M., & Tuchman, Y. 1997, *ApJ*, 491, L47

Auer, L. H., & Woolf, N. J. 1965, *ApJ*, 142, 182

Baud, B., & Habing, H. J. 1983, *A&A*, 127, 73

Beasor, E. R., & Davies, B. 2018, *MNRAS*, 475, 55

Bedijn, P. J. 1987, *A&A*, 186, 136

Belczynski, K., Kalogera, V., Rasio, F. A., et al. 2008, *ApJS*, 174, 223

Bennett, P. D. 2010, in *Astronomical Society of the Pacific Conference Series*, Vol. 425, *Hot and Cool: Bridging Gaps in Massive Star Evolution*, ed. C. Leitherer, P. D. Bennett, P. W. Morris, & J. T. Van Loon, 181

Bernasconi, P. A., & Maeder, A. 1996, *A&A*, 307, 829

Bessell, M. S., Scholz, M., & Wood, P. R. 1996, *A&A*, 307, 481

Biermann, L. 1938, *ZAp*, 16, 29

- Bladh, S., Höfner, S., Aringer, B., & Eriksson, K. 2015, *A&A*, 575, A105
- Bowers, P. F., & Knapp, G. R. 1987, *ApJ*, 315, 305
- Catalán, S., Isern, J., García-Berro, E., & Ribas, I. 2008, *MNRAS*, 387, 1693
- Chiavassa, A., Haubois, X., Young, J. S., et al. 2010, *A&A*, 515, A12
- Choi, J., Dotter, A., Conroy, C., et al. 2016, *ApJ*, 823, 102
- Chugai, N. N., Chevalier, R. A., & Lundqvist, P. 2004a, *MNRAS*, 355, 627
- Chugai, N. N., & Danziger, I. J. 1994, *MNRAS*, 268, 173
- Chugai, N. N., Blinnikov, S. I., Cumming, R. J., et al. 2004b, *MNRAS*, 352, 1213
- Corradi, R. L. M., Sánchez-Blázquez, P., Mellema, G., Gianmanco, C., & Schwarz, H. E. 2004, *A&A*, 417, 637
- Corradi, R. L. M., Steffen, M., Schönberner, D., & Jacob, R. 2007, *A&A*, 474, 529
- Corradi, R. L. M., Rodríguez-Gil, P., Jones, D., et al. 2014, *MNRAS*, 441, 2799
- Cowling, T. G. 1934, *MNRAS*, 94, 768
- Cox, J. P. 1980, *Theory of stellar pulsation* (Princeton University Press)
- Cox, J. P., & Giuli, R. T. 1968, *Principles of Stellar Structure*, 1st edn. (Gordon and Breach, New York)
- Darwin, G. H. 1879, *Proc. R. Soc. Lond.*, 29, 168
- Davies, B. 2017, *Philosophical Transactions of the Royal Society of London Series A*, 375, 20160270
- Davies, B., & Beasor, E. R. 2018, *MNRAS*, 474, 2116
- Davies, B., Crowther, P. A., & Beasor, E. R. 2018, *ArXiv e-print – accepted for publication in MNRAS*, arXiv:1804.06417
- de Jager, C., Nieuwenhuijzen, H., & van der Hucht, K. A. 1988, *A&AS*, 72, 259

- de Kool, M. 1990, *ApJ*, 358, 189
- De Marco, O. 2009, *PASP*, 121, 316
- De Marco, O., & Izzard, R. G. 2017, *PASA*, 34, e001
- De Marco, O., Passy, J.-C., Moe, M., et al. 2011, *MNRAS*, 411, 2277
- de Vries, B. L., Blommaert, J. A. D. L., Waters, L. B. F. M., et al. 2014, *A&A*, 561, A75
- Delgado, A. 1980, *A&A*, 87, 343
- Dermine, T., Jorissen, A., Siess, L., & Frankowski, A. 2009, *A&A*, 507, 891
- Deutsch, A. J. 1956, *ApJ*, 123, 210
- Dewi, J. D. M., & Tauris, T. M. 2000, *A&A*, 476, 63
- Dorfi, E. A. 1998, in *Saas-Fee Advanced Course 27: Computational Methods for Astrophysical Fluid Flow.*, ed. O. Steiner & A. Gautschy, 263
- Eddington, A. S. 1926, *The Internal Constitution of the Stars* (Cambridge University Press)
- Engels, D., & Bunzel, F. 2015, *A&A*, 582, A68
- Engels, D., & Jiménez-Esteban, F. 2007, *A&A*, 475, 941
- Engels, D., Kreysa, E., Schultz, G. V., & Sherwood, W. A. 1983, *A&A*, 124, 123
- Eriksson, K., Nowotny, W., Höfner, S., Aringer, B., & Wachter, A. 2014, *A&A*, 566, A95
- Fadeyev, Y. A. 2017a, *Astronomy Letters*, 43, 602
- Fadeyev, Y. A. 2017b, *Astronomy Letters*, 43, 690
- Ferguson, J. W., Alexander, D. R., Allard, F., et al. 2005, *ApJ*, 623, 585
- Ferraro, F. R., Valenti, E., Straniero, O., & Origlia, L. 2006, *ApJ*, 642, 225
- Filippenko, A. V. 1997, *ARA&A*, 35, 309
- Fleischer, A. J., Gauger, A., & Sedlmayr, E. 1995, *A&A*, 297, 543

- Fox, M. W., & Wood, P. R. 1982, *ApJ*, 259, 198
- Freytag, B., & Höfner, S. 2008, *A&A*, 483, 571
- Freytag, B., Liljegren, S., & Höfner, S. 2017, *A&A*, 600, A137
- Freytag, B., Steffen, M., & Dorch, B. 2002, *Astronomische Nachrichten*, 323, 213
- Gadun, A. S., Hanslmeier, A., Pikalov, K. N., et al. 2000, *A&AS*, 146, 267
- Gamow, G. 1938, *Physical Review*, 53, 595
- Gastine, T., & Dintrans, B. 2011, *A&A*, 530, L7
- Gautschy, A. 1999, *A&A*, 349, 209
- Georgy, C. 2012, *A&A*, 538, L8
- Glebbeeck, E., Gaburov, E., de Mink, S. E., Pols, O. R., & Portegies Zwart, S. F. 2009, *A&A*, 497, 255
- Goffrey, T., Pratt, J., Viallet, M., et al. 2017, *A&A*, 600, A7
- Groenewegen, M. A. T., Whitelock, P. A., Smith, C. H., & Kerschbaum, F. 1998, *MNRAS*, 293, 18
- Habing, H. J., Olsson, F. M., Winnberg, A., Matthews, H. E., & Baud, B. 1983, *A&A*, 128, 230
- Habing, H. J., & Olofsson, H. 2003, *Asymptotic giant branch stars* (Springer)
- Halabi, G. M., Izzard, R. G., Tout, C. A., Jermyn, A. S., & Cannon, R. 2017, *Mem. Soc. Astron. Italiana*, 88, 319
- Han, Z., Podsiadlowski, P., & Eggleton, P. P. 1994, *MNRAS*, 270, 121
- Han, Z., Podsiadlowski, P., & Eggleton, P. P. 1995, *MNRAS*, 272, 800
- Han, Z., Podsiadlowski, P., Maxted, P. F. L., & Marsh, T. R. 2003, *MNRAS*, 341, 669

- Han, Z., Podsiadlowski, P., Maxted, P. F. L., Marsh, T. R., & Ivanova, N. 2002, MNRAS, 336, 449
- Hashimoto, O., Izumiura, H., Kester, D. J. M., & Bontekoe, T. R. 1998, A&A, 329, 213
- Hayashi, C. 1949, Phys. Rev., 75, 1619
- Heger, A., Jeannin, L., Langer, N., & Baraffe, I. 1997, A&A, 327, 224
- Heger, A., Langer, N., & Woosley, S. E. 2000, ApJ, 528, 368
- Helmholtz, H. 1854, Kant Commemoration Lecture
- Henry, L., Vardya, M. S., & Bodenheimer, P. 1965, ApJ, 142, 841
- Henry, L. G., Forbes, J. E., & Gould, N. L. 1964, ApJ, 139, 306
- Henry, L. G., Lelevier, R., & Levée, R. D. 1955, PASP, 67, 154
- Herwig, F. 2005, ARA&A, 43, 435
- Hinkle, K. H., Hall, D. N. B., & Ridgway, S. T. 1982, ApJ, 252, 697
- Hofmeister, E., Kippenhahn, R., & Weigert, A. 1964, ZAp, 59, 215
- Höfner, S. 2015, in Astronomical Society of the Pacific Conference Series, Vol. 497, Why Galaxies Care about AGB Stars III: A Closer Look in Space and Time, ed. F. Kerschbaum, R. F. Wing, & J. Hron, 333
- Höfner, S., Feuchtinger, M. U., & Dorfi, E. A. 1995, A&A, 297, 815
- Höfner, S., Gautschi-Loidl, R., Aringer, B., & Jørgensen, U. G. 2003, A&A, 399, 589
- Höfner, S., & Olofsson, H. 2018, A&A Rev., 26, 1
- Humphreys, R. M. 1978, ApJS, 38, 309
- Humphreys, R. M., & Davidson, K. 1979, ApJ, 232, 409
- Hurley, J. R., Tout, C. A., & Pols, O. R. 2002, MNRAS, 329, 897

- Iaconi, R., Reichardt, T., Staff, J., et al. 2017, MNRAS, 464, 4028
- Iben, Jr., I. 1986, ApJ, 304, 201
- Iben, Jr., I., & Livio, M. 1993, PASP, 105, 1373
- Iben, Jr., I., & Renzini, A., eds. 1981, Astrophysics and Space Science Library, Vol. 88, Physical processes in red giants; Proceedings of the Second Workshop, Advanced School of Astronomy, Erice, Italy, September 3-13, 1980
- Iben, Jr., I., & Renzini, A. 1983, ARA&A, 21, 271
- Iben, I. J., & Tutukov, A. V. 1984, ApJS, 54, 335
- Icke, V., Frank, A., & Heske, A. 1992, A&A, 258, 341
- Iglesias, C. A., & Rogers, F. J. 1993, ApJ, 412, 752
- Iglesias, C. A., & Rogers, F. J. 1996, ApJ, 464, 943
- Ireland, M. J., Scholz, M., & Wood, P. R. 2008, MNRAS, 391, 1994
- Ita, Y., Tanabé, T., Matsunaga, N., et al. 2004, MNRAS, 347, 720
- Ivanova, N. 2002, PhD thesis, University of Oxford
- Ivanova, N. 2017, in IAU Symposium, Vol. 329, The Lives and Death-Throes of Massive Stars, ed. J. J. Eldridge, J. C. Bray, L. A. S. McClelland, & L. Xiao, 199–206
- Ivanova, N., Justham, S., Avendano Nandez, J. L., et al. 2013a, Science, 339, 433
- Ivanova, N., Justham, S., & Podsiadlowski, P. 2015, MNRAS, 447, 2181
- Ivanova, N., & Nandez, J. L. A. 2016, MNRAS, 462, 362
- Ivanova, N., & Podsiadlowski, P. 2003, in From Twilight to Highlight: The Physics of Supernovae, ed. W. Hillebrandt & B. Leibundgut, 19
- Ivanova, N., & Taam, R. E. 2004, ApJ, 601, 1058
- Ivanova, N., Justham, S., Chen, X., et al. 2013b, A&ARv, 21, 59

- Izzard, R. G., Hall, P. D., Tauris, T. M., & Tout, C. A. 2012, in IAU Symposium, Vol. 283, IAU Symposium, 95–102
- Jacobs, M. L., Porter, D. H., & Woodward, P. R. 1998, in Bulletin of the American Astronomical Society, Vol. 30, American Astronomical Society Meeting Abstracts, 1316
- Jeong, K. S., Winters, J. M., Le Bertre, T., & Sedlmayr, E. 2003, A&A, 407, 191
- Jiang, Y.-F., Cantiello, M., Bildsten, L., Quataert, E., & Blaes, O. 2015, ApJ, 813, 74
- John, J. a., & Draper, N. R. 1980, J. R. Statist. Soc. Series C (Applied Statistics), 29, 190
- Jones, T. J., Hyland, A. R., Wood, P. R., & Gatley, I. 1983, ApJ, 273, 669
- Jura, M., & Kleinmann, S. G. 1990, ApJS, 73, 769
- Justtanont, K., Skinner, C. J., Tielens, A. G. G. M., Meixner, M., & Baas, F. 1996, ApJ, 456, 337
- Justtanont, K., Teyssier, D., Barlow, M. J., et al. 2013, A&A, 556, A101
- Kalirai, J. S., Hansen, B. M. S., Kelson, D. D., et al. 2008, ApJ, 676, 594
- Karakas, A. I., & Lattanzio, J. C. 2014, PASA, 31, e030
- Kato, S. 1966, PASJ, 18, 374
- Keeley, D. A. 1970, ApJ, 161, 657
- Keene, J., Young, K., Phillips, T. G., Buettgenbach, T. H., & Carlstrom, J. E. 1993, ApJ, 415, L131
- Keller, S. C., & Wood, P. R. 2006, ApJ, 642, 834
- Kelvin, L. W. T. 1861, British Association for the Advancement of Science Reports, Part II, 27
- Kepler, S. O., Pelisoli, I., Koester, D., et al. 2015, MNRAS, 446, 4078
- Kepler, S. O., Pelisoli, I., Koester, D., et al. 2016, MNRAS, 455, 3413

- Kervella, P., Verhoelst, T., Ridgway, S. T., et al. 2009, *A&A*, 504, 115
- Kiewe, M., Gal-Yam, A., Arcavi, I., et al. 2012, *ApJ*, 744, 10
- Kilic, M., Brown, W. R., Allende Prieto, C., et al. 2011, *ApJ*, 727, 3
- Kim, H., & Taam, R. E. 2012, *ApJ*, 744, 136
- Kiminki, M. M., Reiter, M., & Smith, N. 2016, *MNRAS*, 463, 845
- Kinugawa, T., Inayoshi, K., Hotokezaka, K., Nakauchi, D., & Nakamura, T. 2014, *MNRAS*, 442, 2963
- Kippenhahn, R., Weigert, A., & Hofmeister, E. 1967, *Methods in Computational Physics*, 7, 129
- Kippenhahn, R., Weigert, A., & Weiss, A. 2012, *Stellar Structure and Evolution* (Springer), doi:10.1007/978-3-642-30304-3
- Kiss, L. L., Szabó, G. M., & Bedding, T. R. 2006, *MNRAS*, 372, 1721
- Kleinman, S. J., Kepler, S. O., Koester, D., et al. 2013, *ApJS*, 204, 5
- Knapp, G. R., & Morris, M. 1985, *ApJ*, 292, 640
- Kobulnicky, H. A., & Fryer, C. L. 2007, *ApJ*, 670, 747
- Koester, D., Schulz, H., & Weidemann, V. 1979, *A&A*, 76, 262
- Kolláth, Z., Buchler, J. R., Szabó, R., & Csubry, Z. 2002, *A&A*, 385, 932
- Kuhfuss, R. 1986, *A&A*, 160, 116
- Kutter, G. S., & Sparks, W. M. 1974, *ApJ*, 192, 447
- Kwok, S., Su, K. Y. L., & Hrivnak, B. J. 1998, *ApJ*, 501, L117
- Langer, G. E. 1971, *MNRAS*, 155, 199
- Langer, N., El Eid, M. F., & Fricke, K. J. 1985, *A&A*, 145, 179

- Lau, H. H. B., Gil-Pons, P., Doherty, C., & Lattanzio, J. 2012, *A&A*, 542, A1
- Lebzelter, T., & Wood, P. R. 2005, *A&A*, 441, 1117
- Ledoux, P. 1945, *ApJ*, 102, 143
- Ledoux, P. 1947, *ApJ*, 105, 305
- Ledoux, P., & Walraven, T. 1958, *Handbuch der Physik*, 51, 353
- Levesque, E. M. 2017, *Astrophysics of Red Supergiants* (IOP Publishing), doi:10.1088/978-0-7503-1329-2
- Levesque, E. M., Massey, P., Olsen, K. A. G., et al. 2005, *ApJ*, 628, 973
- Levesque, E. M., Massey, P., Olsen, K. A. G., et al. 2006, *ApJ*, 645, 1102
- Liljegren, S., Höfner, S., Nowotny, W., & Eriksson, K. 2016, *A&A*, 589, A130
- Limoges, M.-M., Bergeron, P., & Lépine, S. 2015, *ApJS*, 219, 19
- Livio, M., & Soker, N. 1988, *ApJ*, 329, 764
- Lucy, L. B. 1967, *AJ*, 72, 813
- MacLeod, M., Macias, P., Ramirez-Ruiz, E., et al. 2017, *ApJ*, 835, 282
- MacLeod, M., & Ramirez-Ruiz, E. 2015, *ApJ*, 803, 41
- Maercker, M., Mohamed, S., Vlemmings, W. H. T., et al. 2012, *Nature*, 490, 232
- Matsuura, M., Barlow, M. J., Zijlstra, A. A., et al. 2009, *MNRAS*, 396, 918
- Mattsson, L. 2016, *Mem. Soc. Astron. Italiana*, 87, 249
- Mauron, N., & Josselin, E. 2011, *A&A*, 526, A156
- McDonald, I., Zijlstra, A. A., Sloan, G. C., et al. 2016, *MNRAS*, 456, 4542
- Meyer, F., & Meyer-Hofmeister, E. 1979, *A&A*, 78, 167
- Meynet, G., Chomienne, V., Ekström, S., et al. 2015, *A&A*, 575, A60

- Miller Bertolami, M. M. 2016, *A&A*, 588, A25
- Mohamed, S., & Podsiadlowski, P. 2007, in 15th European Workshop on White Dwarfs
ASP Conference Series, Vol. 372, Leicester, 397
- Moriya, T. J., & Langer, N. 2015, *Astronomy & Astrophysics*, 573, A18
- Nandez, J. L. A., & Ivanova, N. 2016, *MNRAS*, 460, 3992
- Nandez, J. L. A., Ivanova, N., & Lombardi, J. C. 2015, *MNRAS*, 450, L39
- Noels, A., Grevesse, N., Magain, P., et al. 1991, *A&A*, 247, 91
- Nowotny, W., Höfner, S., & Aringer, B. 2010, *A&A*, 514, A35
- Ohlmann, S. T., Röpke, F. K., Pakmor, R., & Springel, V. 2016, *ApJ*, 816, L9
- Ohnaka, K., Weigelt, G., & Hofmann, K.-H. 2016, *A&A*, 589, A91
- Oke, J. B., & Schwarzschild, M. 1952, *ApJ*, 116, 317
- Olivier, E. A., & Wood, P. R. 2005, *MNRAS*, 362, 1396
- Olofsson, H. 1996, *Ap&SS*, 245, 169
- Olofsson, H., Carlstrom, U., Eriksson, K., Gustafsson, B., & Willson, L. A. 1990, *A&A*,
230, L13
- Ostlie, D. A., & Cox, A. N. 1986, *ApJ*, 311, 864
- Owocki, S. 2010, in *Astronomical Society of the Pacific Conference Series*, Vol. 425, *Hot
and Cool: Bridging Gaps in Massive Star Evolution*, ed. C. Leitherer, P. D. Bennett, P. W.
Morris, & J. T. Van Loon, 199
- Paczynski, B. 1975, *ApJ*, 202, 558
- Paczyński, B. 1976, in *Structure and Evolution of Close Binary Systems*, Vol. 73 (IAU
Symposium No. 73), 75
- Paczyński, B., & Ziółkowski, J. 1968, *Acta Astron.*, 18, 255

- Palla, F., & Stahler, S. W. 1993, *ApJ*, 418, 414
- Passy, J.-C., De Marco, O., Fryer, C. L., et al. 2012, *ApJ*, 744, 52
- Paxton, B., Bildsten, L., Dotter, A., et al. 2011, *ApJS*, 192, 3
- Paxton, B., Cantiello, M., Arras, P., et al. 2013, *ApJS*, 208, 4
- Paxton, B., Marchant, P., Schwab, J., et al. 2015, *ApJS*, 220, 15
- Paxton, B., Schwab, J., Bauer, E. B., et al. 2018, *ApJS*, 234, 34
- Payne, C. H. 1925, PhD thesis, Radcliffe College, Cambridge
- Phillips, J. P., Ramos-Larios, G., Schröder, K.-P., & Contreras, J. L. V. 2009, *MNRAS*, 399, 1126
- Podsiadlowski, P. 2001, *Evolution of Binary and Multiple Star Systems; A Meeting in Celebration of Peter Eggleton's 60th Birthday*. ASP Conference Series, 229, 239
- Podsiadlowski, P., Ivanova, N., Justham, S., & Rappaport, S. 2010, *MNRAS*, 406, 840
- Podsiadlowski, P., & Mohamed, S. 2007, *Baltic Astronomy*, 16, 26
- Podsiadlowski, P., Rappaport, S., & Pfahl, E. D. 2002, *ApJ*, 565, 1107
- Portegies Zwart, S. F., & Meinen, A. T. 1993, *A&A*, 280, 174
- Potekhin, A. Y., & Chabrier, G. 2010, *Contributions to Plasma Physics*, 50, 82
- Prandtl, L. 1925, *Zeitschrift für angewandte Mathematik und Mechanik*, 5, 136
- Puls, J., Vink, J. S., & Najarro, F. 2008, *A&A Rev.*, 16, 209
- Ramos-Larios, G., Santamaría, E., Guerrero, M. A., et al. 2016, *MNRAS*, 462, 610
- Ramstedt, S., Schöier, F. L., & Olofsson, H. 2009, *A&A*, 499, 515
- Reimers, D. 1975, *Memoires of the Societe Royale des Sciences de Liege*, 8, 369
- Renzini, A., & Voli, M. 1981, *A&A*, 94, 175

Ricker, P. M., & Taam, R. E. 2012, ApJ, 746, 74

Riebel, D., Srinivasan, S., Sargent, B., & Meixner, M. 2012, ApJ, 753, 71

Ritter, A. 1898, ApJ, 8, 293

Robinson, F. J., Demarque, P., Li, L. H., et al. 2004, MNRAS, 347, 1208

Roche, É. 1849, Académie des sciences de Montpellier: Mémoires de la section des sciences

Rogers, F. J., & Nayfonov, A. 2002, ApJ, 576, 1064

Romero, A. D., Campos, F., & Kepler, S. O. 2015, MNRAS, 450, 3708

Rose, W. K. 1966, ApJ, 146, 838

Rose, W. K. 1967, ApJ, 150, 193

Roxburgh, I. W. 1967, Nature, 215, 838

Salpeter, E. E. 1955, ApJ, 121, 161

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

Sana, H., de Koter, A., de Mink, S. E., et al. 2013, A&A, 550, A107

Sanchez, E., Montez, Jr., R., Ramstedt, S., & Stassun, K. G. 2015, ApJ, 798, L39

Sandage, A. R., & Schwarzschild, M. 1952, ApJ, 116, 463

Saumon, D., Chabrier, G., & van Horn, H. M. 1995, ApJS, 99, 713

Schlegel, E. M. 1990, MNRAS, 244, 269

Schöier, F. L., Lindqvist, M., & Olofsson, H. 2005, A&A, 436, 633

Schöier, F. L., & Olofsson, H. 2001, A&A, 368, 969

Scholz, M., & Wood, P. R. 2000, A&A, 362, 1065

Schröder, K.-P., & Cuntz, M. 2005, ApJ, 630, L73

Schröder, K.-P., & Cuntz, M. 2007, *A&A*, 465, 593

Schröder, K.-P., Winters, J. M., Arndt, T. U., & Sedlmayr, E. 1998, *A&A*, 335, L9

Schwarzschild, M. 1958, *Structure and evolution of the stars*. (Dover)

Schwarzschild, M., & Härm, R. 1965, *ApJ*, 142, 855

Shiber, S., Kashi, A., & Soker, N. 2017, *MNRAS*, 465, L54

Shklovsky, I. S. 1956, *AZh*, 33, 315

Smartt, S. J. 2015, *PASA*, 32, e016

Smartt, S. J., Eldridge, J. J., Crockett, R. M., & Maund, J. R. 2009, *MNRAS*, 395, 1409

Smartt, S. J., Maund, J. R., Hendry, M. A., et al. 2004, *Science*, 303, 499

Smith, N. 2014, *ARA&A*, 52, 487

Smith, R. L., & Rose, W. K. 1972, *ApJ*, 176, 395

Smolec, R. 2016, *MNRAS*, 456, 3475

Soker, N. 1996, *ApJ*, 460, L53

Soker, N. 2004, *New A*, 9, 399

Soszyński, I. 2007, *ApJ*, 660, 1486

Soszyński, I., Udalski, A., Kubiak, M., et al. 2004, *Acta Astron.*, 54, 129

Soszyński, I., & Wood, P. R. 2013, *ApJ*, 763, 103

Srinivasan, S., Meixner, M., Leitherer, C., et al. 2009, *AJ*, 137, 4810

Staff, J. E., De Marco, O., Macdonald, D., et al. 2016, *MNRAS*, 455, 3511

Stahler, S. W., Shu, F. H., & Taam, R. E. 1980, *ApJ*, 241, 637

Steffen, M., Ludwig, H.-G., & Kruss, A. 1989, *A&A*, 213, 371

- Stein, R. F., & Nordlund, Å. 2000, *Sol. Phys.*, 192, 91
- Stellingwerf, R. F. 1982, *ApJ*, 262, 330
- Stewart, P. N., Tuthill, P. G., Monnier, J. D., et al. 2016, *MNRAS*, 455, 3102
- Stothers, R. B. 1999, *MNRAS*, 305, 365
- Stry, P. E. 1975, *ApJ*, 196, 559
- Sugimoto, D., Nomoto, K., & Eriguchi, Y. 1981, *Progress of Theoretical Physics Supplement*, 70, 115
- Taam, R. E., Bodenheimer, P., & Ostriker, J. P. 1978, *ApJ*, 222, 269
- Taam, R. E., Bodenheimer, P., & Rozyczka, M. 1994, *ApJ*, 431, 247
- Taam, R. E., & Sandquist, E. L. 2000, *ARA&A*, 38, 113
- Tielens, A. G. G. M. 2005, *The Physics and Chemistry of the Interstellar Medium* (Cambridge University Press)
- Timmes, F. X., & Swesty, F. D. 2000, *ApJS*, 126, 501
- Trampedach, R., Stein, R. F., Christensen-Dalsgaard, J., Nordlund, Å., & Asplund, M. 2014, *MNRAS*, 445, 4366
- Tremblay, P.-E., Ludwig, H.-G., Steffen, M., Bergeron, P., & Freytag, B. 2011, *A&A*, 531, L19
- Tremblay, P.-E., Ludwig, H.-G., Steffen, M., & Freytag, B. 2013, *A&A*, 559, A104
- Tuchman, Y., Sack, N., & Barkat, Z. 1978, *ApJ*, 219, 183
- Tuchman, Y., Sack, N., & Barkat, Z. 1979, *ApJ*, 234, 217
- Tylenda, R., Hajduk, M., Kamiński, T., et al. 2011, *A&A*, 528, A114
- Unno, W., Osaki, Y., Ando, H., Saio, H., & Shibahashi, H. 1989, *Nonradial oscillations of stars* (University of Tokyo Press)

- Uttenthaler, S., van Stiphout, K., Voet, K., et al. 2011, A&A, 531, A88
- van Loon, J. T., Cioni, M.-R. L., Zijlstra, A. A., & Loup, C. 2005, A&A, 438, 273
- van Marle, A. J., Smith, N., Owocki, S. P., & van Veelen, B. 2010, MNRAS, 407, 2305
- Vassiliadis, E., & Wood, P. R. 1993, ApJ, 413, 641
- Villaume, A., Conroy, C., & Johnson, B. D. 2015, ApJ, 806, 82
- Vink, J. S., de Koter, A., & Lamers, H. J. G. L. M. 2001, A&A, 369, 574
- Vögler, A. 2004, A&A, 421, 755
- Von Neumann, J., & Richtmyer, R. D. 1950, Journal of Applied Physics, 21, 232
- Voss, R., & Tauris, T. M. 2003, MNRAS, 342, 1169
- Wagenhuber, J., & Weiss, A. 1994, A&A, 290, 807
- Weaver, T. A., Zimmerman, G. B., & Woosley, S. E. 1978, ApJ, 225, 1021
- Webbink, R. F. 1984, ApJ, 277, 355
- Weidemann, V. 2000, A&A, 363, 647
- Weiler, K. W., van Dyk, S. D., Pringle, J. E., & Panagia, N. 1992, ApJ, 399, 672
- Weiss, A., Hillebrandt, W., Thomas, H.-C., & Ritter, H. 2004, Cox and Giuli's Principles of Stellar Structure (Cambridge Scientific Publishers)
- Williams, K. A. 2007, in Astronomical Society of the Pacific Conference Series, Vol. 372, 15th European Workshop on White Dwarfs, ed. R. Napiwotzki & M. R. Burleigh, 85
- Willson, L. A. 2000, ARA&A, 38, 573
- Willson, L. A., Wallerstein, G., & Pilachowski, C. A. 1982, MNRAS, 198, 483
- Winters, J. M., Le Bertre, T., Jeong, K. S., Helling, C., & Sedlmayr, E. 2000, A&A, 361, 641

- Woitke, P. 2006, *A&A*, 452, 537
- Wood, P. R. 1974, *ApJ*, 190, 609
- Wood, P. R. 1975, in *BAAS*, Vol. 7, *Bulletin of the American Astronomical Society*, 527
- Wood, P. R. 2006, *Mem. Soc. Astron. Italiana*, 77, 76
- Wood, P. R. 2007, in *IAU Symposium*, Vol. 239, *Convection in Astrophysics*, ed. F. Kupka, I. Roxburgh, & K. L. Chan, 343–348
- Wood, P. R. 2015, *MNRAS*, 448, 3829
- Wood, P. R., & Faulkner, D. J. 1986, *ApJ*, 307, 659
- Wood, P. R., Whiteoak, J. B., Hughes, S. M. G., et al. 1992, *ApJ*, 397, 552
- Wood, P. R., Alcock, C., Allsman, R. A., et al. 1999, in *IAU Symposium*, Vol. 191, *Asymptotic Giant Branch Stars*, ed. T. Le Bertre, A. Lebre, & C. Waelkens, 151
- Xiong, D.-R. 1989, *A&A*, 209, 126
- Ya'ari, A., & Tuchman, Y. 1996, *ApJ*, 456, 350
- Ya'ari, A., & Tuchman, Y. 1999, *ApJ*, 514, L35
- Yoon, S.-C., & Cantiello, M. 2010, *ApJ*, 717, L62