

White Dwarfs as Probes of Novel Physical Phenomena

by

Leesa Fleury

B.Sc., The University of Western Ontario, 2013
M.Sc., McGill University, 2016

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF
THE REQUIREMENTS FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

in

The Faculty of Graduate and Postdoctoral Studies
(Physics)

THE UNIVERSITY OF BRITISH COLUMBIA

(Vancouver)

December 2024

© Leesa Fleury 2024

The following individuals certify that they have read, and recommend to the Faculty of Graduate and Postdoctoral Studies for acceptance, the thesis entitled:

White Dwarfs as Probes of Novel Physical Phenomena

submitted by Leesa Fleury in partial fulfillment of the requirements for the degree of Doctor of Philosophy
in Physics

Examining Committee:

Jeremy Heyl, Professor, Physics and Astronomy, UBC
Supervisor

Gary Hinshaw, Professor, Physics and Astronomy, UBC
Supervisory Committee Member

Mona Berciu, Professor, Physics and Astronomy, UBC
Supervisory Committee Member

Ingrid Stairs, Professor, Physics and Astronomy, UBC
University Examiner

John Madden, Professor, Electrical and Computer Engineering, UBC
University Examiner

Georg Raffelt, Senior Scientist, Max Planck Institute for Physics
External Examiner

Additional Supervisory Committee Members:

David Morrissey, Research Scientist, TRIUMF
Supervisory Committee Member

Harvey Richer, Professor Emeritus, Physics and Astronomy, UBC
Supervisory Committee Member

Abstract

White dwarfs are the final state of most stars once nuclear burning in the core has finished, leaving a remnant that evolves through a straightforward cooling process that is largely thought to be well-understood. The electron-degenerate core of a white dwarf is an extreme environment that provides a testbed of physics in regimes not achievable with ground-based experiments. In this dissertation, observations of white dwarfs are used to test white dwarf cooling theory and look for evidence of novel physical phenomena. The cooling and kinematics of ultramassive white dwarfs in the solar neighbourhood are analysed using Gaia EDR3 observations to reinvestigate an anomalous cooling delay previously reported based on earlier Gaia DR2 observations, which challenged the conventional understanding of white dwarf cooling. The cooling of white dwarfs in the globular cluster 47 Tucanae is also analysed in detail by comparing cooling models to HST data to both test the implementation of element diffusion in stellar evolution simulations and determine the values of parameters important for modelling white dwarf cooling like the white dwarf mass and envelope thickness. A thorough understanding of these properties enables the cooling of white dwarfs to be used to indirectly search for evidence of novel particles such as axions. The emission of axions produced in the core of a white dwarf via axion bremsstrahlung would provide an additional energy loss mechanism and thus affect the cooling rate. A new bound on the axion-electron coupling of $g_{aee} \leq 0.81 \times 10^{-13}$ is derived from the cooling of white dwarfs in 47 Tucanae. This improves upon previous constraints by nearly a factor of two and excludes the range of values favoured by the hints of axions suggested by galactic white dwarf luminosity functions. Axions could also be produced in the core of a very hot, magnetic white dwarf like ZTF J1901+1458 via the

Abstract

^{57}Fe transition and then convert to photons in the magnetosphere, and it is shown that X-ray observations of ZTF J1901+1458 by NuSTAR to search for the corresponding signal could probe the coupling of axions to nucleons with better sensitivity than both current and future planned ground-based observations.

Lay Summary

White dwarfs are the final state of most stars once they have finished burning nuclear fuel in their cores and are left to simply cool. A cooling anomaly reported for some unusually massive white dwarfs is re-investigated using improved data to validate conventional white dwarf cooling models. The cooling of a population of more typical white dwarfs is also analysed to determine certain properties important for modelling white dwarf cooling and to look for evidence of new particles such as axions, which are hypothesised particles that could affect the cooling rate if they are emitted from white dwarfs. A new constraint is found on the possible strength of the interaction between axions and electrons. Finally, calculations are performed to show that observations of the X-ray radiation emitted by a particular very hot, magnetic white dwarf could also be used to look for evidence of axions interacting with protons and neutrons.

Preface

The research program for this dissertation was developed by my supervisor, Jeremy Heyl (J. H.), and me (the author, Leesa Fleury), and this program was carried out primarily by me. All of the work in this dissertation was done in collaboration with J. H. Some of this work was also done in collaboration with Ilaria Caiazzo (I. C.).

Parts of Chapter 3 follow the treatment of similar material in my master's thesis [1]: L. Fleury, Simulations of axion string-wall networks, Master's thesis, McGill University, Montreal, Canada (2015). Where this is done, I have written new text for this dissertation such that the writing contained in this chapter is original.

Chapter 4 has been adapted from the paper [2]: L. Fleury, I. Caiazzo, and J. Heyl, The cooling of massive white dwarfs from *Gaia* EDR3, MNRAS 511, 5984 (2022). A few paragraphs of text were contributed by my co-authors; the bulk of the text was written by me.

Chapter 5 has been adapted from the paper [3]: L. Fleury, I. Caiazzo, and J. Heyl, The origin of ultramassive white dwarfs: hints from *Gaia* EDR3, MNRAS 520, 364 (2023). Some code for this work was created by J. H.

Chapter 7 uses a module to calculate axion bremsstrahlung emission rates written by Alyssa Obertas.

Chapter 8 has been adapted from the paper [4]: L. Fleury, I. Caiazzo, and J. Heyl, Constraining axions with ZTF J1901+1458, Phys. Rev. D 107, L101303 (2023). This work expanded upon an idea presented in a telescope proposal submitted to *NuSTAR* by the same authors [5]: *NuSTAR* Cycle 8 proposal GO-8204 (PI: I. Caiazzo). The fitting procedure to determine the core temperature and produce Fig. 8.1 was performed by I. C. Some of the text was written by J. H.; the majority of the text was written by me.

Preface

All other parts of the dissertation not mentioned above consist of original, unpublished work by me.

Table of Contents

Abstract	iii
Lay Summary	v
Preface	vi
Table of Contents	viii
List of Tables	xiii
List of Figures	xiv
1 Introduction	1
1.1 White Dwarfs	1
1.2 White Dwarf Evolution	3
1.3 Magnitudes	5
1.3.1 Magnitude Systems	5
1.3.2 Reddening	8
1.4 Overview	9
2 White dwarf cooling	14
2.1 Basic Cooling Theory	14
2.1.1 Energy Transfer Through Surface Layers	15
2.1.2 Cooling of Interior	18
2.2 Neutrino Cooling	20
2.3 Axion Cooling	21
2.4 Convective Coupling	22
2.5 Core Crystallisation	23
2.6 White Dwarf Luminosity Function	26

Table of Contents

3 Axions	30
3.1 The QCD Axion	30
3.2 The Strong CP Problem and Solution	30
3.3 Axion Models	35
3.4 Axion Interactions Important for Stellar Evolution	36
3.4.1 Overview	36
3.4.2 Primakoff Effect	37
3.4.3 Axion Bremsstrahlung	40
3.5 Hints of Axions from White Dwarf Cooling	44
4 Ultramassive white dwarf cooling ages	46
4.1 Introduction	46
4.2 Methods	48
4.2.1 Models	48
4.2.2 Data	50
4.3 Results	55
4.4 Conclusions	72
5 Ultramassive white dwarf kinematics	74
5.1 Introduction	74
5.2 Methods	78
5.2.1 Data Sample	78
5.2.2 Cooling and Atmosphere Models	79
5.2.3 Kinematics	81
5.2.3.1 Streaming Motion	81
5.2.3.2 Random Transverse Motion	83
5.3 Results	85
5.3.1 Summary Statistics	85
5.3.2 Three-Dimensional Velocity Distributions	91
5.4 Discussion	96
5.5 Conclusions	99

Table of Contents

6	The cooling of old white dwarfs in 47 Tucanae	101
6.1	Introduction	101
6.2	Data	103
6.3	Artificial Stars Tests	105
6.4	Data Cleaning Procedures	107
6.4.1	Overview	107
6.4.2	SHARP	107
6.4.3	Proper Motion	115
6.4.3.1	Proper Motion Distribution Model	116
6.4.3.2	SMC Contamination	122
6.4.3.3	Completeness Reduction	130
6.5	Models	133
6.5.1	Creation of Initial Model	133
6.5.2	White Dwarf Cooling Models	134
6.6	Birthrate	139
6.7	Unbinned Likelihood Analysis	147
6.8	Results	155
6.8.1	Likelihood Distribution	155
6.8.2	Best-Fitting Model	159
6.9	Discussion	163
6.10	Conclusions	165
7	Axion constraints from white dwarfs in 47 Tucanae	168
7.1	Introduction	168
7.2	Data	171
7.3	Models	171
7.3.1	MESA Simulations	171
7.3.2	Cooling Curves	173
7.4	Analysis	176
7.5	Results	185
7.5.1	Posterior Distributions	186
7.5.2	Best-Fitting Model	191
7.6	Discussion	197

Table of Contents

7.7	Conclusions	199
8	Prospective axion search using ZTF J1901+1458	201
8.1	Introduction	201
8.2	Calculations	204
8.3	Results	211
8.4	Conclusions	214
9	Conclusions	215
	Bibliography	228

Appendices

A	Supplemental material for the cooling of old white dwarfs in 47 Tucanae	263
A.1	SMC Contamination	263
A.2	Element Diffusion	265
A.2.1	Diffusion in MESA	266
A.2.2	Modified Diffusion	272
B	Supplemental material for axion constraints from white dwarfs in 47 Tucanae	274
B.1	Overview	274
B.2	WFC3/UVIS Data Only	277
B.2.1	Posterior Distributions	277
B.2.1.1	Case 1	277
B.2.1.2	Case 2	280
B.2.1.3	Case 3	283
B.2.2	Best-Fitting Models Comparison	286
B.3	ACS/WFC Data Only	290
B.3.1	Posterior Distributions	290
B.3.1.1	Case 1	290
B.3.1.2	Case 2	293
B.3.1.3	Case 3	296

Table of Contents

B.3.2 Best-Fitting Models Comparison	299
B.4 Comparison of WFC3, ACS, and Combined Analyses	303

List of Tables

1.1	Summary of parameters characterising passbands	6
3.1	Summary of stellar bounds on the axion-electron coupling . .	43
4.1	The p-values from one-sample KS tests	64
4.2	Fitting results for cooling age distribution	69
6.1	Statistics for the SHARP distributions	113
6.2	Results of fitting proper motion distribution	120
6.3	Results of calibrating proper motion cleaning procedure . .	129
6.4	Best-fitting parameter values for piece-wise linear model . .	131
6.5	Maximum likelihood estimates of parameters	159
6.6	Results of KS tests	162
7.1	Conversion between axion mass values	174
7.2	Summary of model parameter values held fixed in analysis .	179
7.3	Parameter values of optimal model for combined analysis .	192
7.4	Results of KS tests for combined analysis	197
8.1	Parameter values used to calculate the X-ray flux	210
B.1	Summary of priors used in each case	275
B.2	Parameter values of optimal model for WFC3 cases	288
B.3	Results of KS tests for WFC3 cases	288
B.4	Parameter values of optimal model for ACS cases	300
B.5	Results of KS tests for ACS cases	300
B.6	Comparison of 95% credible regions	304
B.7	Comparison of best-fitting model parameters	304
B.8	Comparison of KS test results	305

List of Figures

4.1	The <i>Gaia</i> CMD of white dwarf candidates within 200 pc	52
4.2	Completeness limits of white dwarf samples	54
4.3	Joint distribution of mass and cooling age	57
4.4	Distribution of cooling ages	58
4.5	Cumulative distributions of cooling ages	62
4.6	Distribution of white dwarf cooling ages by mass bin	67
4.7	Ultramassive white dwarf formation mechanisms	71
5.1	Fraction of fast-moving massive white dwarfs	86
5.2	Quantile-quantile plots of transverse velocity components . .	93
5.3	AVRs of projected and un-projected velocities	95
6.1	CMDs showing boundaries and effects of cleaning	109
6.2	Distribution of SHARP for 47 Tuc main-sequence stars	112
6.3	Proper motions of SHARP -cleaned data	126
6.4	Completeness reduction factor f_{CR} as a function of F606W .	132
6.5	Theoretical cooling curves for different model parameters .	138
6.6	HST field boundaries over Gaia EDR3 data	141
6.7	CMD selections of RGB stars for birthrate calculations	143
6.8	Data space used in unbinned likelihood analysis	152
6.9	Likelihood locally maximised with respect to birthrate	157
6.10	Joint posterior probability density distribution	158
6.11	Inverse cumulative luminosity function for F606W	160
6.12	Inverse cumulative luminosity function for F814W	161
7.1	Theoretical cooling curves varying axion mass and envelope .	175
7.2	CMDs showing the data space selections	181

List of Figures

7.3	Joint posterior probability density distribution	187
7.4	Two-dimensional joint credible regions	188
7.5	One-dimensional posterior probability density distributions .	190
7.6	Cumulative number distribution of radial distance	193
7.7	Inverse cumulative luminosity function of optimal model . . .	194
8.1	Constraints on the core temperature of ZTF J1901+1458 . .	206
8.2	Projected sensitivity	213
B.1	WFC3, case 1: joint posterior probability density	277
B.2	WFC3, case 1: two-dimensional joint credible regions	278
B.3	WFC3, case 1: one-dimensional posterior density	279
B.4	WFC3, case 2: joint posterior probability density	280
B.5	WFC3, case 2: two-dimensional joint credible regions	281
B.6	WFC3, case 2: one-dimensional posterior density	282
B.7	WFC3, case 3: joint posterior probability density	283
B.8	WFC3, case 3: two-dimensional joint credible regions	284
B.9	WFC3, case 3: one-dimensional posterior density	285
B.10	WFC3, all cases: inverse cumulative luminosity function . . .	286
B.11	WFC3, all cases: cumulative number distribution of R . . .	287
B.12	ACS, case 1: joint posterior probability density	290
B.13	ACS, case 1: two-dimensional joint credible regions	291
B.14	ACS, case 1: one-dimensional posterior density	292
B.15	ACS, case 2: joint posterior probability density	293
B.16	ACS, case 2: two-dimensional joint credible regions	294
B.17	ACS, case 2: one-dimensional posterior density	295
B.18	ACS, case 3: joint posterior probability density	296
B.19	ACS, case 3: two-dimensional joint credible regions	297
B.20	ACS, case 3: one-dimensional posterior density	298
B.21	ACS, all cases: inverse cumulative luminosity function	299

Chapter 1

Introduction

1.1 White Dwarfs

White dwarfs are the remnants of low mass stars that have run out of nuclear fuel to burn in their cores. They are the final state of evolution for most stars in the Universe, with over 98% of stars destined to eventually become white dwarfs [6], including the Sun. Since a white dwarf does not generate energy through nuclear burning in its core, its fate is ultimately to continually cool over time. However, white dwarfs are initially very hot when they are born, and as a white dwarf evolves, it is supported against gravitational collapse by the electron degeneracy pressure of its very dense core. These extreme conditions of white dwarf interiors can lead to the emergence of physical phenomena not seen in other environments and make white dwarfs good testbeds for new physics.

A star will become a white dwarf if its initial mass is less than $8 M_{\odot}$, and even stars with initial masses as large as $11 M_{\odot}$ may become white dwarfs [7, 8]. Progenitor stars that later become white dwarfs lose mass at earlier stages of evolution due to stellar winds and in the planetary nebula ejection that produces the white dwarf. A progenitor star with an initial mass of $\sim 1 M_{\odot}$ typically produces a white dwarf with a mass of $\sim 0.6 M_{\odot}$ due to this mass loss [9]. A mass of $\sim 0.6 M_{\odot}$ ($\sim 0.5 - 0.7 M_{\odot}$) [9–14] is typical for a white dwarf in the solar neighbourhood, though white dwarfs can have masses as large as the Chandrasekhar limit of $\sim 1.4 M_{\odot}$. The Chandrasekhar limit is the maximum possible mass a white dwarf can have; for masses above this limit, electron degeneracy pressure is not sufficient to sustain the white dwarf against gravitational collapse. The mass distribution of white dwarfs in the solar neighbourhood is strongly peaked at $M_{\text{WD}} \sim$

1.1. White Dwarfs

$0.6 M_{\odot}$ but also has a broad shoulder that extends to higher masses with a small secondary peak at about $M_{\text{WD}} \sim 0.8 M_{\odot}$ [15]. White dwarfs with mass $M_{\text{WD}} \gtrsim 0.90 M_{\odot}$ are considered to be massive, and those with even larger masses $M_{\text{WD}} \gtrsim 1.05 M_{\odot}$ are classified as ultramassive white dwarfs.

Stars that are too massive to become white dwarfs reach high enough temperatures in their cores to ignite carbon burning, i.e. the fusion of two ^{12}C nuclei, a process which initiates runaway nuclear reactions that produce heavier elements up to iron and ultimately results in the production of a neutron star or black hole [16, 17]. The less massive stars that become white dwarfs only reach core temperatures hot enough to burn He, and the main products of this He burning are ^{12}C and ^{16}O . The final ratio of ^{12}C to ^{16}O depends on the rate of the α -capture process $^{12}\text{C}(\alpha, \gamma)^{16}\text{O}$ in which ^{12}C fuses with a ^4He nucleus (i.e. an alpha particle, α) to produce ^{16}O and a photon γ . The rate of this reaction is one of the biggest sources of uncertainty in nuclear astrophysics and leads to uncertainty in the ratio of ^{12}C to ^{16}O in white dwarf cores [18–20]. For the most massive stars that are still light enough to form white dwarfs, the core can reach high enough temperatures to also produce ^{22}Ne through the reaction $^{18}\text{O}(\alpha, \gamma)^{22}\text{Ne}$, i.e. α -capture onto the neutron-enriched ^{18}O isotope of oxygen, and some of this neon can further be converted to magnesium [16]. However, most stars that ultimately become white dwarfs do not become hot enough to form these heavier elements in appreciable amounts.

The structure of a white dwarf consists of a core surrounded by layers of different elements. A typical white dwarf consists of a core composed primarily of carbon and oxygen, with the core covered by a shell of helium, which is further covered by an outer envelope of hydrogen. White dwarfs with masses below $\sim 0.4 M_{\odot}$ have helium cores [9], as their progenitors do not reach hot enough temperatures to ignite He burning. Though the Universe is not yet old enough for a single progenitor of this mass to have become a white dwarf, helium-core white dwarfs can be produced in a binary system where rapid mass transfer strips an initially more massive progenitor down to a mass below $0.4 M_{\odot}$. On the other hand, white dwarfs with masses above $\sim 1.05 – 1.15 M_{\odot}$ can have cores composed primarily of oxygen

1.2. White Dwarf Evolution

and neon (and some magnesium) [21–24], as their progenitors can reach hot enough temperatures to convert oxygen to neon. However, most white dwarfs lie in the mass range where carbon-oxygen cores are expected.

Different atmospheric compositions are also possible for white dwarfs. The elements present in the atmosphere of a white dwarf can be determined using spectroscopy to look for the spectral lines associated with relevant elements, and a white dwarf is given a spectral classification based on which elements are present in its atmosphere. The classification “DA” denotes a H-dominated atmosphere (having only H absorption lines in its spectrum), while the classification “DB” denotes a He-dominated atmosphere (having only He I absorption lines in its spectrum). DA is the most common spectral type, followed by DB, and most white dwarfs have one of these two spectral types as their primary spectral classification [9].

1.2 White Dwarf Evolution

While stars in earlier stages of evolution are sustained against gravitational collapse by the outward pressure of nuclear burning, white dwarfs are prevented from completely collapsing under gravity by the electron degeneracy pressure of their dense cores. As a white dwarf does not generate pressure from nuclear burning to counteract gravity, the radius of a white dwarf initially decreases with time as the star cools, and the density and pressure in the core increase, until the electrons in the core become sufficiently degenerate that the white dwarf is sustained against collapse by electron degeneracy pressure.

As white dwarfs are the end state of stars that have finished nuclear burning in their cores, having exhausted all of the available fuel, the evolution of white dwarfs largely consists of a straightforward cooling process that is thought to be well understood. White dwarfs are known to cool by the emission of photons, and they also cool by the emission of neutrinos at very early cooling times. The luminosity of a white dwarf correspondingly decreases over time as the white dwarf cools, and this relation of white dwarf luminosity as a function of cooling time is called a cooling curve. This re-

1.2. White Dwarf Evolution

lation is sensitive to the interior state and energy loss mechanisms of the white dwarf and can be modelled through white dwarf cooling theory. A key tool to test this cooling theory using observations of a population of white dwarfs is the white dwarf luminosity function, which is the number density distribution per unit luminosity interval as a function of luminosity; the fraction of white dwarfs expected to have luminosities within a given range depends on the time it takes white dwarfs to cool through that luminosity range.

Aside from the brief periods of shell burning that may be ignited in young or accreting white dwarfs, which occur over time periods that are negligible compared to the typical cooling ages of white dwarfs, the luminosity of a white dwarf monotonically decreases with time. However, the rate at which the luminosity decreases depends on what stage of evolution the white dwarf is in, which manifests as the slope of the cooling curve. The rate of change of the luminosity differs for cooling dominated by neutrino emission compared to cooling dominated by photon emission, and the rate of cooling can differ from both of these standard cooling regimes if cooling is dominated by the emission of exotic particles such as axions.

As white dwarfs cool over time, they also undergo phase transitions in their core, which affect the rate of cooling. White dwarf cores transition first from a gaseous to liquid stage at early cooling times, then from a liquid to solid crystalline stage at later times. The latter transition of the core from liquid to solid is also referred to as core crystallisation or freezing. The release of latent heat [25], as well as other energy associated with element sedimentation [26–29], during the process of core crystallisation slows down the rate of cooling, resulting in a shallower slope of the cooling curve during this transition. This slowed cooling rate produces a corresponding pile-up of white dwarfs that can be seen in the white dwarf luminosity function [30, 31].

1.3 Magnitudes

1.3.1 Magnitude Systems

Observationally, the luminosity of an object is measured in terms of quantities called magnitudes, which are relative measures of the flux of a source. For a monochromatic flux at wavelength λ , the magnitude m_λ is given with respect to a reference source by the expression

$$m_\lambda - m_{\lambda,\text{ref}} = -2.5 \log_{10} \left(\frac{F_\lambda}{F_{\lambda,\text{ref}}} \right), \quad (1.1)$$

where F_λ is the energy flux of the source for which m_λ is the magnitude, $F_{\lambda,\text{ref}}$ is the energy flux of the reference source, and $m_{\lambda,\text{ref}}$ is the magnitude of the reference source. The star Vega is commonly used as a reference source due to its brightness, with $m_{\lambda,\text{ref}}$ taken to be zero by convention; this defines the Vega magnitude system.

In practice, magnitudes are measured using filters that only allow light spanning a particular range of wavelengths to pass through. A set of particular passbands define a photometric system. A commonly used standardised photometric system is the Johnson-Cousins UBVRI photometric system, which defines magnitudes for light in the ultraviolet (U), blue (B), visible (V), red (R), and infrared (I). Passbands for the UBVRI system are provided by Bessell [32]. For a detailed review of this photometric system that covers the various early contributions to its development, see e.g. Bessell [33].

This dissertation uses photometric data from two telescopes, *Gaia* and the Hubble Space Telescope (HST), which each define their own photometric systems. The *Gaia* data consists of observations from *Gaia* early data release 3 (EDR3) measured using three passbands: a blue passband G_{BP} , a red passband G_{RP} , and a passband G that spans both the blue and red passbands [34]. For the HST data used in this dissertation, six different filters are relevant. Two of these filters (F225W and F336W) are part of

1.3. Magnitudes

System	Passband	λ_{mean}	λ_{eff}	λ_{min}	λ_{max}	W_{eff}
UBVRI	U	360.51	365.99	303.12	417.37	64.04
	B	441.31	438.07	363.33	549.71	95.92
	V	551.21	544.54	473.33	687.50	89.31
	R	658.59	641.15	550.44	883.33	159.10
	I	805.99	798.21	704.17	916.67	149.51
Gaia	G	671.96	582.24	330.91	1,038.17	405.30
	G_{BP}	531.99	503.58	330.18	673.93	215.75
	G_{RP}	793.91	762.00	620.05	1,046.56	292.44
HST	F225W	239.14	237.28	199.00	300.57	46.26
	F336W	336.57	335.90	301.59	370.82	51.24
	F435W	436.01	434.16	361.02	488.38	82.17
	F555W	539.76	533.18	458.43	620.87	109.87
	F606W	603.57	580.93	463.43	718.01	177.19
	F814W	812.92	797.34	686.96	963.20	188.87

Table 1.1: Summary of parameters characterising passbands for different photometric systems (all values given in units of nm).

the WFC3/UVIS¹ instrument, while the other four filters (F435W, F555W, F606W, and F814W) are part of the ACS/WFC instrument. In the HST naming convention, the letter “F” at the start of each bandpass name denotes that it is a filter (as opposed to a grism, which would start with “G”), the three-digit number in the middle of the name gives the nominal effective wavelength of the bandpass in nm (for WFC3/UVIS and ACS/WFC), and the letter “W” at the end of each name denotes a wide bandpass².

Some key parameters (wavelengths and effective width) characterising the relevant passbands for both the *Gaia* and HST photometric systems are summarised in Table 1.1, along with the parameters of the UBVRI passbands provided by Bessell [32] for comparison. All of the values reported in Table 1.1 were retrieved from the Spanish Virtual Observatory (SVO)

¹The UVIS detector is comprised of two chips (called UVIS1 and UVIS2) with slightly different passbands. The values in Table 1.1 are given specifically for UVIS1, though the values for UVIS2 are similar.

²See e.g. the online documentation at <https://hst-docs.stsci.edu/wfc3ihb/chapter-6-uvvis-imaging-with-wfc3/6-5-uvvis-spectral-elements>.

1.3. Magnitudes

Filter Profile Service (FPS) website³ [35, 36]. The parameters given in Table 1.1 are defined as follows.

Let $T(\lambda)$ be the filter transmission, which gives the probability of light with wavelength λ passing through the filter such that the total flux that passes through is $F = \int T(\lambda) F_\lambda d\lambda$. The mean wavelength λ_{mean} is

$$\lambda_{\text{mean}} \equiv \frac{\int \lambda T(\lambda) d\lambda}{\int T(\lambda) d\lambda}, \quad (1.2)$$

and the effective wavelength λ_{eff}

$$\lambda_{\text{eff}} \equiv \frac{\int \lambda T(\lambda) Vg(\lambda) d\lambda}{\int T(\lambda) Vg(\lambda) d\lambda}, \quad (1.3)$$

where $Vg(\lambda)$ is the spectrum of Vega. λ_{mean} and λ_{eff} are both measures of the centre of the passband. The wavelengths λ_{min} and λ_{max} are the minimum and maximum wavelengths, respectively, with a transmission equal to at least 1% of the maximum transmission; these wavelength values are the edges of the passband. Finally, the effective width W_{eff} of the passband is given by

$$W_{\text{eff}} \equiv \frac{\int T(\lambda) d\lambda}{\text{Max}[T(\lambda)]}, \quad (1.4)$$

where $\text{Max}[T(\lambda)]$ is the maximum transmission. W_{eff} is equivalent to the horizontal size of a rectangle with height equal to maximum transmission and with the same area as that covered by the filter transmission curve.

The magnitude measured at the detector is called the apparent magnitude. By convention, the apparent magnitude measured using a particular bandpass filter is simply denoted by the name of that filter, and each telescope has its own naming convention for filters, as described above for *Gaia* and HST. A magnitude can also be defined in terms of the total flux emitted at all wavelengths without applying any bandpass filters (i.e. the total bolometric flux, $F_{\text{bol}} = \int F_\lambda d\lambda$); this is called the bolometric magnitude.

The value of the apparent magnitude depends on the distance between the detector and the source, but it is often useful to discuss magnitudes at

³The SVO FPS is available at <http://svo2.cab.inta-csic.es/theory/fps/>.

1.3. Magnitudes

a fixed distance from the source. This is accomplished by working in terms of the absolute magnitude, which is defined as the apparent magnitude that would be measured for a source at a distance of 10 pc. The absolute magnitude M_λ is related to the apparent magnitude m_λ through the expression

$$m_\lambda - M_\lambda = 5 \log_{10} \left(\frac{d}{10 \text{ pc}} \right) + A_\lambda, \quad (1.5)$$

where d is the distance between to the source and A_λ is the extinction due to the absorption and scattering of light by interstellar dust.

1.3.2 Reddening

As light travels from the source to the detector, it is scattered and absorbed by the dust and gas in the interstellar medium it passes through. Since bluer wavelengths are more strongly attenuated than redder wavelengths, this phenomenon results in photometric colours becoming redder and is called interstellar reddening. The amount of reddening can be quantified in terms of the colour excess $E(B - V)$, which is the difference between the observed and intrinsic values of the $B - V$ colour (in the Johnson photometric system),

$$E(B - V) = (B - V) - (B - V)_0, \quad (1.6)$$

where $(B - V)$ is the value observed by the detector and $(B - V)_0$ is the value the $B - V$ colour index would have if it were not affected by extinction. This colour excess is explicitly expressed in terms of the total extinctions A_B and A_V in the B and V bands, respectively, as

$$E(B - V) = A_B - A_V. \quad (1.7)$$

The extinction A_λ at wavelength λ is given by an extinction curve, which is typically expressed in terms of the absolute extinction A_λ/A_V ⁴. The general shape of this extinction curve is characterised by the relative visibility R_V

⁴Extinction curves can also be given in terms of the quantity $E(\lambda - V)/E(B - V)$, where $E(\lambda - V) = A_\lambda - A_V$ (see e.g. [37]).

[38–40], which compares the total V -band extinction to the colour excess,

$$R_V = \frac{A_V}{E(B - V)}, \quad (1.8)$$

and has a typical value of $R_V = 3.1$ for sources in the Milky Way [41].

1.4 Overview

This dissertation investigates various ways that observations of white dwarfs can be used both as tests of our current understanding of white dwarf cooling and as probes of new physics. The extreme environments of white dwarfs can lead to the emergence of physical phenomena not seen elsewhere, either on Earth or in other astrophysical environments. Ultramassive white dwarfs in particular push the limits of the extreme densities and temperatures that a white dwarf can have, and recent observations of ultramassive white dwarfs have shown evidence that some of them may experience anomalous cooling behaviour during the period of white dwarf crystallisation. Even the more common lower mass white dwarfs exist in density and temperature regimes that cannot be achieved on Earth, and thus observations of white dwarfs can be used to search for evidence of physical phenomena that cannot be seen in laboratory experiments. White dwarfs have in particular been identified as good targets to search for evidence of new particles beyond the Standard Model of particle physics, such as axions and axion-like particles. The emission of axions produced in the core of a white dwarf provides an additional method of cooling that can alter the cooling rate at a measurable level for sufficiently strong axion couplings. If white dwarf cooling is well understood, then observations of the white dwarf luminosity function in comparison to model predictions can be used to search for evidence of axions or other exotic physics. Signatures of axion emission can also be searched for in the photon emission spectrum of certain individual white dwarfs.

More detailed background information on the topics of both white dwarf cooling and axions is given in Chapters 2 and 3. These chapters provide information that is useful for understanding the research presented in later

1.4. Overview

chapters. White dwarf cooling is explained in detail in Chapter 2, covering the key phenomena that are important for modelling the cooling of white dwarfs. Chapter 2 explains the elementary theory of white dwarf cooling, important modifications to this elementary cooling theory, and the relationship between white dwarf cooling and the white dwarf luminosity function. Chapter 3 provides detailed information about axions. It covers the motivation for introducing such particles, the key interactions possible in models of these particles, and how these interactions can affect stellar astrophysics, particularly for white dwarfs. Chapter 3 also provides current bounds on axion couplings from astrophysical observations.

Though white dwarf cooling is generally thought to be a well understood process, some anomalous features reported for ultramassive white dwarfs in the solar neighbourhood, based on observations from the second data release (DR2) of the *Gaia* mission [42], have called into question conventional understanding of white dwarf core crystallisation and the accuracy of current white dwarf cooling models [43]. These anomalous features include a purported excess in the number of young ultramassive white dwarfs and a small fraction of ultramassive white dwarfs moving faster than expected. In Chapters 4 and 5 of this dissertation, these anomalous features of ultramassive white dwarfs are re-investigated using the improved observations of *Gaia* EDR3 [34]. Chapter 4 examines the distribution of cooling ages for massive and ultramassive white dwarfs, while Chapter 5 examines their kinematic properties.

Accurately modelling white dwarf cooling is a particularly important task when searching for evidence of axions or other exotic physics using the white dwarf luminosity function, as anomalies in the cooling behaviour of white dwarfs are precisely the evidence sought. In particular, the increased rate of cooling due to the emission of axions would correspond to a deficit of objects and altered slope in the white dwarf luminosity function over a particular luminosity range that can be predicted using cooling models that account for the extra cooling due to the emission of axions. Evidence of additional cooling compatible with axion emission has been reported for the white dwarf luminosity functions constructed from observations [44–46]

1.4. Overview

of white dwarfs in i) the thin and thick discs and ii) the halo, and it has been suggested that this hints towards the existence of axions that couple to electrons and have a mass of up to $\sim 4 - 10$ meV [47–49]. Chapters 6 and 7 investigate the potential existence of such axions in more rigorous detail by analysing white dwarf cooling in the globular cluster 47 Tucanae (47 Tuc).

A globular cluster is a spherically-distributed group of stars that are tightly bound together by gravity [50]. Globular clusters contain populations of very old stars that were all initially created from the same protostellar material at the same time in a burst of star formation⁵. The main difference between stars in a globular cluster population is simply the initial mass each star is born with. This primarily affects the length of time the star spends on the main sequence burning H in its core, with more massive stars finishing this stage of evolution more quickly. All subsequent stages of evolution for stars in globular clusters are so fast that all but the faintest white dwarfs in the population have approximately the same mass [51, 52]. The more massive white dwarfs lie fainter than those studied in this work. This is in contrast to stars in the solar neighbourhood (i.e. a population of stars in the Galactic disc), for which star formation is a continually ongoing process and the white dwarfs can have a wide range of different masses.

Thus, in contrast to white dwarfs in the Galactic disc or halo, globular clusters provide populations of white dwarfs with fixed values for some key parameters, such as white dwarf mass and distance, that are important when comparing observed cooling behaviour to the predictions of white dwarf cooling models. This makes globular clusters good environments for studying white dwarf cooling, and 47 Tuc is a particularly good target for this work because it is well populated and has been extensively observed. However, while many key parameters have been well quantified for 47 Tuc, a few important parameters are less well constrained, most notably the typical thickness of the white dwarf H envelope. In Chapter 6, the parameters characterising white dwarf cooling in 47 Tuc are determined more precisely by fitting cooling models to observations of old 47 Tuc white dwarfs in a

⁵Though many globular clusters contain multiple populations, each of these populations formed in an isolated star formation event.

1.4. Overview

cooling regime where axion emission would no longer have an observable effect. The results of Chapter 6 are then used in the analysis of Chapter 7 to search for evidence of axions. In Chapter 7, the predictions of cooling models that account for the emission of axions produced by axion bremsstrahlung from electrons colliding with ions are compared to observations of young white dwarfs in the cooling regime that would be dominated by axion emission. In the absence of evidence for extra cooling due to axion emission, the statistical analysis of Chapter 7 provides new constraints on the axion-electron coupling.

Evidence of axions can also be searched for in the emission spectra of white dwarfs, and such searches have been used to constrain the coupling of axions to electrons (and photons). In particular, X-ray observations of the hot, magnetic, white dwarf RE J0317-853 have been used to search for the signature of axions that were produced by axion bremsstrahlung in the core of the white dwarf and then converted to photons in the surrounding magnetosphere [53, 54]. Chapter 8 of this dissertation examines the potential for X-ray observations of hot, magnetic white dwarfs to also probe the coupling of axions to nucleons (and photons) by searching for the signal associated with an analogous mechanism in which axions are produced in the core of the white dwarf by the decay of the first excited state of ^{57}Fe and then convert to photons in the magnetic field surrounding the white dwarf. Though axions produced by this ^{57}Fe transition in the Sun are often the target of helioscope searches [55, 56], white dwarfs have not previously been recognised as viable targets to search for evidence of axions produced in this way. In Chapter 8, the ultramassive white dwarf ZTF J1901+1458 [57] is identified as a particularly good target for indirect searches of such axions, and the projected sensitivity is calculated for a 100 ks observation of this white dwarf.

In summary, the rest of this dissertation is structured as follows. Chapters 2 and 3 provide more detailed background information that is useful for understanding the research presented in later chapters. White dwarf cooling is covered in Chapter 2, while axion and axion-like particle models are covered in Chapter 3. Chapters 4 and 5 present research on ultramassive white

1.4. Overview

dwarfs in the solar neighbourhood investigating evidence for an anomalous cooling delay, with Chapter 4 focusing on the distribution of cooling ages and Chapter 5 focusing on the kinematics. Chapters 6 and 7 present research about white dwarf cooling in 47 Tuc. The work of Chapter 6 uses observations of old white dwarfs to fit for parameters important for modelling white dwarf cooling in 47 Tuc, while the work of Chapter 7 uses observations of younger white dwarfs in 47 Tuc to constrain the coupling of axions to electrons. Chapter 8 presents work showing how observations of hot, magnetic white dwarfs like ZTF J1901+1458 can be used to probe the coupling of axions to nucleons. The results of each of these research chapters are summarised in Chapter 9, and both the implications of the results and possible directions for future research are discussed for all of this work.

Chapter 2

White dwarf cooling

2.1 Basic Cooling Theory

Elementary white dwarf cooling theory was first described by Mestel [58], and this cooling theory has since been summarised in a number of reviews of white dwarf cooling (e.g. [59–61]). The key points of this theory are summarised in this section, largely following chapter 4 of Shapiro and Teukolsky [59]. In this elementary cooling theory, the dominant source of energy loss affecting the white dwarf cooling rate is the emission of photons (i.e. emission of neutrinos and exotic particles is neglected), and the interior of the white dwarf is assumed to be hot enough that the ions can be reasonably modelled as an ideal gas.

The basic model for the internal structure of a white dwarf consists of an electron-degenerate core enveloped by non-degenerate surface layers that are very thin compared to the core. In the electron-degenerate interior of a white dwarf, heat is transferred very efficiently via electron conduction, resulting in an approximately uniform temperature across the core. On the other hand, the non-degenerate surface layers covering this isothermal core are in radiative equilibrium, whereby the diffusion of photons carries an energy flux outward while matter is essentially maintained in thermodynamic equilibrium locally. These surface layers efficiently isolate the hot core of the white dwarf from the cold surrounding space, throttling the energy loss by photon radiation.

The luminosity L of the photons emitted from the surface of the white dwarf is related to the core temperature T_c by radiative heat transfer through the surface layers. The white dwarf is sustained against gravitational collapse by electron degeneracy pressure, so very little gravitational energy can

be released by further gravitational contraction, and there is no appreciable nuclear burning taking place in the core; thus, the only significant source of energy to be radiated is the residual thermal energy of the ions. The change in the core temperature T_c over time is thus related to the change in the total thermal energy U of the core, modelled as an isothermal heat bath. Combining these relations then gives a relation for how the luminosity varies with cooling time.

2.1.1 Energy Transfer Through Surface Layers

The structure of the surface layers relates the luminosity L to the core temperature T_c as follows. The photon diffusion equation for energy transfer through the surface layers is

$$L = -4\pi r^2 \frac{c}{3\kappa\rho} \frac{d}{dr} (aT^4) = -\frac{16\pi a c}{3} r^2 \frac{T^3}{\kappa\rho} \frac{dT}{dr}, \quad (2.1)$$

where κ is the opacity and (aT^4) is the blackbody energy density, with $a = 4\sigma_{\text{SB}}/c$ being a constant given in terms of the Stefan-Boltzmann constant σ_{SB} and the speed of light c . The opacity is assumed to follow Kramer's opacity law

$$\kappa = \kappa_0 \rho T^{-3.5}, \quad (2.2)$$

and the appropriate value for κ_0 is given by [61, p. 237]

$$\kappa_0 = (4.34 \times 10^{24} \text{ cm}^2 \text{ g}^{-1}) (1 + X) Z, \quad (2.3)$$

where X is the mass fraction of hydrogen and Z is the mass fraction of all elements heavier than helium. This law accounts for opacity due to the photoionization of atoms (“bound-free” transitions) and inverse bremsstrahlung of free electrons (“free-free” transitions).

The dependence on the radius r in Eq. (2.1) can be replaced with dependence on the pressure P using the equation of hydrostatic equilibrium, and then P can be replaced by the mass density ρ using the equation of state. The equation of hydrostatic equilibrium for the surface layers can be

2.1. Basic Cooling Theory

written as

$$\frac{dP}{dr} = -G M_{\text{WD}} \frac{\rho}{r^2}, \quad (2.4)$$

where the mass contained inside the radius r has been approximated as constant and equal to the white dwarf mass M_{WD} because the surface layers are much thinner than the core.

Rearranging Eq. (2.1) to get an expression for dT/dr , substituting in Eq. (2.2), and dividing Eq. (2.4) by this expression gives the equation

$$\frac{dP}{dT} = \frac{4ac}{3} \frac{4\pi G M_{\text{WD}}}{\kappa_0 L} \frac{T^{6.5}}{\rho}. \quad (2.5)$$

The density is related to the pressure and temperature through the equation of state, which is simply given by the ideal gas law in the non-degenerate surface layers, yielding

$$\rho = \mu m_u \frac{P}{k_B T}, \quad (2.6)$$

where μ is the mean molecular weight, $m_u = 1.66057 \times 10^{-24}$ g is the atomic mass unit, and k_B is Boltzmann's constant. Using Eq. (2.6) to eliminate the density from Eq. (2.5) gives an equation that can easily be integrated out to the surface boundary conditions $P = 0$ and $T = 0$. The resulting expression for P as a function of T can then be substituted back into Eq. (2.6).

This gives a relation for the density as a function of temperature in the envelope (for a white dwarf of a given luminosity and mass),

$$\rho = \left(\frac{2}{8.5} \frac{4ac}{3} \frac{4\pi G}{\kappa_0} \frac{\mu m_u}{k_B} \right)^{1/2} \left(\frac{M_{\text{WD}}}{L} \right)^{1/2} T^{3.25}. \quad (2.7)$$

This relationship breaks down in the interior where the electrons become degenerate.

The density ρ_* and temperature T_* at the boundary between the non-degenerate surface layers and electron-degenerate core can be estimated by equating the expression for the electron pressure P_e given by the equation

2.1. Basic Cooling Theory

of state for non-degenerate electrons

$$P_e = \frac{\rho_* k_B T_*}{\mu_e m_u} \quad (\text{non-degenerate}) \quad (2.8)$$

to that given by the equation of state for degenerate, non-relativistic electrons

$$P_e = \frac{3^{2/3} \pi^{4/3} \hbar^2}{5 m_e m_u^{5/3}} \left(\frac{\rho_*}{\mu_e} \right)^{5/3} \quad (\text{degenerate}) \quad (2.9)$$

where μ_e is the mean molecular weight per electron. Equating Eq. (2.8) to Eq. (2.9) gives the relation

$$\rho_* = (2.4 \times 10^{-8} \text{ g cm}^{-3}) \mu_e T_*^{3/2} \quad (2.10)$$

at the boundary, and the temperature T_* at this boundary can be taken to be the core temperature T_c .

Combining Eq. (2.10) and Eq. (2.7) with $T_* = T_c$ finally gives an expression relating L and T_c ,

$$L = C M_{\text{WD}} T_c^{3.5}, \quad (2.11)$$

where C is a constant, defined as

$$C = (5.7 \times 10^5 \text{ erg s}^{-1} M_{\odot}^{-1}) \frac{\mu}{\mu_e} \frac{1}{(1+X) Z}, \quad (2.12)$$

related to the composition of the surface layers. For a typical white dwarf consisting of a carbon/oxygen core surrounded by a thin helium layer and much thinner outer hydrogen layer, the mass fraction of helium can be taken to be $Y = 0.9$ and the mass fraction of hydrogen to be $X \approx 0$ as a reasonable example for illustrative purposes. Then $Z = 1 - X - Y$ has a value of $Z = 0.1$, and the values of the remaining parameters are $\mu_e \approx 2$ and $\mu \approx 1.4$. The value of the constant C in this case is then $C \approx 2 \times 10^6 \text{ erg s}^{-1} M_{\odot}^{-1}$.

For a white dwarf of a particular mass and composition, Eq. (2.11) gives the luminosity of photons emitted at the surface of the white dwarf as a function of the temperature of the core. Conversely, it can also be used to give the core temperature in terms of the luminosity, which will be made

use of later. This equation applies as long as heat transfer through the core is radiative (rather than convective), regardless of the state of matter of the core and even if appreciable energy is lost to the emission of other particles such as neutrinos or axions in addition to photons.

2.1.2 Cooling of Interior

The change in total thermal energy of the core is related to the change in the core temperature by the expression

$$\frac{dU}{dT_c} = \left(\frac{M_{\text{WD}}}{A m_u} \right) c_v, \quad (2.13)$$

where c_v is the specific heat capacity of the ions (at constant volume). The factor of $M_{\text{WD}} / (A m_u)$ accounts for the number of ions, where the mass of the core has again been approximated as equal to total mass of the white dwarf and $A m_u$ is the average mass of an ion in the core. Eq. (2.13) follows directly from the definition of the heat capacity combined with the first law of thermodynamics when the volume is constant, which is a condition that applies in the case of a white dwarf core sustained against further gravitational collapse by electron degeneracy pressure. The specific heat capacity in Eq. (2.13) has also been assumed to be independent of pressure and is thus treated as only a function of (core) temperature.

Both T_c and U are functions of time t , so the derivative chain rule

$$\frac{dU}{dT_c} = \frac{dU}{dt} \frac{dt}{dT_c} \quad (2.14)$$

can be used to re-write Eq. (2.13) as

$$\frac{dt}{dT_c} = \left(\frac{M_{\text{WD}}}{A m_u} \right) c_v \left(\frac{dU}{dt} \right)^{-1}. \quad (2.15)$$

If the white dwarf predominantly loses energy via photon radiation, then

$$\frac{dU}{dt} = -L, \quad (2.16)$$

2.1. Basic Cooling Theory

and L is given in terms of T_c by Eq. (2.11). Combining Eqs. (2.11) and (2.16) and substituting the resulting expression for dU/dt in terms of T_c into Eq. (2.15) gives the equation

$$\frac{dt}{dT_c} = - \left(\frac{1}{C A m_u} \right) c_v T_c^{-3.5}. \quad (2.17)$$

The term in brackets in Eq. (2.17) is a constant, while the c_v is in general a function of T_c . If the expression for c_v as a function of T_c is known, then Eq. (2.17) can be integrated.

The appropriate expression for the heat capacity in general depends on the state of matter of the core, which eventually changes over time from gaseous to liquid to solid as the white dwarf cools. However, the ions in the core can be approximated as an ideal gas for most of the white dwarf's evolution while it is still bright enough to be observed. In this ideal gas approximation, the specific heat capacity of the ions is

$$c_v = \frac{3}{2} k_B, \quad (2.18)$$

where k_B is Boltzmann's constant.

Substituting Eq. (2.18) in Eq. (2.17) and integrating from an initial time t_0 and corresponding initial temperature $T_{c,0}$ gives

$$t - t_0 = \frac{3}{5} \frac{k_B}{C A m_u} \left(T_c^{-5/2} - T_{c,0}^{-5/2} \right). \quad (2.19)$$

Taking $T_{c,0} \gg T_c$ and defining the cooling time to be

$$t_{\text{cool}} \equiv t - t_0 \quad (2.20)$$

simplifies this expression, giving

$$t_{\text{cool}} = \frac{3}{5} \frac{k_B}{A m_u C} T_c^{-5/2}. \quad (2.21)$$

Using Eq. (2.11) to express T_c in terms of L then gives the relationship

2.2. Neutrino Cooling

between cooling time and luminosity

$$t_{\text{cool}} = \frac{3}{5} \frac{k_B}{A m_u C^{2/7}} \left(\frac{L}{M_{\text{WD}}} \right)^{-5/7}. \quad (2.22)$$

The key features from these relations are the dependence of T_c and L on cooling time,

$$T_c \propto t_{\text{cool}}^{-2/5} \quad (2.23)$$

and

$$L \propto t_{\text{cool}}^{-7/5}. \quad (2.24)$$

The relationship between luminosity and cooling time derived in this section is the Mestel cooling curve.

2.2 Neutrino Cooling

Very early on in the cooling process, when a white dwarf is at its hottest temperature, neutrino emission is also an important source of energy loss. It is dominant over photon emission during this very early period of white dwarf cooling, before later becoming sub-dominant to photon cooling. Neutrinos are primarily produced in white dwarfs through the plasma neutrino process whereby a plasmon decays into a neutrino and anti-neutrino pair

$$\gamma^* \longrightarrow \nu_e \bar{\nu}_e, \quad (2.25)$$

where γ^* is the plasmon, ν_e is an electron neutrino, and $\bar{\nu}_e$ is an electron anti-neutrino. Itoh *et al.* [62] showed that the plasma neutrino process is the dominant neutrino production process when electrons are strongly degenerate, as is the case in the interior of a white dwarf. Though the improved calculations of Itoh *et al.* [63] modify the exact rate for the plasma neutrino process from the result of Itoh *et al.* [62], the plasma neutrino process was still found to be the dominant process, as summarised in Itoh *et al.* [64].

Calculating the energy loss rate due to the emission of neutrinos is a more

2.3. Axion Cooling

elaborate process than the simple calculations for photon radiation discussed in Section 2.1. For the dense interiors of stars, where many-body effects are important, the energy loss rate as a function of temperature and density for a particular composition typically needs to be calculated numerically. Analytic formulas are then typically fit to the results of such calculations for implementation as input into stellar evolution simulations. The results of such calculations for the energy loss rates of various neutrino processes in stellar interiors have been summarised by Itoh *et al.* [64], who provide analytic fitting formulas for their results.

Unlike photons, which are emitted from the surface of the white dwarf after diffusing through the optically thick outer layers, neutrinos are predominantly emitted from the interior of the white dwarf where they are produced. Because neutrinos interact so weakly with other particles, they have a much longer mean free path than photons and typically exit the white dwarf without further interaction after being produced. This provides another channel through which white dwarfs can lose energy, in addition to photon radiation, and contributes an additional term $-L_\nu$ to the right-hand side of Eq. (2.16), where L_ν is the total energy loss rate due to the emission of neutrinos. The additional energy loss due to neutrino emission causes white dwarfs to cool more quickly than they would by photon emission alone, which reduces the number of white dwarfs expected to be found over the luminosity range where neutrino emission is a dominant energy loss mechanism. This manifests as a deficit in the white dwarf luminosity function compared to Mestel’s cooling law at bright luminosities.

2.3 Axion Cooling

The emission of novel particles like axions can also affect white dwarf cooling in a way similar to the emission of neutrinos. Axions and their effect on stellar evolution are discussed in detail in Chapter 3, but a brief overview of the effect of axion emission on white dwarf cooling is given here for completeness. The main process by which axions can be produced in a white dwarf is axion bremsstrahlung from electrons scattering on nucleons. Axion

2.4. Convective Coupling

bremsstrahlung is analogous to standard bremsstrahlung, but with an axion emitted instead of a photon. Axion bremsstrahlung from electrons in particular relies on an axion-electron interaction, and though not all axion models include an axion-electron coupling, it is a common feature of axion models. The axion emission rate due to this process is sensitive to the strength of the axion-electron coupling, as well as the density and temperature of the white dwarf core. Like in the case of neutrino emission, determining the precise energy loss rate due to axion emission in dense stellar interiors is a complicated calculation that needs to be done numerically. Such computations have been performed by Nakagawa *et al.* [65, 66], who provided fitting formulas for their results.

Like neutrinos, axions produced in the dense interior of a white dwarf would typically leave the white dwarf without further interaction, carrying energy off with them. Axion emission thus likewise contributes an additional term $-L_x$ to the right-hand side of Eq. (2.16), where L_x is the total energy loss rate due to the emission of axions (or other novel particles that would have a similar effect). The emission rate of axions (produced via axion bremsstrahlung from electrons) has a different dependence on density and temperature than the emission rate of neutrinos (produced via the plasma neutrino process). For a typical white dwarf of mass $\sim 0.6 M_\odot$ with a carbon-oxygen core, this results in axion emission becoming a potentially dominant energy loss mechanism at a slightly later cooling time than neutrino emission, though still early in the white dwarf cooling process.

2.4 Convective Coupling

As the white dwarf cools, the heat transport in the envelope eventually becomes convective rather than radiative [67]. When the base of the convection zone that develops in the envelope reaches into the degenerate interior of the white dwarf, coupling the surface to the core, energy can be transferred across the outer envelope at a rate faster than what is possible through radiative transfer alone [68]. This convective coupling of the envelope to the core initially causes the white dwarf to appear more luminous than it would

if the energy continued to be transported radiatively across the envelope.

2.5 Core Crystallisation

As the temperature of the core decreases over time, the ions begin to crystallise into a lattice structure. When the temperature of the core drops to sufficiently low temperatures, the specific heat is due to the lattice vibrations of ions rather than free thermal motion of the ions. The critical temperature below which the lattice vibrations become a more important source of heat than the free thermal motion of the ions is called the Debye temperature θ_D , and c_v drops rapidly below this temperature.

The degree of crystallisation can be quantified by the plasma coupling parameter Γ , which is defined as the ratio of the Coulomb energy between nuclei $E_{\text{Coulomb}} = (Ze)^2 / r_i$ to the thermal energy $E_{\text{thermal}} = k_B T_c$ and is thus given by

$$\Gamma = \frac{(Ze)^2}{r_i k_B T_c}, \quad (2.26)$$

where r_i is the radius of the average spherical volume occupied by an ion and is given by

$$r_i = \left(\frac{3}{4\pi n_i} \right)^{1/3}, \quad (2.27)$$

where n_i is the number density of ions. For $\Gamma \ll 1$, the thermal energy dominates and the ions can be approximated as an ideal gas. For $\Gamma \gg 1$, the Coulomb energy forces dominate and the ion plasma crystallises into a periodic lattice structure that minimizes the Coulomb energy.

The phase transition from gas to liquid occurs when $\Gamma \sim 1$, which happens at a temperature T_g given by

$$T_g \sim (3 \times 10^4 \text{ K}) \left(\frac{\rho}{1 \text{ g cm}^{-3}} \right)^{1/3} Z^{5/3}. \quad (2.28)$$

The liquid to solid phase transitions occurs when $\Gamma \sim 175$, which happens

at the melting temperature T_m given by

$$T_m \approx \frac{1}{\Gamma} \frac{Z^2 e^2}{k_B} \left(\frac{4\pi}{3} \frac{\rho}{2Zm_u} \right)^{1/3} \approx (9 \times 10^2 \text{ K}) Z^{5/3} \left(\frac{\rho}{1 \text{ g cm}^{-3}} \right)^{1/3}, \quad (2.29)$$

where it has been assumed that $\mu_e \approx 2$. The Debye temperature is defined as the temperature at which $k_B\theta_D \equiv \hbar\Omega_p$, where $\Omega_p = (4\pi n_i Z^2 e^2 / m_i)^{1/2}$ is the plasma frequency for an ion with mass m_i ; this temperature is given by

$$\theta_D \approx (4 \times 10^3 \text{ K}) \left(\frac{\rho}{1 \text{ g cm}^{-3}} \right)^{1/2}. \quad (2.30)$$

At temperatures below θ_D , quantum effects become important. For a typical white dwarf with $Z > 2$ in the core, the Debye temperature is cooler than the melting temperature, $\theta_D < T_m$.

As the core of the white dwarf cools and transitions through these different states of matter, the specific heat capacity of the ions that appears in Eq. (2.15) varies as a function of temperature. In general, c_v needs to be calculated numerically. However, expressions can be derived for c_v in certain limiting cases that help to provide a better understanding of how the cooling rate is modified in different temperature regimes compared to Mestel cooling theory (see e.g. [59] for details). The expressions for c_v in some temperature regimes of interest are

$$c_v = \begin{cases} \frac{3}{2}k_B, & T_c \gg T_g \\ 3k_B, & \theta_D \ll T_c \ll T_g \\ \frac{16\pi^4}{5}k_B \left(\frac{T_c}{\theta_D} \right)^3, & T_c \ll \theta_D \end{cases} \quad (2.31)$$

The very high temperature gaseous regime $T_c \gg T_g$ is simply the ideal gas case that was considered in Section 2.1, which has been repeated here for completeness and easy reference. c_v/k_B increases from the value of $3/2$ in the ideal gas case to a maximum value of 3 in the high temperature lattice regime $\theta_D \ll T_c \ll T_g$ then drops dramatically with temperature in the low temperature lattice regime below the Debye temperature θ_D

Using these different expressions for c_v in Eq. (2.15) and solving like in Section 2.1 gives the following expressions for the cooling time

$$\tau = \begin{cases} \tau_{\text{mestel}}, & T_c \gg T_g \\ 2 \tau_{\text{mestel}}, & \theta_D \ll T_c \ll T_g \\ \frac{32\pi^4}{3} \left(\frac{T_c}{\theta_D} \right)^3 \left[\left(\frac{T_{c,0}}{T_c} \right)^{1/2} - 1 \right] \tau_{\text{mestel}}, & T_c \ll \theta_D \end{cases} \quad (2.32)$$

where $T_{c,0} > \theta_D$ is the initial core temperature. The case of $T_c \gg T_g$ is the result that was worked out in Section 2.1, with τ_{mestel} given as a function of T_c by Eq. (2.21). For the regime $\theta_D \ll T_c \ll T_g$, the cooling time is twice as long as would be predicted by the benchmark Mestel case. This is because the lattice potential energy provides an additional $(1/2)k_B T_c$ per vibrational mode to the internal energy per ion, which increases the heat capacity by a factor of 2 compared to the ideal gas case. For the regime $T_c \ll \theta_D$, the initial core temperature $T_{c,0}$ should be taken to be $T_{c,0} \lesssim \theta_D$, so it is always the case that $T_c \lesssim \theta_D$. In this cooling regime, the cooling time is initially longer than the other two regimes, but at even cooler temperatures becomes shorter than either of those regimes. For temperatures $T_c \lesssim 0.1 \theta_D$, the cooling time is shorter than the result for the hot lattice regime $\theta_D \ll T_c \ll T_g$. For temperatures $T_c \lesssim 0.07 \theta_D$, the cooling time is even shorter than the Mestel result for the regime $T_c \gg T_g$.

Note that cooling times given in Eq. (2.32) only account for the effect of different heat capacities in the different cooling regimes; they do not account for the additional energy provided by the latent heat of crystallisation. When the core undergoes the liquid-solid phase transition at a temperature of $T_c \sim T_m$, an amount of latent heat per ion of

$$-q_{\text{latent}} \sim k_B T_m \quad (2.33)$$

is released. This increases the total energy supply of the white dwarf, resulting in the white dwarf taking a longer time to cool than it would without this additional source of energy. As noted earlier, the liquid-solid phase transition typically occurs at temperatures above the Debye temperature,

2.6. White Dwarf Luminosity Function

$\theta_D < T_m < T_g$. The effect of including the latent heat of crystallisation is to further increase the cooling time compared to the classical lattice result for the regime $\theta_D \ll T_c \ll T_g$.

2.6 White Dwarf Luminosity Function

White dwarf cooling properties can be tested for a population of white dwarfs using the white dwarf luminosity function (WDLF), which is the number density of white dwarfs per unit (volume and) magnitude interval. The theoretical WDLF can be calculated as

$$\frac{dN}{dM_{\text{bol}}} = - \int_{M_\ell}^{M_u} \int_0^T \Phi(M_{\text{PS}}) \Psi(t) \frac{dt}{dM_{\text{bol}}} \times \delta\{t - [T - t_{\text{PS}}(M_{\text{PS}}) - t_{\text{cool}}(M_{\text{bol}}, M_{\text{PS}})]\} dt dM_{\text{PS}}, \quad (2.34)$$

where M_{PS} is the mass of a progenitor star that became a white dwarf with mass M_{WD} and absolute bolometric magnitude M_{bol} today, t is the time today since the population first began forming, T is the total age of the population since it first began forming (which is e.g. the age of the disc for a Galactic disc population), t_{PS} is the progenitor lifetime (i.e. the time that it takes the progenitor to evolve through all pre-white dwarf stages of evolution before becoming a white dwarf), and t_{cool} is the white dwarf cooling time (i.e. the time that the white dwarf has spent cooling since the start of the white dwarf stage of evolution). The time a white dwarf has spent cooling in general depends on both the current magnitude of the white dwarf and the mass of the white dwarf, the latter of which is related to the progenitor mass; t_{cool} has simply been labelled with M_{PS} instead of M_{WD} in Eq. (2.34) for convenience. The integration limits M_ℓ and M_u are the minimum and maximum progenitor masses that are able to produce a white dwarf with bolometric magnitude M_{bol} today. The lower mass limit M_ℓ satisfies the condition

$$t_{\text{cool}}(M_{\text{bol}}, M_\ell) + t_{\text{PS}}(M_\ell) = T. \quad (2.35)$$

2.6. White Dwarf Luminosity Function

The function Φ is the initial mass function, which describes the distribution of the initial progenitor masses. The function Ψ is the time-dependent star formation rate of the progenitors as a function of the time since the population first began forming.

Eq. (2.34) applies for the general situation in which the rate at which white dwarfs are born varies with time and the white dwarf population is comprised of white dwarfs of different masses. The time-varying white dwarf birthrate is accounted for indirectly in Eq. (2.34) through the star-formation rate of the progenitors. This general form of the WDLF is applicable, for example, when studying populations of white dwarfs in the Galactic disc. In the case of a globular cluster population where the progenitors all formed at the same time in a single burst of star formation and where a white dwarf population can often be approximated as consisting of white dwarfs that all have the same mass, the expression for the WDLF simplifies.

If the white dwarfs are born at a constant rate and have identical properties, especially the same mass, then the number density of white dwarfs in a given bolometric magnitude interval dM_{bol} is simply proportional to the time interval dt that it takes a white dwarf to cool through that magnitude range. This can be expressed as

$$\frac{dN}{dM_{\text{bol}}} = B \frac{dt}{dM_{\text{bol}}}, \quad (2.36)$$

where N is the number density per unit volume and B is the constant white dwarf birthrate, giving the number of white dwarfs born per unit time per unit volume.

The bolometric magnitude is related to the luminosity (of photon radiation) by the expression

$$M_{\text{bol}} = M_{\text{bol},\odot} - 2.5 \log_{10} \left(\frac{L}{L_{\odot}} \right), \quad (2.37)$$

where the nominal solar values are set by IAU 2015 resolution [69] to be $L_{\odot} = 3.828 \times 10^{26}$ W and $M_{\text{bol},\odot} = 4.74$. The relation between cooling time and luminosity thus determines the WDLF predicted by theory. In

2.6. White Dwarf Luminosity Function

the elementary case where photon radiation is the dominant energy loss mechanism and the white dwarf core is approximated as an isothermal heat bath, with the ions approximated as an ideal gas, this relation is given by Eq. (2.22).

In Section 2.1, this relation was derived for the elementary case where photon radiation is the dominant energy loss mechanism and the white dwarf is still hot enough that the ions in its core can be approximated to be an ideal gas. The resultant relation is given by Eq. (2.22). Combining Eq. (2.22) with Eqs. (2.36) and (2.37) gives Mestel's cooling law [51]

$$\log_{10} \left(\frac{dN}{dM_{\text{bol}}} \right) = \frac{2}{7} M_{\text{bol}} + \text{constant}, \quad (2.38)$$

where the additive constant depends on the mass and composition of the white dwarf.

This relation is modified in regimes where the phenomena discussed in Sections 2.2 to 2.5 are important. Noting from Eq. (2.37) that

$$\frac{dL}{dM_{\text{bol}}} = -\frac{2}{5} \ln(10) L \quad (2.39)$$

and expressing dt/dL in terms of the key relations that determine it

$$\frac{dt}{dL} = \frac{dT_c}{dL} \frac{dU}{dT_c} \frac{dt}{dU}, \quad (2.40)$$

then Eq. (2.36) for the WDLF can be expressed as

$$\frac{dN}{dM_{\text{bol}}} = -\frac{2}{5} \ln(10) B L \frac{(dT_c/dL) (dU/dT_c)}{dU/dt}. \quad (2.41)$$

When photon emission is the dominant energy loss mechanism, dU/dt is given by Eq. (2.16). However, at early cooling times where the neutrino emission and axion emission can be important energy loss mechanisms, Eq. (2.16) is more generally written as

$$\frac{dU}{dt} = - (L + L_\nu + L_x), \quad (2.42)$$

2.6. White Dwarf Luminosity Function

where L_ν denotes the total energy loss rate due to neutrino volume emission and L_x denotes the total energy loss rate due to the emission of other, novel particles like axions. For radiative heat transport through the surface layers, L and T_c are related through Eq. (2.11). This relation is modified at late times when heat transport through the surface layers becomes convective instead of radiative and the envelope becomes convectively coupled to the core. U and T_c are related through the heat capacity, with dU/dT_c given by Eq. (2.13). As discussed in Section 2.5, the heat capacity depends on the state of matter, and in particular changes as the white dwarf core crystallises.

For the early cooling times where neutrino and axion emission are important, heat transfer through the surface can still be assumed to be radiative and the ions in the core can still be approximated as an ideal gas. Following the notation of Raffelt [51] for convenience, let $K \equiv C M_{\text{WD}}$ so that Eq. (2.11) becomes $L = K T_c^{7/2}$, and let $\tilde{C} = (3/2) k_B M_{\text{WD}}/(A m_u)$ so that Eq. (2.13) becomes $dU/dT_c = \tilde{C}$ for an ideal gas, for which $c_v = (3/2) k_B$. Then the expression for the WDLF becomes

$$\frac{dN}{dM_{\text{bol}}} = \frac{4}{35} \ln(10) B \frac{\tilde{C}}{K^{2/7}} \frac{L^{2/7}}{(L + L_\nu + L_x)}. \quad (2.43)$$

This shows how the emission of neutrinos and axions modifies the WDLF function relative to what is expected from photon emission alone at a given luminosity. Non-negligible values of either L_ν or L_x give smaller values of dN/dM_{bol} at a given L . This corresponds to a deficit in the number of white dwarfs expected in a given magnitude range relative to the expectation from Mestel's law.

Chapter 3

Axions

3.1 The QCD Axion

The axion is a hypothesised light pseudo-scalar particle, originally motivated as an explanation for the lack of observed CP violation in the strong force. Quantum chromodynamics (QCD), the theory describing the strong interaction, generically allows CP violation in the strong sector within the Standard Model of particle physics. The lack of observed CP violation despite CP violation being expected in QCD is referred to as the strong CP problem. This problem can be solved by introducing a spontaneously broken global $U(1)$ symmetry, with the axion being the (pseudo-)Nambu-Goldstone boson associated with the spontaneous symmetry breaking.

Axion models that solve the strong CP problem, as was their original motivation, are called QCD axions. The QCD axion is a particularly well-motivated dark matter candidate, as it arises naturally in a solution to one of the few problems with the Standard Model. There is also a much larger class of axion-like particle (ALP) models of light pseudo-scalar bosons similar to the QCD axion, but that do not solve the strong CP problem. These particles are often referred to simply as axions, and are also viable dark matter candidates.

3.2 The Strong CP Problem and Solution

The full QCD Lagrangian density can be expressed as

$$\mathcal{L}_{\text{QCD}} = -\frac{1}{4}G_a^{\mu\nu}G_{a\mu\nu} - \sum_f \bar{q}_f (\gamma^\mu D_\mu + m_f) q_f + \mathcal{L}_{\text{CP}}, \quad (3.1)$$

3.2. The Strong CP Problem and Solution

$$\mathcal{L}_{\text{CP}} = -\bar{\theta} \frac{\alpha_s}{8\pi} G_a^{\mu\nu} \tilde{G}_{a\mu\nu}, \quad (3.2)$$

where q_f are the quark fields (with the index f labelling the quark flavours), m_f are the corresponding quark masses, $G_a^{\mu\nu}$ is the gluon field strength tensor (with the index a labelling the gluon fields), $\tilde{G}_{a\mu\nu} \equiv \frac{1}{2} \epsilon_{\mu\nu\rho\phi} G_a^{\rho\phi}$ is the dual gluon field strength tensor, $\bar{\theta}$ is a dimensionless parameter that parametrizes the degree of CP violation in the strong sector, and α_s is the QCD fine-structure constant (analogous to the fine structure constant of electromagnetism). The covariant derivative in Eq. (3.1) is given by

$$D_\mu = \partial_\mu - i g_s G_\mu^a \frac{\lambda_a}{2}, \quad (3.3)$$

where g_s is the coupling constant of the strong interactions, G_μ^a denotes the gluon fields, and λ_a are the Gell-Mann matrices⁶. The gluon field strength tensor is defined in terms of the gluon fields through the expression

$$G_{\mu\nu}^a = \partial_\mu G_\nu^a - \partial_\nu G_\mu^a + g_s f_{bc}^a G_\mu^b G_\nu^c, \quad (3.4)$$

where f_{bc}^a are the structure constants of the group $SU(3)$. Together with the strong force generators $T_a = (1/2) \lambda_a$, these structure constants characterise the Lie algebra describing QCD such that the generators satisfy the commutation relations $[T_a, T_b] = i f_{ab}^c T_c$. Note that the fine-structure constant of the strong force is distinct from the structure constants of the Lie algebra of $SU(3)$. The QCD fine-structure constant is simply defined in terms of the strong coupling through the expression $\alpha_s = g_s^2 / (4\pi)$.

The first two terms in Eq. (3.1) for the QCD Lagrangian are invariant under CP transformation, but the term \mathcal{L}_{CP} violates CP symmetry⁷. The amount of CP violation is parametrised by $\bar{\theta}$, which is generically expected to have a value of order ~ 1 . Though CP violation has been observed in the weak interactions, it has not been observed in the strong interactions. The strong CP-violating term given by Eq. (3.2) results in a neutron electric

⁶These are a set of 3×3 matrices for which the rows and columns correspond to the three quark colours. There are eight Gell-Mann matrices and likewise eight gluon fields.

⁷This term violates P and T but not C, so it also violates CP.

3.2. The Strong CP Problem and Solution

dipole moment $d_n \simeq -3.9 \times 10^{-16} \bar{\theta} e \cdot \text{cm}$ [70], which has been constrained by experimental searches to be $|d_n| \leq 2.9 \times 10^{-26} e \cdot \text{cm}$ [71, 72], resulting in $\bar{\theta}$ being constrained to $|\bar{\theta}| \lesssim 7.4 \times 10^{-11}$. This strongly constrained lack of CP violation in the strong interactions despite being a prediction of QCD, with no *a priori* reason for $\bar{\theta}$ to be so small, is referred to as the strong CP problem.

Though Eq. (3.2) corresponds to a total derivative, which integrates to a boundary term, the non-trivial vacuum topology of QCD prevents this term from vanishing. More specifically, QCD has an infinite number of degenerate, topologically distinct vacuum states, with the topology of each state characterised by the Chern-Simons number $n \in \mathbb{Z}$. The QCD vacuum is a superposition of these states parametrised by the phase θ such that

$$|\theta\rangle = \sum_n e^{i\theta n} |n\rangle, \quad (3.5)$$

and instantons allow tunnelling between states labelled by different n . In addition to the non-trivial vacuum of QCD, the strong CP-violating term also receives a contribution arising from the phase transformations of the quark fields needed to make the quark mass matrix M_q real and diagonal in the full Standard Model. The total amount of CP violation in the strong sector is parametrised by $\bar{\theta} = \theta + \arg \det M_q$. The two terms contributing to $\bar{\theta}$ have completely different physical origins with no *a priori* reason why these terms should cancel to give a value of $\bar{\theta}$ small enough to be in keeping with experimental constraints [51].

Peccei and Quinn presented a solution to the strong CP problem in 1977 [73, 74] by postulating that the Standard Model was invariant under a spontaneously broken global $U(1)$ chiral symmetry, now known as the PQ symmetry $U(1)_{\text{PQ}}$. To better understand the strong CP problem and its solution, first consider the case of massless quarks. If quarks were massless, with n_f being the number of quarks with non-trivial PQ charge, then $U(1)_{\text{PQ}}$ could be an exact symmetry and the strong CP problem would be trivially solved. Under a $U(1)_{\text{PQ}}$ transformation parametrised by α , the left- and

3.2. The Strong CP Problem and Solution

right-handed quark fields (q_L and q_R) transform as

$$q_R \rightarrow e^{i\alpha} q_R, \quad q_L \rightarrow e^{-i\alpha} q_L, \quad (3.6)$$

which changes the vacuum angle $\bar{\theta}$ through the QCD chiral anomaly [75]

$$\bar{\theta} \rightarrow \bar{\theta} + 2n_f\alpha. \quad (3.7)$$

Thus $\bar{\theta}$ could be set to zero through a simple redefinition of the quark fields [76]. However, since quarks are massive, their mass term in the Lagrangian explicitly breaks the chiral symmetry. To see this, write the mass term $\sum_f m_f \bar{q}_f q_f$ in terms of the left- and right-handed components $\bar{q}_f q_f = \bar{q}_{fL} q_{fR} + \bar{q}_{fR} q_{fL}$. Each mass term transforms as

$$\begin{aligned} m_f \bar{q}_f q_f &= m_f (\bar{q}_{fL} q_{fR} + \bar{q}_{fR} q_{fL}) \\ &\rightarrow m_f (e^{2i\alpha} \bar{q}_{fL} q_{fR} + e^{-2i\alpha} \bar{q}_{fR} q_{fL}) \neq m_f \bar{q}_f q_f. \end{aligned} \quad (3.8)$$

So $\bar{\theta}$ is physical and cannot be removed by simplying redefining the quark fields.

The PQ solution instead introduces a new complex scalar field that carries PQ charge and has a non-zero vacuum expectation value (vev). In order for the Standard Model to conserve $U(1)_{\text{PQ}}$ while still solving the strong CP problem, PQ charge must be carried by at least one flavour of quark (either a Standard Model quark or new species), which acquires its mass during the spontaneous breaking of $U(1)_{\text{PQ}}$ from its Yukawa interaction with the new PQ-charged scalar field, rather than from the Standard Model Higgs field. The axion is the (pseudo-)Nambu-Goldstone boson associated with the spontaneous breaking of PQ symmetry [77, 78], corresponding to the angular component of this new complex scalar field. Since $U(1)_{\text{PQ}}$ is also explicitly broken by the QCD chiral anomaly at energies at and below the QCD scale $\Lambda_{\text{QCD}} \sim 200$ MeV, the axion is not a true Nambu-Goldstone boson, though it is approximately so at high energies.

Peccei and Quinn showed that if such conditions are met, then after

3.2. The Strong CP Problem and Solution

PQ symmetry breaking, the parameter $\bar{\theta}$ in Eq. (3.2) is replaced by the pseudo-scalar axion field, converting the CP-violating term to a CP-invariant interaction between axions and gluons. To see this, let Φ be the new scalar field carrying PQ charge. Neglecting the radial component of the field, which is not important for the following discussion, Φ can be written as

$$\Phi(x) = \frac{f_a}{\sqrt{2}} e^{ia(x)/f_a}, \quad (3.9)$$

where f_a is the PQ symmetry-breaking scale and $a(x)$ is the axion field. Under a $U(1)_{\text{PQ}}$ transformation, $\Phi \rightarrow e^{i\alpha} \Phi$, the axion field transforms as

$$a(x) \rightarrow a(x) + \alpha f_a. \quad (3.10)$$

Though classically the Lagrangian describing this theory should only have the axion field appear through derivative couplings to conserve PQ symmetry, quantum Standard Model anomalies break this symmetry and allow tree-level couplings of the axion to the Standard Model gauge bosons. In particular, the QCD chiral anomaly permits a term that couples the axion and gluon fields

$$\mathcal{L}_{aG\tilde{G}} = -\frac{a}{f_a} \xi \frac{\alpha_s}{8\pi} G_b^{\mu\nu} \tilde{G}_{b\mu\nu}, \quad (3.11)$$

where ξ is a parameter that depends on the particular axion model being implemented.

For spontaneous breaking of PQ symmetry, we write $a(x)$ as the sum of its vev $\langle a \rangle$ and excitation \tilde{a} above its vev, $a(x) = \langle a \rangle + \tilde{a}(x)$, and see that

$$\mathcal{L}_{\text{CP}} + \mathcal{L}_{aG\tilde{G}} = -\left(\bar{\theta} + \langle a \rangle \frac{\xi}{f_a}\right) \frac{\alpha_s}{8\pi} G_a^{\mu\nu} \tilde{G}_{a\mu\nu} - \tilde{a} \frac{\xi}{f_a} \frac{\alpha_s}{8\pi} G_b^{\mu\nu} \tilde{G}_{b\mu\nu}. \quad (3.12)$$

The term containing $\langle a \rangle$ effectively shifts the QCD vacuum angle so that $\bar{\theta}_{\text{eff}} = \bar{\theta} + \langle a \rangle \xi / f_a$ parametrises the total CP-violating term, while the term for the interaction of \tilde{a} with gluons provides a potential for the axion. This potential is temperature-dependent and vanishes at energy scales well above the QCD scale ($\Lambda_{\text{QCD}} \sim 200$ MeV, which corresponds to a temperature $T_{\text{QCD}} \sim 2 \times 10^{12}$ K), so $\langle a \rangle$ is unconstrained in the high-temperature limit.

3.3. Axion Models

However, as the potential becomes non-negligible at energy scales below the QCD scale, $\langle a \rangle$ is driven to the energy minimum of the effective axion potential. In the low-temperature limit, the vev of the axion field takes the value $\langle a \rangle = -\bar{\theta}f_a/\xi$, thus dynamically driving $\bar{\theta}_{\text{eff}}$ to zero and causing the strong CP-violating term to vanish. What remains is the CP-invariant term giving the interaction between the gluon fields and the physical (excited) axion field \tilde{a} , which from now on will just be called a .

3.3 Axion Models

There exist multiple models for implementing the QCD axion, depending on which particles are assigned non-trivial PQ charge. In hadronic models, like the Kim-Shifman-Vainshtein-Zakharov (KSVZ) model [79, 80] new quark species are introduced to carry the PQ charge. In fermionic models, such as the Dine-Fischler-Srednicki-Zhitnitsky (DFSZ) model [81, 82] PQ charge can be carried by Standard Model fermions.

The most general Lagrangian density that can be written for the axion field, including all possible interactions with Standard Model particles, can be expressed as [83]

$$\begin{aligned} \mathcal{L}_{\text{axion}} = & -\frac{1}{2}\partial_\mu a\partial^\mu a + J^\mu\partial_\mu a \\ & -\frac{a}{f_a}\left[\xi\frac{\alpha_s}{8\pi}G_b^{\mu\nu}\tilde{G}_{b\mu\nu} - k_1W_b^{\mu\nu}\tilde{W}_{b\mu\nu} - k_2B^{\mu\nu}\tilde{B}_{\mu\nu}\right], \end{aligned} \quad (3.13)$$

with the current J^μ given by

$$J^\mu f_a = i\bar{\psi}_{i,m}(X_{imn} + iY_{imn}\gamma_5)\psi_{i,n} + X_{\phi_H}\left(\phi_H^\dagger iD^\mu\phi_H + \text{h.c.}\right), \quad (3.14)$$

where $W_b^{\mu\nu}$ and $B^{\mu\nu}$ are the gauge bosons of $SU(2)$ and $U(1)$ and mediate the electro-weak interaction, ϕ_H is the Standard Model Higgs field, and ψ denotes all fermions that carry PQ charge in the model, with i indexing all flavours of quarks and leptons and m and n indexing all generations of a given flavour. The abbreviation “h.c.” denotes Hermitian conjugate terms. The quantities ξ , k_1 , k_2 , X_{imn} , Y_{imn} , and X_{ϕ_H} are model dependent

3.4. Axion Interactions Important for Stellar Evolution

parameters. There can also be additional terms if new particle content is introduced in a particular axion model.

At temperatures well above the QCD scale, the QCD chiral anomaly term that explicitly breaks PQ symmetry is negligible, so the axion behaves like a true Nambu-Goldstone boson and is massless. Below the QCD scale, PQ symmetry is explicitly broken and the axion becomes massive. At temperatures well below the QCD scale, as is relevant for stellar astrophysics, the axion mass is

$$m_a \approx f_\pi m_\pi^2 \frac{m_u m_d}{(m_u + m_d)^2} \frac{1}{f_a}, \quad (3.15)$$

where m_π , m_u , and m_d are the masses of the pion, up quark, and down quark, and f_π is the pion symmetry-breaking scale.

3.4 Axion Interactions Important for Stellar Evolution

3.4.1 Overview

There are a large variety of ways that axions could have an observable effect in astrophysical environments, and in particular have observable consequences for stars. This section focuses on axion phenomena most relevant for stellar evolution, including the evolution of white dwarfs, and the current constraints on axion parameters pertaining to those phenomena. For a more extensive review of the various ways axions could affect astrophysical observables and the corresponding constraints on axion parameters, including subdominant constraints, see the recent review by Di Luzio *et al.* [84]. For a review that also includes axion constraints from experimental searches, see e.g. Di Luzio *et al.* [85].

The main axion production processes relevant for stellar evolution are the axion Primakoff effect and axion bremsstrahlung. The Primakoff effect is a process by which a photon in an external electric field (such as the Coulomb field of a nucleus in a stellar plasma) can convert into a particle with a two-photon vertex interaction (such as a neutral pion, for which this

3.4. Axion Interactions Important for Stellar Evolution

effect was first noted [86], or an axion, as is relevant here), with the external field providing a virtual photon for the interaction. Axion bremsstrahlung is a process analogous to standard bremsstrahlung but with an axion emitted instead of a photon, and it can occur via the scattering on nucleons of either electrons (through the axion-electron interaction) or other nucleons (through axion-nucleon interactions).

The axion Primakoff effect is most important for the horizontal branch stage of evolution during which stars are hot but less dense than stages where the core is degenerate, such as the red giant branch (RGB), asymptotic giant branch (AGB), or white dwarf stages of evolution. Axion bremsstrahlung via electrons becomes an important effect for stellar cooling when a star is dense enough to have an electron-degenerate core, in particular during the RGB phase of low mass stars and in white dwarfs. Axion bremsstrahlung via nucleons is important in neutron stars but is not relevant for white dwarfs or the low-mass progenitors from which they form, so it will not be discussed further. It should be noted that all axion models include the axion-photon interaction through which the axion Primakoff process occurs. Axion bremsstrahlung by electrons, on the other hand, occurs via the axion-electron interaction, which only exists for models where electrons carry PQ charge. In particular, such an interaction does exist for the DFSZ axion model, though there is no such interaction (at tree-level) in the KSVZ axion model.

3.4.2 Primakoff Effect

Axions generically have an interaction with photons given by

$$\mathcal{L}_{a\gamma\gamma} = -\frac{g_{a\gamma\gamma}}{4} F_{\mu\nu} \tilde{F}^{\mu\nu} = g_{a\gamma\gamma} a \mathbf{E} \cdot \mathbf{B}, \quad (3.16)$$

where $g_{a\gamma\gamma}$ is the axion-photon coupling constant, $F_{\mu\nu}$ is the electromagnetic field strength tensor, $\tilde{F}^{\mu\nu}$ is the dual of $F_{\mu\nu}$, \mathbf{E} is the electric field, \mathbf{B} is the magnetic field, and a is the axion field.

This axion-photon interaction enables the Primakoff effect, whereby a plasmon in the presence of the Coulomb field of a nucleus converts into an

3.4. Axion Interactions Important for Stellar Evolution

axion

$$\gamma^* + (Z, A) \longrightarrow a + (Z, A), \quad (3.17)$$

where γ^* denotes the plasmon and (Z, A) denotes a nucleus with proton number Z and atomic mass number A .

The energy loss rate per unit mass via Primakoff conversion is [51, 52, 84, 85]

$$\epsilon = (2.8 \times 10^{21} \text{ erg g}^{-1} \text{ s}^{-1}) \left(\frac{g_{a\gamma\gamma}}{\text{GeV}^{-1}} \right)^2 \frac{T_8^7}{\rho_4} F(\xi^2), \quad (3.18)$$

where the temperature T and density ρ have been scaled by typical values for horizontal branch stars,

$$\rho_4 \equiv \frac{\rho}{10^4 \text{ g cm}^{-3}}, \quad (3.19)$$

$$T_8 \equiv \frac{T}{10^8 \text{ K}}, \quad (3.20)$$

and F is a function that depends on the Debye-Huckle screening wavenumber κ through the variable ξ , which is defined to be

$$\xi \equiv \frac{\kappa}{2T}. \quad (3.21)$$

The function $F(\xi^2)$ is given explicitly by the expression

$$F(\xi^2) = \frac{\xi^2}{2\pi} \int_0^\infty dx \left[(x^2 + \xi^2) \ln \left(1 + \frac{x^2}{\xi^2} \right) - x^2 \right] \frac{x}{e^x - 1}, \quad (3.22)$$

and it can be approximated by the simpler analytic expression [87]

$$F(\xi^2) \simeq \left(\frac{1.037\xi^2}{1.01 + \xi^2/5.4} + \frac{1.037\xi^2}{44 + 0.628\xi^2} \right) \ln \left(3.85 + \frac{3.99}{\xi^2} \right). \quad (3.23)$$

Typically the value of F is of order unity in regimes where the Primakoff effect is important. For horizontal branch stars with $\rho = 10^4 \text{ g cm}^{-3}$ and $T = 10^8 \text{ K}$, the values of κ^2 and F are $\kappa^2 \approx 2.5$ and $F \approx 0.98$.

Energy loss via the Primakoff process is particularly important for stars

3.4. Axion Interactions Important for Stellar Evolution

on the horizontal branch. These are stars hot enough to burn He in their cores, and their cores are not dense enough to be electron-degenerate. They are hotter than main-sequence stars (which are stars burning H in their cores), but not as dense as white dwarfs or stars on the RGB or AGB (which are all dense enough to have electron-degenerate cores). Note that the energy emission rate for the Primakoff process given by Eq. (3.18) is suppressed at high densities by the factor of ρ in the denominator. Eq. (3.18) specifically applies in the non-degenerate case (which is applicable for horizontal branch stars), but the Primakoff emission rate in the high-density degenerate case is also suppressed.

The additional avenue of energy loss provided by the emission of axions produced through the Primakoff process reduces the amount of time stars spend on the horizontal branch, while not appreciably affecting the time spent on the RGB. The number of stars observed at a particular stage of stellar evolution is proportional to the length of time spent in that phase, so the additional axion energy loss from the Primakoff process would decrease the ratio of stars on the horizontal branch to those on the RGB [87–89].

To discuss constraints of $g_{a\gamma\gamma}$ from this phenomenon, one can define [89] the R parameter to be the ratio of the number N_{HB} of stars found on the horizontal branch over the number N_{RGB} of stars found on the RGB:

$$R \equiv \frac{N_{\text{HB}}}{N_{\text{RGB}}}. \quad (3.24)$$

Larger values of $g_{a\gamma\gamma}$ will result in smaller theoretical values of R . The most recent bound on $g_{a\gamma\gamma}$ from measurements of the R parameter is [89, 90]

$$g_{a\gamma\gamma} \leq 0.65 \times 10^{-10} \text{ GeV}^{-1} \quad (3.25)$$

at 95% confidence level. This bound was obtained by analysing measurements of the R parameter in comparison to theoretical predictions for a sample of globular clusters.

Constraints on $g_{a\gamma\gamma}$ can also be derived from a related parameter called the R_2 parameter, which is the ratio of the number N_{AGB} of stars on the

3.4. Axion Interactions Important for Stellar Evolution

AGB to the number of stars on the horizontal branch such that

$$R_2 \equiv \frac{N_{\text{AGB}}}{N_{\text{HB}}}. \quad (3.26)$$

The R_2 parameter has been used to derive a bound of [91]

$$g_{a\gamma\gamma} < 0.47 \times 10^{-10} \text{ GeV}^{-1}, \quad (3.27)$$

which is an improvement on the bound from the R parameter given above.

3.4.3 Axion Bremsstrahlung

The DFSZ model contains an axion-electron interaction term of the form

$$\mathcal{L}_{aee} = i g_{aee} \bar{\psi}_e \gamma_5 \psi_e a, \quad (3.28)$$

where g_{aee} is the axion-electron coupling constant, a is the axion field, and ψ_e is the electron field. The axion-electron coupling can be expressed as

$$g_{aee} = \frac{C_e m_e}{f_a}, \quad (3.29)$$

where C_e is a model-dependent constant, m_e is the electron mass, and $f_a = f_{PQ}/N_f$ is the Peccei-Quinn (symmetry breaking) scale.

Note that hadronic axion models such as the KSVZ model do not contain axion-electron tree-level interactions, as fermions do not carry PQ charge in these models. Nevertheless, this term is present in the popular DFSZ model and other grand-unified axion models. To summarise for the benchmark models, C_e has the form

$$C_e = \begin{cases} \frac{\sin^2 \beta}{N_f} & (\text{DFSZ}) \\ 0 & (\text{KSVZ}) \end{cases}, \quad (3.30)$$

where N_f is the number of standard fermion families, which can be taken to be $N_f = 3$. For the DFSZ model, C_e has simply been recast in terms of the parameter β , where $\tan \beta$ is defined as the ratio of the vacuum ex-

3.4. Axion Interactions Important for Stellar Evolution

pectation values of the two electroweak doublet scalar fields that carry PQ charge in the DFSZ model. Note that it is simply a matter of convention which vev appears in the numerator or denominator when defining $\tan \beta$. The definition used here is $\tan \beta \equiv \langle H_u \rangle / \langle H_d \rangle$, where $\langle H_u \rangle$ is the vev of the Higgs doublet H_u that couples to the up-type quarks and $\langle H_d \rangle$ is the vev of the Higgs doublet H_d that couples to the down-type quarks. This definition has been used in some recent reviews of axion models and constraints, e.g. [84, 85]. If β is alternatively defined such that $\langle H_u \rangle$ and $\langle H_d \rangle$ are swapped in the expression for $\tan \beta$, then $\sin \beta$ is simply replaced by $\cos \beta$ in the expression for C_e . This alternative definition was used in some earlier reviews that have been extensively cited, e.g. [51, 52]. For the definition of β used here, $\tan \beta$ is restricted to the interval $\tan \beta \in [0.25, 170]$ to remain in the perturbative domain.

This axion-electron coupling leads to a process analogous to standard bremsstrahlung, with the photon of standard bremsstrahlung replaced by an axion

$$e^- + (Z, A) \rightarrow e^- + (Z, A) + a. \quad (3.31)$$

This process is commonly referred to as “axion bremsstrahlung”. The energy loss rate due to axion bremsstrahlung in dense stellar interiors has been calculated to be [65, 66]

$$\epsilon = (1.08 \times 10^{23} \text{ ergs g}^{-1} \text{ s}^{-1}) \alpha_a \frac{Z^2}{A} T_7^4 F, \quad (3.32)$$

where ϵ is the energy loss rate per unit mass, $\alpha_a = g_{aee}^2 / (4\pi)$ is a fine-structure constant (analogous to the fine-structure constant of electromagnetism), Z and A are respectively the number of protons and total number of nucleons of the nucleus, $T_7 = T/(10^7 \text{ K})$ is the temperature, and F is a dimensionless form factor with a prescription given by Nakagawa *et al.* [65, 66]. The value of F depends on the phase and composition of material in the star, though its maximum value is typically of order unity ($F \sim 1$).

Axion bremsstrahlung (via electrons) will affect stellar evolution during the RGB stage of low-mass stars and during white dwarf cooling, where the

3.4. Axion Interactions Important for Stellar Evolution

star has an electron-degenerate core. The effect on RGB evolution is to delay the onset of He burning and thus increase the luminosity and temperature of the tip of the RGB⁸. The bound on the axion-electron coupling due to delayed He ignition on the RGB is [92]

$$g_{aee} \leq 1.5 \times 10^{-13} \quad (3.33)$$

at 95% confidence level. This bound was determined by an analysis of the tip of the RGB for 22 globular clusters. A similar bound of $g_{aee} \leq 1.6 \times 10^{-13}$ was also found by an independent analysis of the RGB tip of the globular cluster NGC 4258 [93].

As for white dwarfs, energy loss through axion bremsstrahlung affects the rate at which white dwarfs cool. This affects both the shape and cutoff of the WDLF, which is the number of white dwarfs found as a function of luminosity. Hints of axions have been suggested based on observed WDLFs of white dwarfs in the Galactic thin and thick discs, as well as the halo, to explain features of these WDLFs that indicate a cooling anomaly [47–49]. In addition to finding a favoured range of values for g_{aee} , the analysis of Miller Bertolami *et al.* [48] investigating these axion hints also provided a bound on the axion-electron coupling of $g_{aee} \leq 2.8 \times 10^{-13}$ at 3σ confidence (and $g_{aee} \leq 2.1 \times 10^{-13}$ at 2σ confidence), though this bound is not as stringent as the more recent bound from the tip of the RGB. The hints of axions suggested by WDLFs is discussed in further detail in Section 3.5 below.

The effect of axion emission on white dwarf cooling can also be searched for using observations of white dwarf variables, a class of white dwarfs whose luminosity changes periodically with time. The pulsation period P of these variable white dwarfs is typically on the order of a few minutes and decreases over time at a rate \dot{P} that depends on the cooling efficiency of the white dwarf. To a good approximation, \dot{P}/P is directly proportional to the cooling

⁸The number of stars on the RGB is also increased, which further decreases the value of the R parameter given by Eq. (3.24). The bound on $g_{a\gamma\gamma}$ reported in Eq. (3.25) based on the R parameter is a conservative bound that assumes no axion-electron coupling. Joint bounds on $g_{a\gamma\gamma}$ and g_{aee} can be derived from the R parameter for axion models that include a coupling to electrons, but such bounds are model-dependent.

3.4. Axion Interactions Important for Stellar Evolution

Observation	$g_{aee} / 10^{-13}$	Reference
Tip of RGB (22 clusters)	≤ 1.5	Straniero <i>et al.</i> [92]
Tip of RGB (NGC 4258)	≤ 1.6	Capozzi and Raffelt [93]
WDLF (Galactic disc, halo)	≤ 2.1	Miller Bertolami <i>et al.</i> [48]
White Dwarf Variables	≤ 4.1	Giannotti <i>et al.</i> [95]
WDLF hint	$\sim 1.1 - 2.8$	Isern <i>et al.</i> [49]

Table 3.1: Summary of stellar bounds on the axion-electron coupling. The hint from Galactic WDLFs (discussed in Section 3.5) is also given on the final line for comparison. The bounds are all reported to 2σ (i.e. 95% confidence), while the hint is reported to 3σ .

rate \dot{T}/T [94]. Observations of the period decrease (\dot{P}/P) for particular white dwarf variables have shown discrepancies with the expected behaviour, with a systematic tendency for the observed value of \dot{P} to be larger than expected. This indicates that these white dwarfs are cooling faster than expected, which could be explained by the emission of axions.

The period of white dwarf variables changes so slowly ($\dot{P}/P \approx 10^{-18} \text{ s}^{-1}$ in most measured cases [85]) that accurately measuring the period change requires observing the white dwarf over a span of decades. The rate of period change has thus only been measured for a few white dwarf variables, despite many white dwarf variables being known. A global analysis of the available measurements of white dwarf variables was found to give a good fit for $g_{aee} = 2.9 \times 10^{-13}$ (with a reduced chi-squared value of 1.1) and to favour an axion interpretation of the cooling anomaly for the white dwarf variables at 2σ [95]. The hint at 1σ has been expressed by Di Luzio *et al.* [84] and Di Luzio *et al.* [85] as $g_{aee} = 2.9_{-0.9}^{+0.6} \times 10^{-13}$, and the 2σ bound from the global analysis of the white dwarf variable data is $g_{aee} \leq 4.1 \times 10^{-13}$. This bound from the white dwarf variables is currently less constraining than either the bound from the tip of the RGB or the Galactic WDLFs.

For easy reference, all of the axion-electron coupling bounds discussed in this section are summarized in Table 3.1. The bounds are all given to

the same confidence of 2σ , which is equivalent to the 95% confidence level. For comparison, the axion hint from Galactic WDLFs is also included. This hint is reported to 3σ in Table 3.1, which is how it was reported by Isern *et al.* [49]. The confidence level of the hint reported by Isern *et al.* [49] is not actually well defined, but the upper limit of the range of values given in Table 3.1 for this hint corresponds to the 3σ bound from Miller Bertolami *et al.* [48]. Note that the 3σ threshold used to report the hint differs from the 2σ threshold used to report the bounds in Table 3.1.

3.5 Hints of Axions from White Dwarf Cooling

Hints of axion emission affecting the cooling of white dwarfs have been suggested based on empirical Galactic WDLFs [47–49]. In particular, it has been suggested that the emission of axions produced via axion bremsstrahlung from electrons could explain a slight excess in the observed white dwarf cooling rates, as indicated by Galactic WDLFs, compared to theoretical models without axions (where cooling is due only to Standard Model neutrinos and photons). The most recent analysis [49] suggests this cooling anomaly could be explained by a DFSZ axion with a mass (and angular parameter) of $m_a \sin^2 \beta \sim 4 - 10$ meV, corresponding to an axion-electron coupling of $g_{aee} \sim (1.1 - 2.8) \times 10^{-13}$ more generally for any axion or axion-like particle model that includes an axion-electron interaction term.

This work argued that including cooling from axion emission in white dwarf cooling models produces theoretical WDLFs that fit observational WDLFs better than current theoretical WDLFs over certain magnitude ranges. Isern *et al.* [49] in particular compared theoretical model predictions of WDLFs to the observational WDLFs for white dwarfs in the Galactic thin disc, thick disc, and stellar halo based on data from the Sloan Digital Sky Survey (SDSS) [44, 45] and SuperCOSMOS Sky Survey (SCSS) [46]. The WDLF excess studied by Isern *et al.* [49] specifically occurs over the bolometric magnitude range $M_{\text{bol}} \sim 8 - 10$, and Isern *et al.* [49] limited their investigation to the range $6 \lesssim M_{\text{bol}} \lesssim 13$. This magnitude range has less theoretical uncertainty than if fainter white dwarfs with $M_{\text{bol}} \gtrsim 13$ are

3.5. *Hints of Axions from White Dwarf Cooling*

included, but it excludes a wealth of data available at larger magnitudes, which are actually more populated with white dwarfs than the limited range of magnitudes investigated in that work. As discussed by Isern *et al.* [49], there are a number of uncertainties in constructing luminosity functions for Galactic white dwarfs, particularly associated with the star formation rate. Isern *et al.* [49] used an improved treatment of the star formation rate compared to the earlier work of Isern *et al.* [47] and Miller Bertolami *et al.* [48], but the star formation rate is still an important source of uncertainty in that analysis. Globular clusters provide populations of white dwarfs with better-controlled white dwarf formation rates, making them good targets for further investigation into these hints of axions.

Chapter 4

Ultramassive white dwarf cooling ages

4.1 Introduction

The *Gaia* mission [42] has revealed the largest number of white dwarfs ever observed in our Galaxy. Before *Gaia*, less than 40,000 confirmed white dwarfs were known, mostly from the spectroscopic Sloan Digital Sky Survey [96, 97]. This sample was largely incomplete and concentrated in the northern hemisphere. Thanks to *Gaia*'s precise parallaxes and photometric measurements, Gentile Fusillo *et al.* [98] were able to identify \sim 260,000 high-confidence white dwarf candidates simply from their position in the HR diagram. The latest data release of *Gaia*, *Gaia* EDR3 [34], provided 20 to 30% better parallax measurements on average, as well as twice as accurate proper motions and better measurements of colour, especially for the bluest objects. These improvements allowed Gentile Fusillo *et al.* [99] to update the catalogue, which now includes \sim 359,000 high-confidence white dwarfs.

The wealth of new data provided by *Gaia* DR2 has challenged our understanding of white dwarf cooling, a process that was thought to be well understood. The *Gaia* DR2 colour-magnitude diagram (CMD) unveiled the presence of the Q branch, a transversal sequence of white dwarfs with mass above $\sim 1 M_{\odot}$ that was not aligned with theoretical cooling sequences [100], which Tremblay *et al.* [31] identified as observational evidence of a white dwarf cooling delay due to core crystallization. While a pile-up of white dwarfs in the region of the Q branch is expected in standard crystallization models due to the cooling delay resulting from the release of latent

4.1. Introduction

heat in the phase transition from a liquid to solid state [31], Cheng *et al.* [43] presented evidence of an additional anomalous cooling delay on the Q branch in *Gaia* DR2 data experienced by massive ($1.08 - 1.23 M_{\odot}$) white dwarfs that could not be explained by conventional models of core crystallization. This has resulted in much research into the details of white dwarf crystallization, particularly ^{22}Ne sedimentation, in an attempt to explain this cooling anomaly [101–105].

Another significant finding arising from *Gaia* DR2 data was the discovery of a star formation burst in the Galactic disc 2 – 3 Gyr ago based on observations of main sequence stars [106]. In addition to presenting evidence of such an event, Mor *et al.* [106] mathematically characterized the time dependent star formation rate (SFR), modelling it as a bounded exponential function plus a Gaussian component that accounts for the star formation burst. Isern [107] noted the significance of this star formation history for the white dwarf luminosity function, reconstructing the SFR for $0.9 - 1.1 M_{\odot}$ *Gaia* DR2 white dwarfs within 100 pc of the Sun, using the luminosity function of Tremblay *et al.* [31] and the BaSTI cooling models for DA white dwarfs and qualitatively comparing this to the SFR of Mor *et al.* [106].

The catalogue of *Gaia* EDR3 white dwarfs recently produced by Gentile Fusillo *et al.* [99] provides us with a powerful tool to re-investigate the cooling of massive white dwarfs in the solar neighbourhood with improved statistical power. In this chapter, we determine the distribution of cooling ages for massive high-probability white dwarf candidates recently identified by the Gentile Fusillo *et al.* [99] catalogue that are within 200 pc of the Sun. Using the publicly available `WD_models` code [108], we use different sets of white dwarf models, including both the carbon/oxygen (C/O) core cooling models of Bédard *et al.* [109] and the oxygen/neon (O/Ne) core cooling models of Camisassa *et al.* [110], to determine the masses and cooling ages of each white dwarf candidate from the *Gaia* EDR3 photometric observations. For each set of models, we construct the distribution of cooling ages for objects with mass in the range $0.95 - 1.25 M_{\odot}$ and binned into subdivisions by mass within this range. We quantitatively assess the statistical significance

of the similarity of the cooling rate distribution of massive *Gaia* EDR3 white dwarfs to the star formation rate of main sequence stars, as characterized by Mor *et al.* [106]. We also illustrate how the presence of a significant fraction of mergers among the most massive white dwarfs, with single progenitors created at the Mor *et al.* [106] rate of star formation, can explain the nearly uniform distribution that we find for their cooling ages.

4.2 Methods

4.2.1 Models

Throughout this chapter, we compare multiple sets of white dwarf cooling models to *Gaia* EDR3 white dwarf data. For our analysis, we use the publicly available `WD_models` package provided by Cheng [108]⁹. We consider the cooling models of the Montreal group [109] and the La Plata group [110–112]. For each set of cooling models, we consider both H-atmosphere and He-atmosphere models. We use the evolutionary models in conjunction with the publicly available Montreal group synthetic colours¹⁰ for pure H and pure He atmosphere models with *Gaia* EDR3 band-pass filters to determine the masses and cooling ages of each white dwarf candidate from the *Gaia* EDR3 photometric observations.

For the Montreal group models [109], we consider three cases of atmosphere composition: thick H envelope evolution models with a pure H atmosphere model (“thick H envelope models”), thin H envelope evolution models with a pure H atmosphere model (“thin H envelope models”), and thin H envelope evolution models with a pure He atmosphere model (“He envelope models”). We collectively refer to the set of thick H envelope models plus He envelope models as the “thick Montreal” models, and analogously refer to the set of thin H envelope models plus He envelope models as the “thin Montreal” models. In all of these cases, the Montreal models

⁹The `WD_models` package is publicly available at https://github.com/SihaoCheng/WD_models.

¹⁰The Montreal group synthetic colours are publicly available at <http://www.astro.umontreal.ca/~bergeron/CoolingModels>.

4.2. Methods

have a C/O core composed of a uniform mixture of C and O in equal parts by mass ($X_C = X_O = 0.5$). It should be noted that this core composition differs from typical model predictions for single stellar evolution (e.g. [21, 113, 114]), which predict lower fractions of C for white dwarfs with mass in the range that we consider. Furthermore, the fixed uniform mixture used for the cores in the Montreal models does not account for element diffusion, which separates the elements into stratified layers.

For the La Plata group set of models, we combine the higher mass O/Ne core models of Camisassa *et al.* [110] with the lower mass C/O core models of Renedo *et al.* [112] and Camisassa *et al.* [111]. While the Camisassa *et al.* [110] models are available for both DA and DB white dwarfs, the Renedo *et al.* [112] models are only available for DA white dwarfs and the Camisassa *et al.* [111] are only available for DB white dwarfs. Thus, for the La Plata group H-atmosphere models, we used the DA white dwarf models of Camisassa *et al.* [110] for masses $\geq 1.1 M_\odot$ and the models of Renedo *et al.* [112] with metallicity $Z = 0.01$ for masses $\leq 0.93 M_\odot$, linearly interpolating between these models for masses between 0.93 and $1.1 M_\odot$. To make the full set of He-atmosphere La Plata group models, we used the DB white dwarf models of Camisassa *et al.* [110] for masses $\geq 1.1 M_\odot$ and the models of Camisassa *et al.* [111] for masses $\leq 1 M_\odot$. We collectively refer to all of these models as the “La Plata” models.

The composition of ultramassive white dwarfs (i.e. white dwarfs with mass above $1.05 M_\odot$) is not well known. The single stellar evolution simulations of Siess [21, 22] and Doherty *et al.* [23] indicate that white dwarfs transition from having C/O cores to O/Ne for masses $\gtrsim 1.05 M_\odot$, and Doherty *et al.* [24] find this transition to correspond to masses between $\sim 1.07 - 1.15 M_\odot$ for solar metallicity. However, Althaus *et al.* [114] showed that ultramassive C/O-core white dwarfs with masses in excess of $1.25 M_\odot$ can readily be produced by models with either high rotation rates or low rates of mass loss on the AGB. Massive white dwarfs can also be formed through the merger of two lighter white dwarfs, and this process could result in white dwarfs with different composition than those produced from the evolution of a single progenitor [114]. While the evolutionary models of

4.2. Methods

Schwab [115] for the merger remnant of two C/O white dwarfs indicate that merger products with final mass $\gtrsim 1.05 M_{\odot}$ have O/Ne cores, other work (e.g. [116, 117]) has found scenarios in which ultramassive remnants with C/O cores may be produced from double white dwarf mergers.

For ultramassive C/O white dwarf models, Althaus *et al.* [114] noted that the composition of those produced by single stellar evolution and those produced by double white dwarf mergers is expected to be different. White dwarfs produced by single stellar evolution are expected to be born with the most abundant elements already largely separated into distinct layers due to element diffusion during pre-white-dwarf evolution, while the merger of two white dwarfs mixes the elements and thus causes the resulting white dwarf to (re)start cooling with a more uniformly mixed composition. Even for single stellar evolution models, there is significant uncertainty in the expected ratio of carbon to oxygen in C/O-core white dwarfs due to factors such as the uncertainty in the $^{12}C(\alpha, \gamma)^{16}O$ rate and the treatment of convection in progenitor models (e.g. [19, 20, 118]), in addition to the choice of rotation and mass loss on the AGB [114].

We adopt an agnostic approach to the white dwarf core composition by considering models with a variety of compositions and performing the same analysis for each set of models to determine the best-fitting model for each mass bin. The models considered in this work are particularly relevant for comparison with the work of Cheng *et al.* [43], which found an anomalous cooling delay on the Q branch for massive *Gaia* DR2 white dwarfs. The Bédard *et al.* [109] models are the most recent version of the older Fontaine *et al.* [68] models used by Cheng *et al.* [43], who used the thick H models of Fontaine *et al.* [68] for white dwarf masses $\leq 1.05 M_{\odot}$ and the models of Camisassa *et al.* [110] for white dwarf masses $\geq 1.1 M_{\odot}$.

4.2.2 Data

We used the main catalogue of Gentile Fusillo *et al.* [99], which consists of *Gaia* EDR3 white dwarfs with reliable parallax measurements. From this catalogue, we selected sources within a distance of 200 pc and for which the

4.2. Methods

probability of being a white dwarf exceeded 90%, as determined by the `Pwd` parameter of the catalogue (for details concerning how this probability was determined, see [99]).

For each source in the catalogue, Gentile Fusillo *et al.* [99] provide the three parameters `chisq_H`, `chisq_He`, and `chisq_mixed`, indicating the goodness-of-fit of models with different atmosphere compositions: pure H, pure He, and mixed H/He atmospheres, respectively. In our work, we classified each white dwarf as having an atmosphere composition corresponding to the atmosphere model with the smallest (non-empty) chi-squared value. White dwarfs for which all three chi-squared parameters were empty were classified as having an “unknown” atmosphere. For a given set of models, this classification was used to determine whether the H-atmosphere or He-atmosphere models were used to determine the mass and cooling age of each source from the observed photometry. He-atmosphere models were used for both the pure He and mixed H/He atmosphere white dwarfs. Following the recommendations of Gentile Fusillo *et al.* [99], we used H-atmosphere models for white dwarfs with unknown atmosphere composition.

Figure 4.1 depicts our observational sample of white dwarfs identified by Gentile Fusillo *et al.* [99] within 200 pc. Superimposed on the diagram are contours of constant mass and cooling age calculated using the thin H envelope models of Bédard *et al.* [109]. The other models considered in our work show similar trends. Furthermore, the thresholds and completion of core crystallisation for carbon/oxygen and oxygen/neon white dwarfs as identified by Bauer *et al.* [101] are shown by solid and dashed black lines. These contours are similar to those identified by Tremblay *et al.* [31] for carbon/oxygen white dwarfs and to those by Camisassa *et al.* [110] for oxygen/neon white dwarfs. For the masses that we consider, the crystallisation occurs between the absolute G-band magnitude of 12 and 14. The magnitude and colour have been de-reddened using the mean A_V values given in the catalogue and the prescription given by Gentile Fusillo *et al.* [99] to convert this extinction in the Johnson V band to extinctions in *Gaia* bands. To illustrate the shift in the CMD caused by reddening, we show the *Gaia* extinction vector for a colour excess of $E(B - V) = 0.1$ (the median extinc-

4.2. Methods

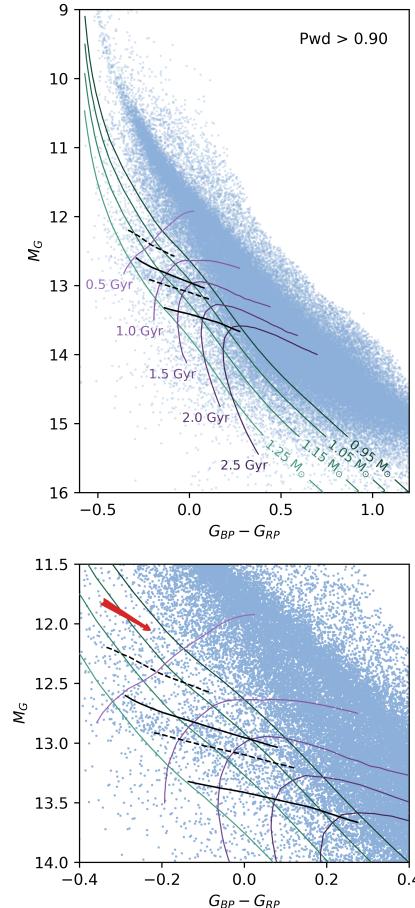


Figure 4.1: Upper: the *Gaia* CMD of white dwarf candidates within 200 pc from Gentile Fusillo *et al.* [99]. Lower: focus on the high-mass white dwarfs. The effect of interstellar extinction with $E(B - V) = 0.1$ is indicated by the red arrow. Superimposed on the observations are contours of equal mass, from 0.95 to 1.25 solar masses (right to left), and contours of equal cooling age, from 0.5 to 2.5 Gyr (top to bottom), according to the thin H atmosphere white dwarf cooling models of Bédard *et al.* [109]. The four black lines show significant points in core crystallization, as identified by Bauer *et al.* [101]. The dashed black lines indicate the region of O/Ne crystallisation, showing the points at which 20% (top) and 80% (bottom) of the O/Ne core is frozen. Analogously, the solid black lines show the points at which 20% (top) and 80% (bottom) of the C/O core is frozen.

4.2. Methods

tion for stars within 200 pc from the *Gaia* EDR3) as a red arrow in the lower panel of Fig. 4.1. De-reddening undoes the effect illustrated by this arrow, shifting the white dwarfs back in the opposite direction. As the extinction vector is approximately parallel to the phase transition lines for the range of white dwarf masses that we will consider, the location of a white dwarf relative to the phase transition lines is approximately independent of extinction. While reddening will make a white dwarf appear slightly less massive by a few hundredths of a solar mass for $E(B - V) = 0.1$, since we will be studying white dwarfs binned into samples one tenth of a solar mass wide, our conclusion should not be very sensitive to the exact treatment of reddening.

A key concern in determining the distribution of white dwarf cooling ages is to determine the volume sampled as a function of the luminosity of the white dwarf [119]. Although we could estimate this using the magnitude limits of the Gentile Fusillo *et al.* [99] catalogue, we would still not have a reliable estimate of the completeness rate at the faint end of the catalogue; consequently, we will study the completeness of the catalogue as a function of absolute G-band magnitude and distance using a variant of the Schmidt [120] estimator. If the population that we are sampling is uniform surrounding the Sun (which is a good approximation within 200 pc), the number of white dwarfs that we detect should increase linearly with the volume sampled.

Figure 4.2 depicts the cumulative number of white dwarfs detected as a function of the volume sampled (normalised by the total volume V_{\max} , which is the volume of a sphere with a 200 pc radius). The solid curves of different colours correspond to samples binned by absolute G-band magnitude into bins with a width of one magnitude and centred on the values indicated by the legend, with the values of these bin centres ranging from $M_G = 9$ to $M_G = 17$. White dwarfs with absolute magnitudes of 14 and brighter are detected with high completeness all the way out to 200 pc. For inherently fainter white dwarfs, the samples become incomplete closer to the Sun. The incompleteness manifests itself as the downward curvature of the cumulative distribution from the linear relation.

We choose the limiting volume to be where this curvature develops. In

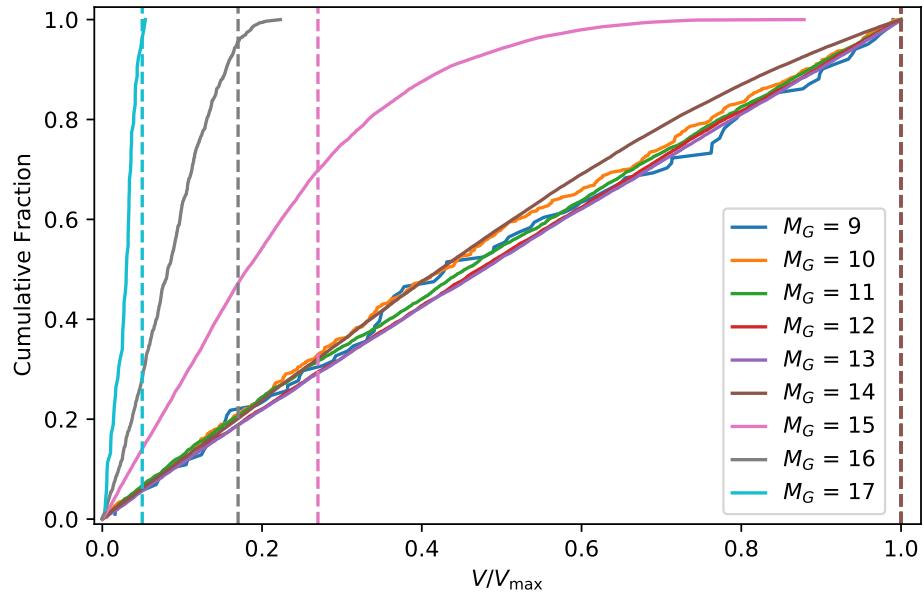


Figure 4.2: Completeness limits of white dwarf samples as a function of absolute magnitude. The legend indicates the central value of each magnitude bin. The vertical dashed lines indicate the volume limits of the complete samples that we consider as a function of absolute magnitude.

4.3. Results

particular we choose the limiting volume to be where the slope of the cumulative distribution deviates by more than 5% relative to its value for small volumes. Consequently, within these limiting volumes, the samples are nearly complete. We only perform this adjustment for samples that we estimate to be less than 85% complete within 200 pc, that is, for $M_G \geq 15$. The three leftmost vertical lines from left to right show the limiting volume that we choose for magnitude values¹¹ of $M_G = 17, 16$ and 15 . For $M_G \leq 14$, we do not reduce the volume relative to the total volume of the 200 pc sample, and we illustrate this choice as a limiting volume equal to V_{\max} shown by the rightmost vertical line in Fig. 4.2. We restrict our sample to objects within their magnitude dependent completeness-limiting volume, $V_{\lim}(M_G)$. The reduced sampling volume of the faint white dwarfs is then corrected by assigning a weight of $V_{\max} / V_{\lim}(M_G)$ to objects for which $V_{\lim}(M_G) < V_{\max}$ and a weight of 1 to objects for which $V_{\lim}(M_G) \geq V_{\max}$.

4.3 Results

In Fig. 4.3 we show the joint distribution of masses and cooling ages of *Gaia* EDR3 white dwarfs according to different sets of models. To generate these distributions, the white dwarfs were weighted to correct for the reduced sampling volume required to make the sample complete. For each set of models, the mass and cooling age of each object in our reduced sample was determined by linearly interpolating the models, using either the H-atmosphere or He-atmosphere models depending on the atmosphere classification of the source (as determined by the `chisq_H`, `chisq_He`, and `chisq_mixed` parameters of the Gentile Fusillo *et al.* [99] catalogue). Later in this work we sort the white dwarfs into three mass bins of width $0.1 M_{\odot}$, centred on the values 1.0 , 1.1 , and $1.2 M_{\odot}$. The edges of these mass bins are indicated in Fig. 4.3 by light grey dotted vertical lines. For illustrative purposes and to help orient the reader, we also show the core crystallization lines of Bauer *et al.*

¹¹These magnitude values correspond to the centres of the magnitude bins for the solid curves plotted in Fig. 4.2. We linearly interpolate between the limiting volumes chosen for these reference points to determine the limiting volume for (and thus assign a weight to) each white dwarf in the full sample.

4.3. Results

[101] as red lines in Fig. 4.3. The solid lines indicate C/O crystallization thresholds of 20% (bottom, solid) and 80% (top, solid) frozen, while the dashed lines similarly indicate O/Ne crystallization thresholds of 20% (bottom, dashed) and 80% (top, dashed) frozen. Bauer *et al.* [101] calculated these crystallization thresholds using models with pure He atmospheres and *Gaia* DR2 filters. For consistency with the Bauer *et al.* [101] models used to calculate these threshold, we used He-atmosphere models to convert those (*Gaia* DR2) colour-magnitude crystallization thresholds to curves in mass and cooling age. The models used for this conversion consisted of the He-envelope evolution models of each of the three sets of models considered in this work along with the synthetic colours for *Gaia* DR2 filters, rather than those for the EDR3 filters used to determine the masses and cooling ages of objects in our observational sample. By comparison with the crystallization lines in Fig. 4.1, we can identify where the Q branch region of the CMD has been mapped to in Fig. 4.3, as the bulk of the Q branch region occurs between the lines of 20% C/O frozen and 80% O/Ne frozen.

Several trends are immediately apparent from Fig. 4.3, regardless of which set of models is considered. The decreasing density from left to right across each plot indicates that there are more low-mass white dwarfs than high-mass white dwarfs. Equally striking is that, as the cooling age increases (moving from bottom to top in each plot), the density of white dwarfs gradually increases until it reaches a peak well after the age at which 80% of a C/O core is frozen. This corresponds to an excess of stars on the Q branch and below it on the CMD. This effect is much more pronounced for the lower mass white dwarfs, with masses $\sim 0.95 - 1.15 M_{\odot}$. Since white dwarfs with mass $\gtrsim 1.15 M_{\odot}$ are much less numerous than white dwarfs on the lower end of the mass range shown in Fig. 4.3, the scale of the plots makes it difficult to see the trend with cooling age for these most massive white dwarfs; however, our analysis later in this work will reveal that their cooling ages are much closer to uniformly distributed.

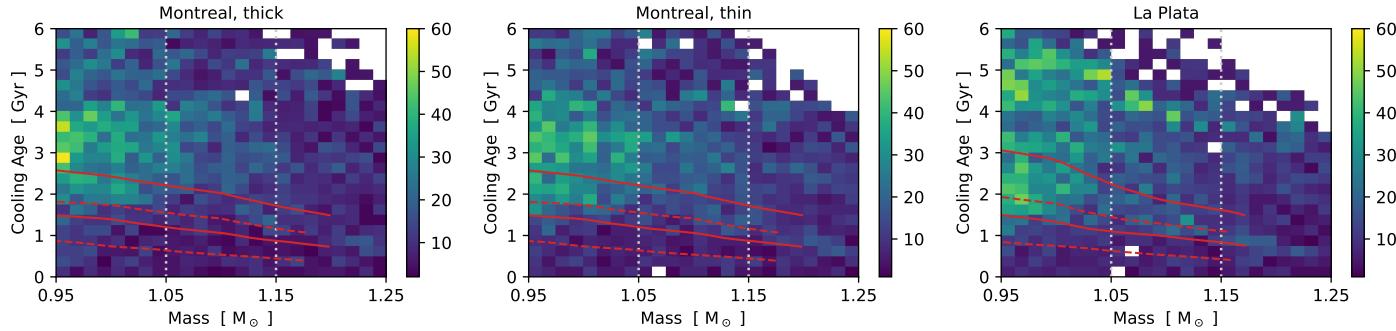


Figure 4.3: Joint distribution of mass and cooling age for massive *Gaia* EDR3 white dwarfs according to different sets of models, weighted to correct for reduced sampling volume. For each subplot, the set of models used to calculate the cooling ages and masses from *Gaia* EDR3 photometry is indicated at the top of the subplot. From left to right, these are the thick Montreal models, the thin Montreal models, and the La Plata models (see main text for details). The fill colour of each bin indicates the corrected count found in that bin, with corresponding value given by the colour bar and white denoting that no objects were found in that bin. The red lines indicate the C/O (solid) and O/Ne (dashed) freezing lines of Bauer *et al.* [101], shown only for reference. From bottom to top, these lines indicate 20% O/Ne frozen, 20% C/O frozen, 80% O/Ne frozen, and 80% C/O frozen. The vertical dotted lines indicate the boundaries between the three mass bins considered in this work.

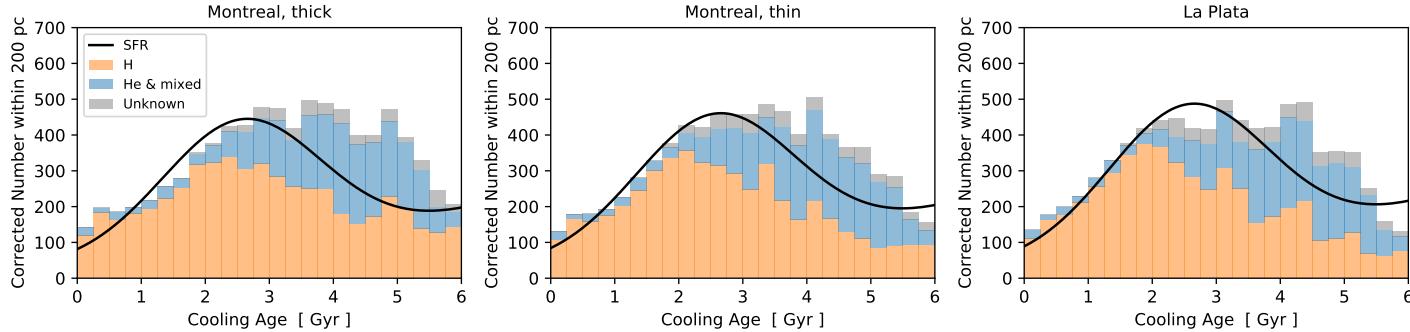


Figure 4.4: Distribution of cooling ages (in Gyr) of $0.95 - 1.25 M_{\odot}$ white dwarfs within 200 pc of the Sun, weighted to correct for reduced sampling volume. This sample was sorted into subgroups based on atmospheric composition, as determined by the fits of Gentile Fusillo *et al.* [99]. White dwarfs with atmospheric composition classified as either pure He or mixed H/He were combined into one group (shown in blue), and He-atmosphere models were used to determine their ages and masses. H-atmosphere models were used to determine the ages and masses of white dwarfs with atmospheric composition classified as either pure H or unknown, but the H-atmosphere white dwarfs (shown in orange) and the unknown-atmosphere white dwarfs (shown in grey) were kept as separate groups. The (corrected) counts for each of these groups are stacked such that the top of the stacked distribution indicates the total (corrected) count from combining all three groups. The black line indicates the star formation rate of main sequence stars determined by Mor *et al.* [106], normalized to the number of white dwarfs with cooling age ≤ 3 Gyr.

4.3. Results

If the models are correct and the white dwarf birthrate were constant, then the distribution of cooling ages at a given mass would be expected to be uniform. The non-uniform density seen in Fig. 4.3 (for $0.95 - 1.15 M_{\odot}$ white dwarfs) indicates that either the white dwarf birthrate is not constant or that these white dwarfs are cooling more slowly than predicted by the models in the region where the density increases. In the latter scenario, the observed pile-up of white dwarfs would imply a cooling delay that spans both the Q branch region and well beyond it, with the apparent delay in fact being much more prominent for white dwarfs that have already finished their Q branch stage of evolution. The peak of this white dwarf pile-up actually roughly coincides with the burst of star formation found by Mor *et al.* [106] for *Gaia* DR2 main sequence stars, with the general trend of the density gradually increasing with cooling age to a peak for the first ~ 3 Gyr seeming to follow the corresponding time-varying star formation rate found in that work. This strongly suggests that the non-uniform distribution of the white dwarf cooling ages seen in Fig. 4.3 is simply the imprint of the time-varying star formation history of their main sequence progenitors rather than an indication of a cooling delay relative to the standard white dwarf cooling models. Progenitor stars that produce massive white dwarfs in the mass range that we consider ($0.95 - 1.25 M_{\odot}$) evolve through their pre-white-dwarf stages of evolution on timescales that are negligible compared to the cooling ages we consider, with pre-white-dwarf lifetimes of less than 100 Myr [121, 122]. Thus, if the massive white dwarfs that we consider are born from single progenitors, then the distribution of their cooling ages should follow the star formation rate of their main sequence progenitors.

The distribution of cooling ages for white dwarfs in the mass range $0.95 - 1.25 M_{\odot}$, according to different sets of models and weighted to correct for the completeness-limiting reduced sampling volume, are shown in Fig. 4.4. These plots are the result of marginalizing the joint distributions of Fig. 4.3 over white dwarf mass for the given mass range. The different colours indicate the contribution of white dwarfs with different atmosphere classifications to the total distribution. Nearly all of the youngest white dwarfs, with cooling age $\lesssim 2$ Gyr, are classified as having a H-

4.3. Results

atmosphere; while there are significant numbers of both H-atmosphere and He-atmosphere white dwarfs with older cooling ages. White dwarfs with unknown atmosphere composition only make up a very small fraction of the total white dwarfs at any cooling age, so our treatment of these white dwarfs should not significantly impact our results.

In each panel of Fig. 4.4, we also show the star formation rate determined by Mor *et al.* [106] for *Gaia* DR2 main sequence stars, which we have normalized to the number of white dwarfs with a cooling age of ≤ 3 Gyr. Up to about 3 Gyr, the distribution of massive *Gaia* EDR3 white dwarfs tracks the shape of the SFR distribution, up to a small lag that should be expected due to the time required for a main sequence star to become a white dwarf. Heyl *et al.* [122] have measured the lag between star formation and the formation of a one-solar-mass white dwarf to be less than 100 Myr from the Pleiades, so neglecting this lag is a reasonable approximation. If the white dwarf birthrate were constant, this distribution should be approximately uniform. Instead we see an accumulation of massive white dwarfs at cooling ages that coincide with the look-back time for the star formation burst found by Mor *et al.* [106]. Beyond about 3 Gyr, the white dwarf cooling age distribution appears to track the SFR less well. This could indicate that the star formation burst is actually broader than found for the fiducial case of Mor *et al.* [106]. The discrepancy for old white dwarfs could alternately be due to the poorer statistics of the oldest white dwarfs or an imperfect correction for the reduced sampling volume of these stars, which are typically fainter and less complete than the younger white dwarfs. The weighting only begins to affect the distributions for cooling ages older than about 3 Gyr, so we can reliably compare the cooling age distribution to the SFR up to at least 3 Gyr even if the latter situation is the issue. We examine the divergence between the weighted and unweighted distribution more closely in our analysis of the cumulative cooling age distributions below.

To investigate the effect of white dwarf mass, M_{WD} , on the distribution of cooling ages, we sort the massive white dwarfs into three mass bins: $0.95 - 1.05 M_{\odot}$, $1.05 - 1.15 M_{\odot}$, and $1.15 - 1.25 M_{\odot}$. The cumulative distributions for white dwarfs in each of these mass bins, according to the three sets

4.3. Results

of models we consider, are shown in Fig. 4.5. We show the white dwarf distribution both with and without weighting to correct for the reduced sampling volume for which the *Gaia* EDR3 observations are complete. In each panel of Fig. 4.5, we show the cumulative distribution corresponding to the SFR of Mor *et al.* [106], normalized to the unweighted white dwarf distribution at 3 Gyr. For reference, we also show a uniform distribution with the same normalization. From the top and middle rows of this figure, we see that the distributions for white dwarfs in the two lightest mass bins, $0.95 - 1.05 M_{\odot}$ and $1.05 - 1.15 M_{\odot}$, approximately follow the time-dependent SFR of main sequence stars [106]. In stark contrast, the distribution of the heaviest white dwarfs, in the mass bin $1.15 - 1.25 M_{\odot}$ (bottom row), does not seem to follow the SFR at all; instead, the cooling ages of the heaviest white dwarfs are nearly uniformly distributed.

To statistically assess the similarity between the observed cumulative cooling age distributions and the cumulative SFR of [106], we perform a series of one-sample Kolmogorov–Smirnov (KS) tests. We test the null hypothesis that the empirical cumulative distribution of a particular sample of white dwarf cooling ages is equal to the analytic cumulative distribution function of the SFR. These cumulative distributions will be equivalent if the models are correct and the white dwarf birthrate is proportional to the SFR of the main sequence progenitors, *i.e.* if the probability density of producing a white dwarf with a particular cooling age is proportional to the SFR at the equivalent look-back time. Viewing the cooling age of a white dwarf as a random variable drawn from an underlying probability distribution, which here is taken to be proportional to the SFR under the null hypothesis, the p-value from the one-sample KS tests quantifies the probability of randomly drawing a sample from the proposed underlying distribution for which the sample distribution differs from the underlying distribution by at least as much as what is seen for the observed empirical distribution. A small p-value indicates that the observed white dwarf cooling ages are unlikely to have been drawn from the proposed SFR distribution. Viewed as a goodness-of-fit test, a small p-value indicates a poor fit of the SFR to the empirical distribution of cooling ages.

4.3. Results

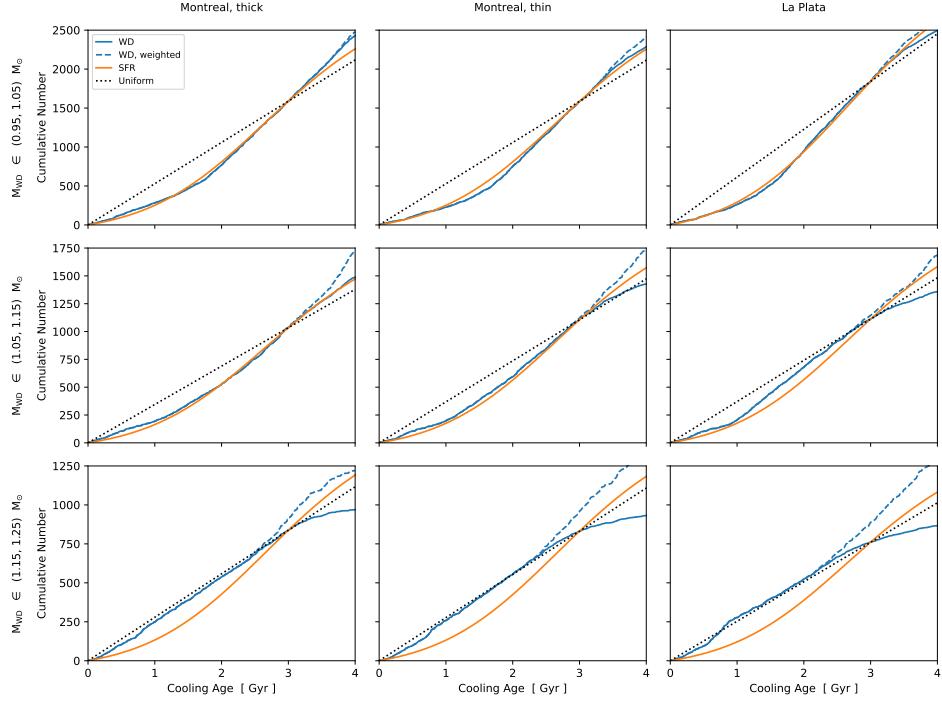


Figure 4.5: Cumulative distributions of cooling ages (in Gyr) of massive white dwarfs in different mass bins according to the different sets of cooling models considered in this work. Each column corresponds to a particular set of models: thick Montreal (left), thin Montreal (centre), and La Plata (right). Each row corresponds to a particular white dwarf mass bin: $0.95 - 1.05 M_{\odot}$ (top), $1.05 - 1.15 M_{\odot}$ (middle), and $1.15 - 1.25 M_{\odot}$ (bottom). For each subplot, we show both the unweighted (solid blue line) and weighted (dashed blue line) white dwarf distributions. For comparison, we also show the cumulative distribution for the star formation rate of Mor *et al.* [106] (solid orange line) and a uniform distribution (dotted black line), both normalized to the (unweighted) white dwarf distribution at 3 Gyr.

4.3. Results

We performed separate KS tests for each mass bin and each set of models, as well as for the full mass range $0.95 - 1.25 M_{\odot}$ encompassing all bins. As can be seen in Fig. 4.5, the weighted distribution starts to deviate from the unweighted distribution at cooling ages $\gtrsim 3$ Gyr for the first two mass bins, and slightly earlier for the most massive bin. The KS tests do not account for the weighting of the white dwarfs, so we can only use these tests to draw statistically meaningful conclusions for distributions of white dwarfs young enough that the weighted distribution does not yet differ significantly from the unweighted distribution. For each sample we consider in a KS test, the sample includes only white dwarfs with cooling age $\leq t_{\max}$. For each combination of model set and mass range, we perform KS tests for three choices of maximum cooling age: 2 Gyr, 3 Gyr, and 4 Gyr. The p-values resulting from all of these tests are summarized in Table 4.1. These different samples allow us to consider different options for balancing the improved statistical power of larger sample numbers with the error for the oldest white dwarfs from not accounting for the reduced sampling volume. For the lightest two mass bins, we consider the sub-samples for which the white dwarf cooling ages are ≤ 3 Gyr to give the best balance between these two considerations.

For each of the two lightest mass bins, there are two sets of models that give results consistent with the SFR for the sample with $t_{\max} = 3$ Gyr, though the optimal models are different for each bin. The La Plata models give the best results for the $0.95 - 1.05 M_{\odot}$ mass bin, with a p-value of 0.0266 for $t_{\max} = 3$ Gyr, while the thick Montreal models give the best results for the $1.05 - 1.15 M_{\odot}$ mass bin, with a p-value of 0.0138 for $t_{\max} = 3$ Gyr (and 0.0707 for $t_{\max} = 4$ Gyr). For white dwarfs in the heaviest mass bin, $1.15 - 1.25 M_{\odot}$, none of the models give white dwarf distributions consistent with the SFR (all p-values $< 10^{-4}$). This result is in line with Fig. 4.5, which shows that these heaviest white dwarfs have approximately uniformly distributed cooling ages, particularly for cooling ages above ~ 1 Gyr, regardless of which set of models is considered. To assess this apparently uniform trend, we performed another series of one-sample KS tests for just white dwarfs in this heaviest bin to compare their cooling age distribution to a uniform

4.3. Results

Mass (M_{\odot})	t_{\max} (Gyr)	Montreal, thick	Montreal, thin	La Plata
0.95 – 1.05	2	0.0005	$< 10^{-4}$	$< 10^{-4}$
	3	0.0177	$< 10^{-4}$	0.0266
	4	$< 10^{-4}$	$< 10^{-4}$	$< 10^{-4}$
1.05 – 1.15	2	$< 10^{-4}$	0.0103	0.0179
	3	0.0138	0.0208	$< 10^{-4}$
	4	0.0707	$< 10^{-4}$	$< 10^{-4}$
1.15 – 1.25	2	$< 10^{-4}$	$< 10^{-4}$	$< 10^{-4}$
	3	$< 10^{-4}$	$< 10^{-4}$	$< 10^{-4}$
	4	$< 10^{-4}$	$< 10^{-4}$	$< 10^{-4}$
0.95 – 1.25	2	$< 10^{-4}$	$< 10^{-4}$	0.0033
	3	$< 10^{-4}$	0.0008	$< 10^{-4}$
	4	$< 10^{-4}$	$< 10^{-4}$	$< 10^{-4}$

Table 4.1: The p-values from one-sample KS tests comparing the (unweighted) cumulative distribution of cooling ages of white dwarfs in different mass bins, according to different models, to the cumulative SFR of Mor *et al.* [106] for samples restricted to different maximum cooling ages, t_{\max} (in Gyr).

4.3. Results

distribution. For distributions over the cooling age range $1 - 2.5$ Gyr, all of the models give distributions that are highly consistent with a uniform distribution, with p-values for KS tests comparing these distributions to a uniform distribution of 0.9803, 0.6321, and 0.2523 for the thick Montreal, thin Montreal, and La Plata models, respectively. The distributions for the first 1 Gyr (*i.e.* restricted to the cooling age range $0 - 1$ Gyr) are in tension with a uniform distribution, having p-values of 0.0014, 0.0005, and $< 10^{-4}$. This indicates that it could take ~ 1 Gyr for the uniform distribution to develop.

A possible explanation of the approximately uniform distribution of the most massive white dwarfs is that a large fraction of these white dwarfs formed due to mergers. Temmink *et al.* [123] have used binary population synthesis simulations to show that the assumption of single stellar evolution for white dwarfs that have undergone mergers typically leads to an underestimate of the white dwarf age, particularly for white dwarfs more massive than $0.9 M_{\odot}$ that are produced by double white dwarf mergers. Temmink *et al.* [123] further estimated that 30 – 50% of white dwarfs with mass above $0.9 M_{\odot}$ are the product of binary mergers. The earlier binary population synthesis simulations of Bogomazov and Tutukov [124] predicted even higher merger fractions for the most massive white dwarfs, with over 50% of white dwarfs more massive than $1.1 M_{\odot}$ predicted to be the product of double white dwarf mergers. These theoretical predictions of a high fraction of merger products among massive white dwarfs have been supported by observations. Based on the kinematic, magnetic, and rotational properties of white dwarfs in the Montreal White Dwarf Database 100 pc sample, Kilic *et al.* [125] found direct observational evidence that at least 32% of the ultra-massive white dwarfs in their sample, with mass $> 1.3 M_{\odot}$, were the product of binary mergers. Cheng *et al.* [126] have shown strong evidence of the presence of double white dwarf merger products among *Gaia* DR2 high mass white dwarfs and determined the cooling delay time distribution of double white dwarf merger products for multiple mass bins using population synthesis simulations.

For the mass ranges we consider, if most white dwarfs in the sample

4.3. Results

were formed from the evolution of a single progenitor, then the distribution of white dwarf cooling ages would be expected to directly track the SFR of their main sequence progenitors, similar to what we see for the lightest two mass bins. However, white dwarfs that are the result of a double white dwarf merger event will have a younger apparent cooling age than predicted by single white dwarf evolution models. If a large fraction of white dwarfs in the sample are the result of double white dwarf merger events, then their cooling ages will be approximately distributed according to the convolution of the SFR of the single progenitors and the distribution of cooling delay times due to mergers.

To estimate the cooling age distribution if most of the white dwarfs in our most massive bin are the result of double white dwarf mergers, we convolve the Cheng *et al.* [126] distribution of cooling delay times for $1.14 - 1.24 M_{\odot}$ (approximately corresponding to our mass bin) with the Mor *et al.* [106] SFR. The result of this convolution is shown as the solid black curve in Fig. 4.6, normalized to the number of white dwarfs with apparent cooling age ≤ 3 Gyr in our heaviest bin according to the thick Montreal models. The dashed coloured curves in Fig. 4.6 show the SFR by itself, without convolution with the merger delay time distribution, with each curve normalized to the number of white dwarfs with cooling age ≤ 3 Gyr in the mass bin of the corresponding colour. Each normalization was determined using the most statistically consistent model for that mass bin according to the KS tests. Note that normalizing the SFR to a different mass bin only affects the height of the curve, not the shape.

The cooling age distributions for the observed *Gaia* EDR3 white dwarfs in each mass bin, according to the most consistent set of models for each bin as determined by the earlier KS tests, are shown in Fig. 4.6 as histograms with 0.5 Gyr bin width. These distributions were calculated using the same weighting as the previous figures in this work to correct for the reduced sampling volume of the complete samples. The La Plata models were used for the lightest bin and the thick Montreal models for the two heaviest bins. Comparison of the dashed green and solid black curves in Fig. 4.6 illustrates the effect on the cooling age distribution of accounting

4.3. Results

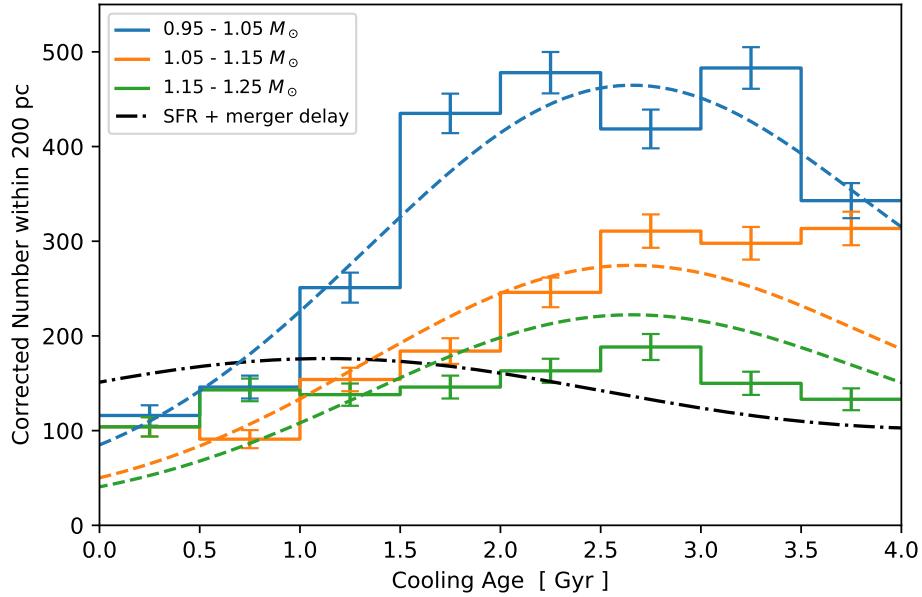


Figure 4.6: Distribution of white dwarf cooling ages by mass bin, weighted to correct for the reduced sampling volume of the complete sample. For the lightest two mass bins, we show the distributions made using the set of models that was most consistent with the star formation rate based on the results of KS tests. For the heaviest bin, we used the thick Montreal models. The dash-dot black curve indicates the distribution resulting from convolving the star formation rate of Mor *et al.* [106] with the merger delay distribution of Cheng *et al.* [126] appropriate for our heaviest mass bin; this curve has been normalized to the number of white dwarfs in the heaviest bin with cooling age ≤ 3 Gyr. The dashed lines indicate the star formation rate of Mor *et al.* [106] analogously normalized to the number of white dwarfs with cooling age ≤ 3 Gyr in the mass bin with the corresponding colour. From top to bottom these are the lightest to heaviest mass bins.

4.3. Results

for the merger delay, shifting the peak associated with the star formation burst event to younger cooling ages and overall flattening the distribution, resulting in a distribution much closer to uniform and in better agreement with the observed distribution for white dwarfs in the most massive bin.

We model the cooling age distribution of ultramassive white dwarfs as a linear combination of the distribution for white dwarfs produced by direct stellar evolution and the distribution for white dwarfs produced by the merger of two lighter white dwarfs. The distribution for white dwarfs that evolved from a single progenitor is taken to be the SFR found by Mor *et al.* [106] for main sequence stars, while the distribution for merger products is taken to be the convolution of the SFR of Mor *et al.* [106] and the cooling delay distribution determined by Cheng *et al.* [126]. We determine the empirical cooling age distribution for $1.15 - 1.25 M_{\odot}$ white dwarfs according to each set of white dwarf cooling models by constructing weighted histograms of 0.5 Gyr bin width for cooling ages spanning 0 to 4 Gyr, and we fit this model to each empirical distribution using non-linear least squares regression. The results of these fits are summarized in Table 4.2. We also perform fits for the two limiting cases in which the ultramassive white dwarfs are i) only the product of mergers and ii) only the product of single stellar evolution, labelled respectively as “Only Mergers” and “Only Singles” in Table 4.2.

The goodness-of-fit is assessed using the reduced chi-squared statistic, χ_{ν}^2 , which allows us to compare the fits for the different distributions despite differences in the degrees of freedom, ν . Fits for the limiting cases of only merger products and only single stellar evolution products each have $\nu = 7$, while fits for the fiducial combination scenario have $\nu = 6$. Even after accounting for the different degrees of freedom, the combination scenario gives a better fit than either of the limiting cases, regardless of which white dwarf cooling model is used to determine the masses and cooling ages. Among the different white dwarf cooling models, the thick Montreal models give the best fit for all of the cooling age distributions.

For the best-fitting combination of merger and direct formation products according to each set of white dwarf cooling models, we determine the

4.3. Results

	Montreal, thick	Montreal, thin	La Plata
χ^2_ν	Combination	1.4	4.9
	Only Mergers	10.5	17.6
	Only Singles	11.3	12.6
f_m	Today	0.805	0.756
	Average	0.511	0.439

Table 4.2: Fitting results for cooling age distribution of $1.15 - 1.25 M_\odot$ white dwarfs. The goodness-of-fit, scaled to the degrees of freedom (ν), for each combination of origin model and white dwarf cooling model is indicated by the reduced chi-squared value, χ^2_ν . We fit origin scenarios in which all of the white dwarfs are the result of mergers (“Only Mergers”, $\nu = 7$), all are the result of single stellar evolution (“Only Singles”, $\nu = 7$), and some fraction are merger products (“Combination”, $\nu = 6$). The fraction of merger products, f_m , forming today and averaged over 4 Gyr is given for each best-fit combination.

fraction of merger products, f_m , as a function of cooling age. In Table 4.2, we summarize f_m by reporting the fraction of merger products among white dwarfs forming today (at 0 Gyr cooling age) and the fraction of merger products among white dwarfs that have formed over the past 4 Gyr. It should be noted that f_m is not a parameter of the fit; it is a function that is determined from the result after the fit has been performed. The function $f_m(t_{\text{cool}})$ is then evaluated at the particular instant $t_{\text{cool}} = 0$ Gyr to get the value today (reported in the row labelled “Today” in Table 4.2), and it is averaged over the cooling age range $0 - 4$ Gyr to get the the fraction of merger products for white dwarfs that have formed over the past 4 Gyr (labelled “Average” in Table 4.2). The best-fitting cooling models, the thick Montreal models, give merger product fractions of 0.805 ± 0.037 for white dwarfs forming today and 0.511 ± 0.059 for white dwarfs that have formed over the last 4 Gyr, in line with the theoretical estimates of Bogomazov and Tutukov [124] and the observational results of Kilic *et al.* [125], though somewhat larger than the observational results of Cheng *et al.* [126]. From the χ^2_ν values summarized in Table 4.2, we see that the case of the combination formation scenario

4.3. Results

with the thick Montreal cooling models fits the data well, with $\chi^2_\nu = 1.4$, but none of the other cases produce a good fit.

In Fig. 4.7, we show the results of fitting the different ultramassive white dwarf formation scenarios to the empirical distribution according to the thick Montreal cooling models. The empirical distribution is shown as the green histogram. The solid red curve shows the cooling age distribution for the best-fitting combination of merger and direct formation products. Also shown are the best-fitting cooling age distributions for the limiting cases in which all of the ultramassive white dwarfs are the result of direct formation (dashed green curve) and all are the result of mergers (dash-dot black curve), which are simply proportional to the SFR and the convolution of the SFR and merger delay distribution, respectively. Using a separate y-axis scale, Fig. 4.7 shows f_m for the best-fitting combination as a solid grey line. The scale for all of the cooling age distributions is indicated by the y-axis on the left side of the plot, while the scale for the fraction of merger products is indicated by a separate y-axis on the right side of the plot. Comparison of the cooling age distributions visually illustrates the better fit of the optimal combination formation scenario relative to the limiting cases, as quantified by the χ^2_ν values listed in Table 4.2. For the optimal combination, we also see from Fig. 4.7 that the fraction of merger products is a function of the white dwarf cooling age, with a maximum value for white dwarfs that are forming today.

4.3. Results

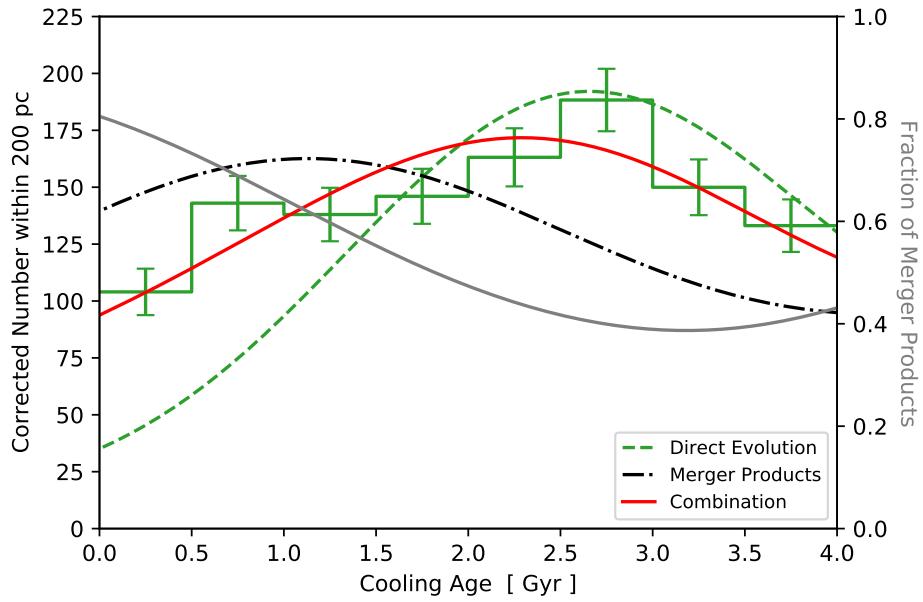


Figure 4.7: Ultramassive white dwarf ($1.15 - 1.25 M_{\odot}$) formation mechanisms. In the best-fitting combination (red solid curve, $\chi^2_{\nu} = 1.4$ with six degrees of freedom) of mergers and direct formation, 51% of ultramassive white dwarfs that formed over the past 4 Gyr, formed through mergers. Among the ultramassive white dwarfs forming today, 80% are the result of mergers. The limiting cases of only mergers (black dash-dot curve) and only single stellar evolution (dashed green curve) are shown for comparison. The green histogram is the empirical distribution. The fraction of merger products in the best-fitting combination is shown as the solid grey curve, with y-axis on the right.

4.4 Conclusions

In this chapter, we have shown that the number density of *Gaia* EDR3 white dwarfs in the mass range $0.95 - 1.15 M_{\odot}$ gradually increases with cooling age over the first ~ 3 Gyr of cooling time. This pile-up of white dwarfs extends over the entire age range associated with the Q branch and continues well beyond it, reaching a peak after the end of core crystallization. If the white dwarf birthrate were constant, this apparent excess of white dwarfs would indicate an anomalous cooling delay relative to the models. However, the discovery of a time-varying star formation rate for *Gaia* DR2 main sequence stars [106] suggests that the white dwarf birthrate is likely to also vary with time, with the expectation that the distribution of cooling ages for the massive white dwarfs that we consider closely follows this star formation rate if the white dwarfs predominantly originate from single progenitor stars and if the models are correct. We indeed find that white dwarfs in the mass bins $0.95 - 1.05 M_{\odot}$ and $1.05 - 1.15 M_{\odot}$ have cumulative cooling age distributions that are statistically consistent with the expectation from the star formation rate observed for main sequence stars [106]. We do not see statistical evidence for an anomalous cooling delay in these results.

For slightly more massive white dwarfs, in the mass bin $1.15 - 1.25 M_{\odot}$, we find that the distribution of their cooling ages is consistent with a uniform distribution for all sets of models considered over the cooling age range $1 - 2.5$ Gyr, and this is not consistent with the star formation rate of Mor *et al.* [106]. This could indicate that a large fraction of these most massive white dwarfs formed from double white dwarf mergers. The binary population synthesis results of Cheng *et al.* [126] for white dwarfs in this mass range indicate a merger delay time distribution that, when convolved with the star formation rate of the main sequence progenitors, pushes the peak density to earlier white dwarf cooling ages and overall flattens the distribution.

We model the cooling age distribution of $1.15 - 1.25 M_{\odot}$ white dwarfs as a linear combination of the distribution of single stellar evolution products and the distribution of double white dwarf merger products, and we show that this formation model provides a good fit to the empirical distribution of

4.4. Conclusions

the photometric cooling ages determined by C/O-core white dwarf cooling models. Depending on the cooling model, we find that about 69 – 80% of the white dwarfs forming today and 36 – 51% of the white dwarfs that have formed over the last 4 Gyr are the product of double white dwarf mergers.

We note that we find no evidence for a substantial cooling delay in the numbers of massive white dwarfs detected in *Gaia* EDR3 when one takes the star formation history of the Galaxy into consideration. For the most massive white dwarfs we find that the bulk were formed in mergers. The C/O white dwarf models seem to fit better to the number of white dwarfs than O/Ne models. This may indicate that these merger remnants are actually C/O white dwarfs. It alternately could be that the merger delay distribution proposed by Cheng *et al.* [126] is not realized in the solar neighbourhood, that the star formation history is not precisely that given by Mor *et al.* [106], or that the models for O/Ne white dwarfs require some minor revision.

Chapter 5

Ultramassive white dwarf kinematics

5.1 Introduction

The kinematics of stars can encode important information about their history and origin. The dynamical history of the Galaxy is imprinted on the kinematics-age relations of stars in the solar neighbourhood, such as the trend for the velocity dispersion of stars from the local Galactic disc to increase with age due to disc heating, as seen in the age-velocity dispersion relation (AVR; [127–140]). Since the epoch of the HIPPARCOS mission [141], the kinematics of stars in the solar neighbourhood have been extensively studied [131, 142–148], and the precise astrometry of the *Gaia* mission [42], the successor to HIPPARCOS, has further advanced these studies [139, 140, 149–153].

Gaia EDR3 [34] has been used to study the kinematics of white dwarfs in the solar neighbourhood, showing that the kinematic structure of local white dwarfs [153] is similar to that of local main-sequence stars [150, 151, 154]. For example, Raddi *et al.* [140] performed a detailed three-dimensional kinematic analysis of a sample of 3133 white dwarfs from *Gaia* EDR3 with radial velocity measurements from either *Gaia* or cross-matched spectroscopic observations, finding that their sample, which was mostly located within ~ 300 pc of the Sun, consisted of $\sim 90 - 95\%$ thin disc stars, $\sim 5 - 10\%$ thick disc stars, and a few isolated white dwarfs and halo members. For the thin disc members, Raddi *et al.* [140] determined their AVRs and found them to agree with previous AVR results that were determined from dif-

5.1. Introduction

ferent samples of LAMOST-*Gaia* FGK-type stars without radial velocity information [139].

Gaia has also revealed structure in the CMD of white dwarfs in the solar neighbourhood, which has implications for our understanding of white dwarf cooling. *Gaia* DR2 revealed the so-called Q branch [100], an observed sequence of massive white dwarfs in the CMD that is transversal to the theoretical cooling sequences for white dwarfs with mass $\gtrsim 1 M_\odot$ and coincides with the region of white dwarf crystallization [31]. The phase transition from liquid to solid state during the process of crystallization results in the release of latent heat [25] and other energy associated with element sedimentation [26–29] that slows the rate of white dwarf cooling, and this cooling delay results in a pile-up of white dwarfs that has been identified with the Q branch [31].

While a cooling delay due to white dwarf crystallization had been predicted for over 50 yr prior to *Gaia* DR2, Cheng *et al.* [43] presented evidence from *Gaia* DR2 for an extra cooling delay of ~ 8 Gyr experienced by $\sim 6\%$ of ultramassive ($1.08 - 1.23 M_\odot$) white dwarfs on the Q branch that is not explained by standard crystallization models. By looking at the number density of white dwarfs as a function of the age inferred by white dwarf cooling models, Cheng *et al.* [43] noted an excess number of ultramassive white dwarfs on the Q branch above the expected pile-up from the crystallization delay. Furthermore, Cheng *et al.* [43] noted that for some ultramassive *Gaia* DR2 white dwarfs on the Q branch there was a discrepancy between the dynamic age, inferred from the transverse velocity using the AVR^s of Holmberg *et al.* [133] and [138] for the local thin and thick discs, and the photometric isochrone age, inferred from the photometry using white dwarf cooling models. Among $\sim 1.08 - 1.23 M_\odot$ white dwarfs, they found a population moving anomalously fast relative to the velocity expected from the AVR for their photometric ages. Cheng *et al.* [43] argued that this discrepancy could only be explained by an extra cooling delay, in addition to the delays due to crystallization and double white dwarf mergers, and suggested modifications to the treatment of ^{22}Ne sedimentation in white dwarf cooling models as a possible mechanism. This prompted many efforts to explain

5.1. Introduction

this cooling anomaly through modifications to white dwarf cooling models [101–105, 155–157].

In Chapter 4, we used the improved *Gaia* EDR3 data to re-investigate this cooling anomaly by analysing the number density distribution of the photometric cooling ages for massive ($0.95–1.25 M_{\odot}$) white dwarfs identified in the white dwarf catalogue of Gentile Fusillo *et al.* [99]. We considered a variety of publicly available white dwarf cooling models, including both the oxygen-neon core models of Camisassa *et al.* [110] and carbon-oxygen core models of Bédard *et al.* [109] with different envelope thicknesses. These are all standard cooling models, by which we mean that they do not attempt to reproduce the anomalous ~ 8 Gyr cooling delay proposed by Cheng *et al.* [43]. For each set of cooling models considered, the masses and ages of the white dwarfs were determined from the *Gaia* photometry, and the resultant cooling age distributions were compared to the expectation from the time-varying star formation rate observed by Mor *et al.* [106] for *Gaia* DR2 main-sequence stars. Mor *et al.* [106] discovered a star formation burst in the local Galactic disc 2 – 3 Gyr ago, which, if a uniform white dwarf birthrate were assumed, would appear as an excess number of massive white dwarfs around this age produced through single stellar evolution. In Chapter 4, it was indeed found that, under the assumption of a uniform birthrate, there was an apparent excess of white dwarfs both along the Q branch and below it, coinciding with the burst of star formation seen in the star formation history of main-sequence stars.

As part of the detailed analysis in Chapter 4, the sample was further subdivided into three equally spaced mass bins, and the number density distribution of photometric cooling ages according to each set of cooling models was statistically compared to the Mor *et al.* [106] star formation rate. For the two lightest mass bins, $0.95–1.05$ and $1.05–1.15 M_{\odot}$, it was found that standard cooling models could produce cooling age distributions that were statistically consistent with the expected distribution from the star formation rate observed by Mor *et al.* [106]. For the most massive bin, $1.15–1.25 M_{\odot}$, it was found that the cooling age distribution was well fitted by a linear combination of the distribution expected for single

5.1. Introduction

stellar evolution products (taken to be proportional to the star formation rate) and the distribution expected for the product of double white dwarf merger products (taken to be the convolution of the star formation rate and the merger delay time distribution calculated by Cheng *et al.* [126] using population synthesis simulations) when approximately 40–50% of the $1.15 - 1.25 M_{\odot}$ white dwarfs that formed over the past 4 Gyr were produced through double white dwarf mergers. From this analysis, it was found in Chapter 4 that the photometric cooling age distribution of ultramassive white dwarfs could be explained by accounting for the time-dependent star formation rate and the presence of a large fraction of merger products among $1.15 - 1.25 M_{\odot}$ white dwarfs, without needing to invoke an anomalous ~ 8 Gyr cooling delay.

In this chapter, we follow up the work of Chapter 4 with a kinematic analysis of the transverse motions of the same sample of $0.95 - 1.25 M_{\odot}$ white dwarfs in the solar neighbourhood that was considered in Chapter 4, subdivided into the same mass bins of $0.95 - 1.05$, $1.05 - 1.15$, and $1.15 - 1.25 M_{\odot}$ and with masses and cooling ages determined using white dwarf cooling models found in Chapter 4 to produce photometric cooling age distributions consistent with the expectation from the star formation history. For each of these mass bins, we estimate the transverse velocity dispersion as a function of age using different estimators and analyse the distributions of the separate transverse velocity components for several age ranges. We compare the empirical distributions to the expectation from the AVRs measured for main-sequence stars in the local thin disc and find a population of anomalously fast-moving ultramassive white dwarfs on the Q branch with kinematic features that are not explained by an extra cooling delay in white dwarfs originating from the local thin disc.

The structure of this chapter is as follows. The data and models used in this chapter, along with the formal procedure for determining transverse motions, are described in Section 5.2. The results are presented in Section 5.3, and the implications for the origin of anomalously fast-moving ultramassive white dwarfs on the Q branch are discussed in Section 5.4. Finally, the results and discussion are summarized in Section 5.5.

5.2 Methods

5.2.1 Data Sample

We use the data and white dwarf cooling models described in Chapter 4. The data consist of *Gaia* EDR3¹² white dwarfs identified in the Gentile Fusillo *et al.* [99] catalogue with at least 90% probability, as determined by the `Pwd` parameter of the catalogue, and located within 200 pc of the Sun. The magnitude and colour observations were de-reddened following the prescription of Gentile Fusillo *et al.* [99]. We used extinctions of $A_G = 0.835A_V$, $A_{G_{\text{BP}}} = 1.139A_V$, and $A_{G_{\text{RP}}} = 0.650A_V$ for the G -, G_{BP} -, and G_{RP} -band magnitudes, respectively, with A_V values given by the `meanAV` parameter of the catalogue. We calculated the absolute G -band magnitude, M_G , according to the expression

$$M_G = G - 5 \log_{10} \left(\frac{100 \text{ mas}}{\bar{\omega}} \right) - A_G , \quad (5.1)$$

where $\bar{\omega}$ is the parallax in mas, using the values of the catalogue parameters `parallax` and `phot_g_mean_mag` for $\bar{\omega}$ and G , respectively.

The completeness of this 200 pc sample as a function of M_G was analysed in Chapter 4 and found to be volume limited for white dwarfs with $M_G \geq 15$. The limiting volume as a function of magnitude, $V_{\text{lim}}(M_G)$, was determined in Chapter 4 using a variant of the Schmidt [120] estimator as follows. The cumulative number distribution as a function of sampling volume was constructed for magnitude-binned sub-samples, and V_{lim} was determined for each magnitude bin by finding the volume above which the distribution began to deviate from the linear relation expected for a complete sample (and which was realized at smaller volumes). The limiting volume for each white dwarf was then calculated by linearly interpolating between these magnitude-binned V_{lim} values as a function of M_G , taking the limiting volume for magnitudes $M_G \geq 15$ to simply be the total volume, V_{max} , of the 200 pc sample (i.e. the volume of a sphere with a radius of

¹²Although there is a more recent data release, *Gaia* Data Release 3 [158], it does not contain new astrometric information compared to EDR3.

200 pc).

Following the procedure of Chapter 4, we restrict our sample to only include white dwarfs within their completeness-limiting volumes and correct for the magnitude-dependent reduced sampling volume by assigning each remaining white dwarf a weight of

$$w_i = \begin{cases} V_{\max} / V_{\lim,i} & \text{if } V_{\lim,i} < V_{\max} \\ 1 & \text{otherwise} \end{cases}, \quad (5.2)$$

where i is a data index denoting the i th white dwarf, $V_{\lim,i}$ is the limiting volume for the magnitude of that white dwarf, and w_i is the weight assigned to that white dwarf.

5.2.2 Cooling and Atmosphere Models

Using the publicly available `WD_MODELS`¹³ package provided by Cheng [108], we use white dwarf cooling models to determine the photometric age and mass of each white dwarf from its absolute G -band magnitude, M_G , and colour, $G_{\text{BP}} - G_{\text{RP}}$. To determine masses and cooling ages from these photometry measurements, the atmosphere composition of the white dwarf is an important factor to consider. As described in detail in Chapter 4, we classified each source in our sample as having either a pure-H, pure-He, or mixed atmosphere based on the best-fitting results from Gentile Fusillo *et al.* [99], as determined by the chi-squared values provided in the catalogue for their model fits: `chisq_H`, `chisq_He`, and `chisq_mixed`. Each source was classified as having an atmosphere composition corresponding to whichever model gave the smallest chi-squared value. Sources for which all three chi-squared values were empty were assumed to have pure-H atmospheres.

Each set of models from `WD_MODELS` considered in Chapter 4 included both H-atmosphere and He-atmosphere models. The evolutionary models were combined with the corresponding publicly available Montreal group

¹³The `WD_MODELS` package is publicly available at https://github.com/SihaoCheng/WD_models.

5.2. Methods

synthetic colours¹⁴ for pure-H and pure-He atmosphere models with *Gaia* EDR3 band-pass filters to convert the photometry measurements of each source into mass and cooling age values. The synthetic colours map the magnitude and colour to the effective temperature and surface gravity, respectively T_{eff} and $\log g$, and the evolutionary models map T_{eff} and $\log g$ to cooling age and mass. Models of the appropriate atmosphere composition were used for each source as determined by the atmosphere classification from the Gentile Fusillo *et al.* [99] catalogue described above. He-atmosphere models were used for sources classified as having pure-He or mixed atmospheres, while H-atmosphere models were used for sources classified as having pure-H or unknown atmospheres.

In Chapter 4, we assessed the consistency of the photometric cooling age distribution determined using various models available through WD_MODELS with the star formation history observed for *Gaia* DR2 main-sequence stars [106]. In this chapter, we use only the thick Montreal white dwarf cooling models [109]. These models gave the best-fitting cooling age distribution for the $1.05 - 1.15$ and $1.15 - 1.25 M_{\odot}$ mass bins in Chapter 4. For the $0.95 - 1.05 M_{\odot}$ mass bin, the La Plata models [110–112] and thick Montreal models were both determined to produce cooling age distributions consistent with the star formation history and yielded similar p -values, so for consistency we use the thick Montreal models for all three mass bins. Using the La Plata models instead of the thick Montreal models for the lightest mass bin produces similar results.

After the publication of the results of Chapter 4 (see Fleury *et al.* [2], the La Plata group produced cooling models for ultramassive white dwarfs ($\gtrsim 1.15 M_{\odot}$) with carbon-oxygen cores [159]. These models implement more realistic physics in the core of the white dwarfs than the Bédard *et al.* [109] models, such as element diffusion and abundance ratios determined by progenitor evolution simulations, though there is still a great deal of uncertainty regarding the abundances expected for carbon-oxygen white dwarfs [19, 20, 114, 118]. While an analysis of the photometric cooling age distribu-

¹⁴The Montreal group synthetic colours are publicly available at <http://www.astro.umontreal.ca/~bergeron/CoolingModels>.

tions determined by the Camisassa *et al.* [159] models in comparison to the star formation rate (as was done in Chapter 4) is worthwhile, our choice of cooling model will not change the salient results presented in this chapter. Using the Camisassa *et al.* [159] models instead of the Bédard *et al.* [109] models would shift the particular ages at which the key features of the velocity dispersion appear, but it would not remove these features. Likewise, using oxygen-neon core models (e.g. [110]) instead of carbon-oxygen core models would shift the inferred cooling ages and, to a lesser extent, masses of our white dwarf sample, but it would not affect the values of the velocities and therefore would not eliminate the key features presented in this chapter.

5.2.3 Kinematics

5.2.3.1 Streaming Motion

The observed velocity of a star in the solar neighbourhood consists of the local (average) streaming velocity, which includes both the average motion due to Galactic rotation and reflex solar motion relative to the local standard of rest, plus the peculiar velocity of the star due to random motion. In this work, we are interested in the random motion of white dwarfs about the local mean, so we correct the observed velocities for both Galactic rotation and solar motion.

Olling and Dehnen [160] give a concise derivation of the local streaming motion in the Galactic plane to first order in a two-dimensional Taylor series expansion of the mean velocity field about the location of the Sun. To first order in this expansion, the Galactic rotation can be parametrized by the four Oort constants A , B , C , and K , with $C = K = 0$ in the case of axisymmetry. The proper motions corresponding to the streaming velocity in Galactic coordinates are

$$\bar{\mu}_{\ell*} = (U_0 \sin \ell - V_0 \cos \ell) \bar{\omega} + \cos b (A \cos 2\ell - C \sin 2\ell + B) \quad (5.3)$$

$$\begin{aligned} \bar{\mu}_b = & [(U_0 \cos \ell + V_0 \sin \ell) \sin b - W_0 \cos b] \bar{\omega} \\ & - \sin b \cos b (A \sin 2\ell + C \cos 2\ell + K) , \end{aligned} \quad (5.4)$$

5.2. Methods

where the proper motions $\bar{\mu}_{\ell*}$ and $\bar{\mu}_b$ are given in units of $\text{km s}^{-1} \text{kpc}^{-1}$, $\bar{\mu}_{\ell*} \equiv \bar{\mu}_\ell \cos b$, $\bar{\omega}$ is the inverse distance in units of kpc^{-1} (corresponding to a parallax angle measured in mas), ℓ is the Galactic longitude, b is the Galactic latitude, and (U_0, V_0, W_0) is the velocity of the Sun relative to the local streaming in a heliocentric Cartesian coordinate system with $\hat{\mathbf{x}}$ pointing in the direction $(\ell, b) = (0^\circ, 0^\circ)$, $\hat{\mathbf{y}}$ in the direction $(\ell, b) = (90^\circ, 0^\circ)$, and $\hat{\mathbf{z}}$ in the direction of $b = 90^\circ$. The components of the transverse streaming velocity in the $\hat{\ell}$ and \hat{b} directions, in units of km s^{-1} , are

$$(\bar{v}_\ell, \bar{v}_b) = (\bar{\mu}_{\ell*}, \bar{\mu}_b) / \bar{\omega}. \quad (5.5)$$

We calculate the streaming velocities using values reported in the literature for both the solar motion and the Oort constants. For the solar motion, we use the peculiar solar velocity measurements of Wang *et al.* [161], with values of $U_0 = 11.69 \text{ km s}^{-1}$, $V_0 = 10.16 \text{ km s}^{-1}$, and $W_0 = 7.67 \text{ km s}^{-1}$. For the Galactic rotation, we assign the Oort constants the values measured by Bovy [149] using *Gaia* DR1 main-sequence stars from the Tycho-*Gaia* Astrometric Solution catalogue, with values of $A = 15.3 \text{ km s}^{-1} \text{kpc}^{-1}$, $B = -11.9 \text{ km s}^{-1} \text{kpc}^{-1}$, $C = -3.2 \text{ km s}^{-1} \text{kpc}^{-1}$, and $K = -3.3 \text{ km s}^{-1} \text{kpc}^{-1}$.

A conversion factor of $\kappa = 4.74047$ is needed to convert proper motions given in mas yr^{-1} to units of $\text{km s}^{-1} \text{kpc}^{-1}$. Including this conversion factor, the observed velocities (in km s^{-1}) in the $\hat{\ell}$ and \hat{b} directions are

$$(v_\ell^{(\text{obs})}, v_b^{(\text{obs})}) = \kappa \left(\mu_{\ell*}^{(\text{obs})}, \mu_b^{(\text{obs})} \right) / \bar{\omega}. \quad (5.6)$$

Gaia only provides proper motions in ICRS coordinates, so we convert the observed proper motions from ICRS coordinates to Galactic coordinates using the ASTROPY¹⁵ package for PYTHON [162, 163].

We calculate the peculiar velocities, v_ℓ and v_b , of the white dwarfs by

¹⁵<http://www.astropy.org>

5.2. Methods

correcting the observed velocities for solar motion and Galactic rotation

$$v_\ell = v_\ell^{(\text{obs})} - \bar{v}_\ell \quad (5.7)$$

$$v_b = v_b^{(\text{obs})} - \bar{v}_b . \quad (5.8)$$

These two velocities, together with the sky coordinates and parallax, completely specify the random transverse motion of our sample relative to the local standard of rest.

5.2.3.2 Random Transverse Motion

Since most sources in *Gaia* EDR3 do not have radial velocity measurements, we perform our analysis of the kinematics using only the transverse motion. Following the procedure of Dehnen and Binney [143], we relate the transverse motion to the full three-dimensional motion through use of a projection operator \mathbf{P} that projects the three-dimensional space velocity \mathbf{v} onto the celestial sphere.

We write this formulation explicitly using matrix notation with vectors as column matrices and work in Cartesian Galactic coordinates in which $\hat{\mathbf{x}}$ points towards the Galactic centre ($\ell = 0^\circ, b = 0^\circ$), $\hat{\mathbf{y}}$ points in the direction of Galactic rotation ($\ell = 90^\circ, b = 0^\circ$), $\hat{\mathbf{z}}$ points towards the North Galactic Pole ($b = 90^\circ$), and $\mathbf{v} = (U, V, W)^\text{T}$. Let $\hat{\mathbf{r}} = (\cos b \cos \ell, \cos b \sin \ell, \sin b)^\text{T}$ be the unit position vector pointing to the location of the star. Then the projection operator in matrix notation is

$$\mathbf{P} \equiv \mathbf{I} - \hat{\mathbf{r}} \hat{\mathbf{r}}^\text{T} , \quad (5.9)$$

where \mathbf{I} is the 3×3 identity matrix. The tangent velocity vector \mathbf{v}_\perp is then related to \mathbf{v} by

$$\mathbf{v}_\perp = \mathbf{P} \mathbf{v} , \quad (5.10)$$

where \mathbf{v}_\perp is the component of \mathbf{v} in the tangent plane at the angular position of the star.

The rectangular components of the tangential velocity vector can be

5.2. Methods

written explicitly as

$$\mathbf{v}_\perp = \begin{pmatrix} U_\perp \\ V_\perp \\ W_\perp \end{pmatrix} = \begin{pmatrix} -\sin \ell v_\ell - \cos \ell \sin b v_b \\ \cos \ell v_\ell - \sin \ell \sin b v_b \\ \cos b v_b \end{pmatrix}, \quad (5.11)$$

and the magnitude of the transverse velocity, $v_t = |\mathbf{v}_\perp|$, is simply given by

$$v_t = \sqrt{v_\ell^2 + v_b^2}. \quad (5.12)$$

We calculate v_t and the components of \mathbf{v}_\perp for the white dwarfs in our sample using v_ℓ and v_b as given by Eqs. (5.7) and (5.8). Since the mean local motion due to both Galactic rotation and solar motion has already been removed, the velocities here are the peculiar velocities, corresponding to the random motion about the mean.

If the full three-dimensional space velocities (including radial motion) have a diagonal dispersion tensor $\boldsymbol{\sigma}^2 = \text{diag}(\sigma_U^2, \sigma_V^2, \sigma_W^2)$ and the positions of the white dwarfs are uncorrelated with the velocities, then the dispersion of the transverse velocity can be quantified by

$$\sigma_t^2 = E [\text{Tr}(\mathbf{P} \boldsymbol{\sigma}^2 \mathbf{P}^T)], \quad (5.13)$$

where Tr denotes the matrix trace, which is equivalent to the sum of eigenvalues of the matrix, and E denotes the expectation value for a particular distribution of sky coordinates, which can be estimated by a statistic such as the sample mean.

Throughout this chapter, we compare the results we find for the white dwarf transverse kinematics relations with the expectation assuming the dispersions of their full three-dimensional velocities follow the AVRs determined by Holmberg *et al.* [133] for main-sequence stars from the local thin disc. While more recent determinations of these AVRs exist (e.g. [139]), they are similar to the AVRs of Holmberg *et al.* [133] and do not change our results. As evidence for a cooling anomaly among ultramassive white dwarfs was presented by Cheng *et al.* [43] with reference to the thin disc AVRs

5.3. Results

of Holmberg *et al.* [133], using the Holmberg *et al.* [133] AVR_s facilitates comparison of our results with earlier work on this topic.

The sample for which the Holmberg *et al.* [133] AVR_s were determined was volume-complete to a distance of ~ 40 pc [131–133], so both the Holmberg *et al.* [133] sample and our 200 pc sample are local enough that bias against populations with large scale height is not a major concern in comparing our results to the Holmberg *et al.* [133] AVR_s. This was verified by comparing the Holmberg *et al.* [133] AVR_s to the Yu and Liu [139] AVR_s for a sample of thin disc (metal-rich) stars with $|z| < 270$ pc, which we found to be similar enough that the main features in the comparison of our results to the AVR_s were unchanged. The similarity of the Yu and Liu [139] AVR_s to the Holmberg *et al.* [133] AVR_s demonstrates the negligible impact of scale height bias for the results presented in this chapter.

5.3 Results

5.3.1 Summary Statistics

Fig. 5.1 shows various probes of the AVR for *Gaia* EDR3 white dwarfs in three different mass bins, $0.95 - 1.05 M_{\odot}$ (blue), $1.05 - 1.15 M_{\odot}$ (orange), and $1.15 - 1.25 M_{\odot}$ (green). The masses and cooling ages were calculated from the observed photometry using the single stellar evolution white dwarf cooling models of Bédard *et al.* [109] which were determined in Chapter 4 to produce cooling age distributions that are statistically consistent with the star formation rate observed for *Gaia* DR2 main-sequence stars [106].

The upper panel of Fig. 5.1 shows the fraction of fast-moving white dwarfs with $v_t > 60$ km s⁻¹ as a function of photometric cooling age. The empirical distributions, shown as histograms, were constructed using cooling age bins of 0.5 Gyr width over the range 0 – 4 Gyr. The fraction of fast-movers in each age bin was calculated using weighted counts for that bin, where each white dwarf was assigned a weight according to the procedure of Chapter 4 to correct for the reduced sampling volume that ensures completeness.

5.3. Results

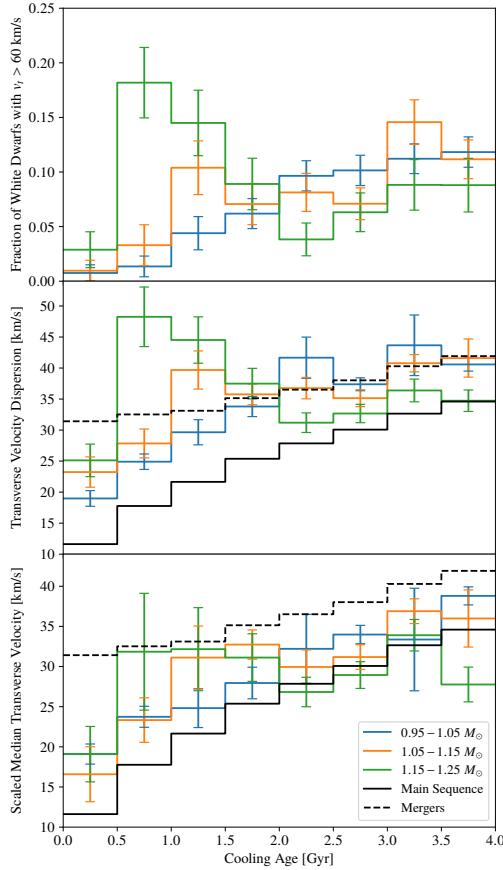


Figure 5.1: Upper panel: Fraction of fast-moving massive white dwarfs for three mass bins. The fast-moving white dwarfs are defined as those with transverse velocity greater than 60 km s^{-1} . Middle panel: Transverse velocity dispersion accounting for uncertainties in proper motion. Lower panel: Scaled median transverse velocity (robust estimator of velocity dispersion). In the middle and lower panels, the solid black line traces the results for main-sequence stars from Holmberg *et al.* [133] with the same sky distribution as the white dwarfs, and the dashed black lines trace the results from main-sequence stars including a merger delay described by the merger delay distribution of Cheng *et al.* [126]. The main-sequence curve (solid black) is the expectation if the $1.15 - 1.25 M_\odot$ white dwarfs are all produced by single progenitors, while the mergers curve (dashed black) is the expectation in the limit that 100% of those white dwarfs are produced by double white dwarf mergers.

5.3. Results

The fraction of fast-moving ultramassive white dwarfs in *Gaia* EDR3 as a function of photometric cooling age shows a trend with cooling age that is similar to what Cheng *et al.* [43] found for *Gaia* DR2 white dwarfs, though we see a less pronounced peak for young $1.05 - 1.15 M_{\odot}$ white dwarfs on the Q branch than what was found by Cheng *et al.* [43] for this bin. For the $0.95 - 1.05 M_{\odot}$ white dwarfs, the fraction of fast-movers tends to increase with cooling age, which is the expected behaviour if the white dwarf velocities follow an AVR similar to what has been observed for main-sequence (G and F dwarf) stars in the thin disc [133]. Of particular note is the distribution for $1.15 - 1.25 M_{\odot}$ white dwarfs, for which the fraction of young fast-movers in the age range $0.5 - 1.5$ Gyr is larger than the fraction in either the younger or older age bins. This corresponds to an excess of young fast-movers on the Q branch relative to the expectation based on the AVR observed for stars in the thin disc [133]. While we see a hint of a similar excess in the $1.5 - 2.0$ Gyr bin of the $1.05 - 1.15 M_{\odot}$ white dwarfs, we do not see a clear signal of an excess of fast-movers within the uncertainties. This difference between our distribution and that of Cheng *et al.* [43] for the $1.05 - 1.15 M_{\odot}$ white dwarfs could be due to the difference in colours and proper motions between the *Gaia* EDR3 and DR2 measurements. The updated colour measurements for very blue objects in particular affect the masses inferred for those objects.

The middle and lower panels of Fig. 5.1 show estimates of the transverse velocity dispersion as a function of age using different estimators of σ_t , one based on the sample mean in each bin (middle panel) and the other based on the sample median (lower panel), the latter of which is more robust to outliers. Random variation in the observed values of v_t will arise both from the intrinsic randomness in the motion of the white dwarfs about the local mean (i.e. the peculiar motion of the white dwarfs) that we are interested in and from random measurement error. We correct for the latter effect in both methods of estimating σ_t from the observations. The random error for each value of v_t was calculated by propagating the standard errors and correlations provided by *Gaia* for the proper motions and parallax from which the observed transverse velocity (without corrections for Galactic rotation

5.3. Results

or solar motion) was calculated. The error was propagated using the expression $e_{v_t,(\text{obs})}^2 = \mathbf{J} \boldsymbol{\Sigma} \mathbf{J}^T$, where $\boldsymbol{\Sigma}$ is the covariance matrix, \mathbf{J} is the Jacobian matrix, and $e_{v_t,(\text{obs})}$ is the random error for the observed transverse velocity.

The middle panel of Fig. 5.1 shows the transverse velocity dispersion calculated using the mean for each age (and mass) bin with a correction for random measurement error according to the expression

$$\sqrt{\langle v_t^2 \rangle - \langle e_{v_t,(\text{obs})}^2 \rangle}, \quad (5.14)$$

where v_t is the transverse velocity after correcting for Galactic rotation and reflex solar motion (i.e. calculated from the sky-projected velocity after subtracting the mean velocity), $e_{v_t,(\text{obs})}$ is the standard error for the transverse velocity calculated from observed proper motions and parallax without subtracting the mean Galactic rotation and reflex solar motion, and the angle brackets $\langle \rangle$ denote a weighted sample average (e.g. $\langle v_t^2 \rangle = (\sum_i w_i v_t^2) / (\sum_i w_i)$ with i an index labeling the data points in the sample).

The lower panel of Fig. 5.1 shows the scaled median velocity corrected for random measurement error, which is calculated for each joint mass and photometric cooling age bin using the formula

$$\sqrt{\frac{\tilde{v}_t^2}{\ln 2} - \langle e_{v_t,(\text{obs})}^2 \rangle}, \quad (5.15)$$

where \tilde{v}_t is the sample median of v_t , determined for each bin using the weighted cumulative fraction for white dwarfs in that bin. The motivation for scaling the median is that we empirically find the the transverse velocities to be approximately Rayleigh-distributed. A Rayleigh distribution parametrized by mode σ has a mean of $\sigma\sqrt{\pi/2}$, median of $\sigma\sqrt{2\ln 2}$, and variance of $\sigma^2(2 - \pi/2)$, yielding a transverse velocity dispersion of $\sigma\sqrt{2}$ and $\tilde{v}_t^2 = (\ln 2)\langle v_t^2 \rangle$.

In both the middle and lower panels of Fig. 5.1, the solid black line shows the expected relation according to the AVRs of main-sequence (G and F dwarf) stars found by Holmberg *et al.* [133] for the sky distribution of the sample of $1.15 - 1.25 M_\odot$ white dwarfs in each age bin. This curve assumes

5.3. Results

that all the white dwarfs formed through single stellar evolution, and thus that none of the white dwarfs formed through mergers. The transverse velocity dispersion value for each bin was calculated by taking the sample average of the sum of the eigenvalues of $\mathbf{P}\boldsymbol{\sigma}^2\mathbf{P}^T$ (equivalent to the trace of this matrix),

$$\langle \text{Tr}(\mathbf{P}\boldsymbol{\sigma}^2\mathbf{P}^T) \rangle, \quad (5.16)$$

where \mathbf{P} is the projection operator that projects the velocity \mathbf{v} onto the tangent plane of the celestial sphere such that $\mathbf{v}_\perp = \mathbf{P}\mathbf{v}$, $v_t = |\mathbf{v}_\perp|$, and $\boldsymbol{\sigma}^2 = \text{diag}(\sigma_U^2, \sigma_V^2, \sigma_W^2)$ is the dispersion tensor of the full three-dimensional space velocities before projection. The values of σ_U , σ_V , and σ_W are calculated for each white dwarf using the AVR^s determined by Holmberg *et al.* [133] with the photometric cooling age of that white dwarf. The AVR^s of Holmberg *et al.* [133] for the three velocity components and the total velocity (giving σ_U , σ_V , σ_W , and σ_{tot} as a function of age) each follow a power law with exponents of 0.39, 0.40, 0.53, and 0.40 for U , V , W , and total dispersion, respectively.

We similarly determined the expected relation for the Holmberg *et al.* [133] AVR^s if all of the white dwarfs are the product of double white dwarf mergers, which is shown as the dashed black curve in the middle and lower panels of Fig. 5.1. This corresponds to the limit in which the merger fraction is 100% for white dwarfs that formed over the last 4 Gyr. Instead of using the photometric cooling age inferred directly using the white cooling models applicable for single stellar evolution to calculate the predicted σ_U , σ_V , and σ_W values from the AVR^s, the relation for merger products was calculated using the expected true age after accounting for a merger delay time, with the expectation value calculated using the merger delay distribution of Cheng *et al.* [126]. Delay time distributions associated with double white dwarf mergers were calculated by Cheng *et al.* [126] using binary population synthesis simulations for resultant white dwarfs of various mass ranges; we use the distribution for $1.14 - 1.24 M_\odot$ white dwarfs, which spans approximately the same mass range as our most massive bin and is the merger delay time distribution that was used in Chapter 4.

5.3. Results

Double white dwarf mergers can produce photometrically young white dwarfs with transverse velocities faster than what would be expected for single stellar evolution products of the same photometric age because the true age of the white dwarf merger products (accounting for the lifetime of the merger progenitors) is typically older than the age inferred from photometry [123]. Estimates of the merger fraction from both simulations [123, 124] and data [2, 125, 126, 164] indicate that double white dwarf merger products constitute an appreciable fraction of ultramassive white dwarfs, with estimates ranging from $\sim 15\%$ to over 50%. The population synthesis simulations of Temmink *et al.* [123] indicate that 30 – 50% of white dwarfs with mass $\geq 0.9 M_{\odot}$ form through binary mergers, with about 45% of these mergers occurring between two white dwarfs, while the population synthesis simulations of Bogomazov and Tutukov [124] indicate that over 50% of white dwarfs with mass $\geq 1.1 M_{\odot}$ are produced through double white dwarf mergers.

Empirically, based on the magnetism, kinematics, rotation, and composition of the 25 most massive white dwarfs identified by Kilic *et al.* [125] in the Montreal White Dwarf Database 100 pc sample, Kilic *et al.* [164] found a merger fraction of $56^{+9}_{-10}\%$ among ultramassive white dwarfs with mass $\gtrsim 1.3 M_{\odot}$. Theoretical results indicate that nearly half of the identified merger products would have been produced through the double white dwarf merger channel [123], and over half of the merger products identified in Kilic *et al.* [125, 164] showed kinematic and/or rotation features indicative of double white dwarf mergers in particular. Cheng *et al.* [126] estimated a double white dwarf merger fraction of about 20% for *Gaia* DR2 white dwarfs within 250 pc with mass in the range $0.8 – 1.3 M_{\odot}$ and a larger fraction of about 35% for the subset in the mass range $1.14 – 1.24 M_{\odot}$ (assuming carbon-oxygen cores to determine the mass). For our *Gaia* EDR3 white dwarf sample within 200 pc in the $1.15 – 1.25 M_{\odot}$ mass bin, it was found in Chapter 4 that the photometric cooling age distribution was well fitted when the double white dwarf merger fraction was about 40 – 50%.

The expectation for the limiting case of 100% double white dwarf merger products is shown in Fig. 5.1 as a simple check of the feasibility of a double

5.3. Results

white dwarf merger delay explanation of the fast-movers. The lower panel of Fig. 5.1 shows that a sufficiently large fraction of merger products among $1.15 - 1.25 M_{\odot}$ white dwarfs can reach the excess observed for the median estimate of the transverse velocity dispersion, which is representative of the (more slowly moving) bulk of the white dwarfs as the median estimator is robust to the outlying fast-movers. However, even in the extremum limit that 100% of the $1.15 - 1.25 M_{\odot}$ white dwarfs are the product of double white dwarf mergers, the middle panel of Fig. 5.1 shows that the excess observed for the mean estimate of the transverse velocity dispersion, which is sensitive to the outlying fast-movers, cannot be achieved. In keeping with Cheng *et al.* [43], we thus find that double white dwarf mergers (in the local Galactic disc) cannot be the sole explanation of the fast-movers.

5.3.2 Three-Dimensional Velocity Distributions

The behaviour of the AVR_s that we measure in Section 5.3.1 for the transverse velocities aligns with the behaviour of the observed fraction of fast-movers as a function of cooling age. For both the non-robust and robust estimators of transverse velocity dispersion, we see that the lightest ($0.95 - 1.05 M_{\odot}$) white dwarfs tend to become more dispersed with age, as inferred from the fraction of fast-movers and expected from the AVR_s of main-sequence stars. The $1.05 - 1.15 M_{\odot}$ white dwarfs mostly follow this trend as well, but with a small peak for cooling ages of $\sim 1.0 - 2.0$ Gyr. The most massive ($1.15 - 1.25 M_{\odot}$) white dwarfs show a prominent peak for cooling ages of $\sim 0.5 - 2.0$ Gyr. The peak in transverse velocity dispersion of ultramassive white dwarfs inferred using the non-robust estimator (mean) is much larger than that inferred using the robust estimator (median), suggesting that the fast-movers in these age bins are outliers or comprise a separate population that is much more dispersed than the main population.

To better understand the kinematic features of the population of young fast-movers, we consider the distribution of each Cartesian component of the transverse velocities in each mass and age bin. Fig. 5.2 shows quantile-quantile (Q-Q) plots comparing the sample distribution of transverse Car-

5.3. Results

sian velocity components U_{\perp} (upper row), V_{\perp} (middle row), and W_{\perp} (bottom row) to the quantile function for a normal distribution with zero mean. The quantile function of a distribution is the inverse of the cumulative distribution; for an input percentile p , it gives the value of the random variable for which there is p probability of the random variable being less than or equal to the output value. For a 1D normal distribution with mean μ and variance σ^2 , the quantile function is $\mu + \sigma \sqrt{2} \operatorname{erfc}^{-1}(2[1-p])$, where $p \in [0, 1]$ is the percentile (i.e. cumulative probability) and erfc^{-1} is the inverse complimentary error function. In Fig. 5.2, the measured value of the velocity component is used as the x -axis variable, while the y -axis variable is the value of the quantile function in units of standard deviation calculated from the weighted cumulative fraction. Written explicitly, the y -axis variable is $y_i = \sqrt{2} \operatorname{erfc}^{-1}(2[1-p_i])$ for the i th ordered data point (in order of increasing x) with $p_i = \sum_{j \leq i} w_j$ where j runs over all points for which $x_j \leq x_i$. For reference, we also show quantile functions for normal distributions with zero mean and $\sigma = 10, 20, 30$, and 40 km s^{-1} in order of decreasing slope as black lines. If the transverse velocity in a particular direction is normally distributed for a particular sample, the corresponding Q-Q plot in Fig. 5.2 should appear as a straight line. A non-zero mean will simply cause a shift of the distribution along the horizontal axis, as can be seen in the plots for V_{\perp} due to asymmetric drift.

Observations of main-sequence stars (e.g. [133]) indicate that the velocity dispersion of stars in the local thin disc increases with age. If the white dwarf velocities follow the AVRs of main-sequence stars [133], the slope of the distribution in Fig. 5.2 should get progressively less steep with increasing cooling age for a given mass bin and velocity component. For the age bins considered in Fig. 5.2, this effect should be more pronounced for younger ages due to the power-law form of the AVRs with exponents < 1 . This is the approximate behaviour seen for the lightest two bins, $0.95 - 1.05 M_{\odot}$ and $1.05 - 1.15 M_{\odot}$ (left and middle columns of Fig. 5.2). For the ultramassive $1.15 - 1.25 M_{\odot}$ white dwarfs, however, we see that the curves for the $0.5 - 1.5 \text{ Gyr}$ white dwarfs have a shallower slope and thus larger dispersion than what would be predicted by the Holmberg *et al.* [133] AVRs, particularly for

5.3. Results

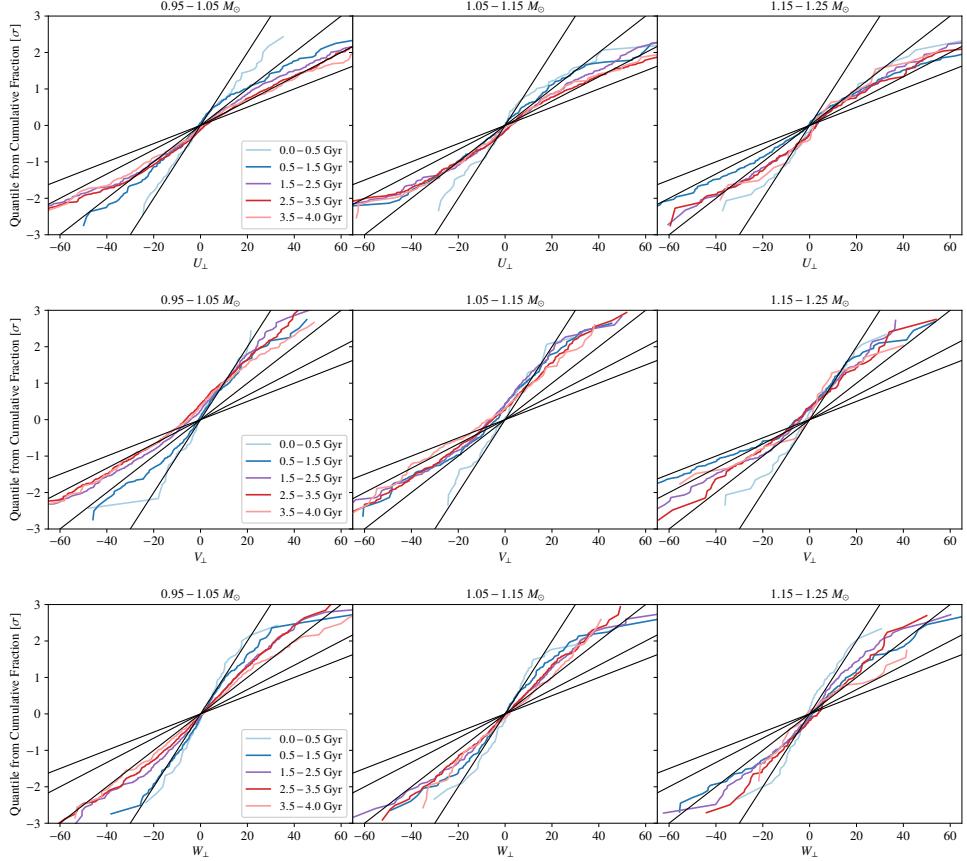


Figure 5.2: Quantile-quantile plots of transverse velocity components. The vertical axis is the quantile function in units of standard deviation σ for a normal distribution with zero mean, calculated from the empirical cumulative fraction corresponding to the velocity component shown on the horizontal axis. The black lines correspond to normal distributions with $\sigma = 10, 20, 30$, and 40 km s^{-1} in order of decreasing slope.

5.3. Results

negative U_{\perp} and V_{\perp} . The dispersions for these 0.5–1.5 Gyr white dwarfs are notably larger than the dispersions for the older 1.5–2.5 and 2.5–3.5 Gyr white dwarfs with mass in the range $1.15 - 1.25 M_{\odot}$. Furthermore, for the 0.5–1.5 Gyr age bin of $1.15 - 1.25 M_{\odot}$ white dwarfs, the white dwarfs with negative V_{\perp} are much more dispersed in V_{\perp} than in U_{\perp} , indicating that these white dwarfs likely do not originate from the local thin disc.

The AVRs of Holmberg *et al.* [133] indicate that local thin disc stars are more dispersed in U than V , and we show in Fig. 5.3 that this expectation holds for the tangential velocity components U_{\perp} and V_{\perp} as well. Since the Holmberg *et al.* [133] AVRs describe the full three-dimensional space velocities, we converted those relations to the expected AVRs of the projected velocities for the sky coordinate distribution of our $1.15 - 1.25 M_{\odot}$ sample with 0.5–2.5 Gyr photometric ages. This enables us to properly compare the results of Holmberg *et al.* [133] for main-sequence stars from the local thin disc to the dispersion values of the tangential velocity components inferred from Fig. 5.2 for our ultramassive white dwarf sample. The expected AVRs of the projected velocities are given by the square root of the diagonal elements of the expectation of the matrix $\mathbf{P}\sigma^2\mathbf{P}^T$. Fig. 5.3 shows the AVRs from Holmberg *et al.* [133] for the full three-dimensional space velocities for each of the U , V , and W components (as solid blue, orange, and green curves, respectively), compared to the computed AVRs expected from the projected velocities for our sample’s empirical sky distribution (dashed lines). We also analytically evaluated the expected AVRs of the projected velocities for an isotropic distribution (shown as the dotted curves in Fig. 5.3) and found them to be nearly the same as the AVRs determined using the empirical sky distribution of our sample.

From Fig. 5.3 we see that the expected dispersions for the projected velocities are typically notably smaller than the corresponding dispersion of the un-projected velocities, but that the relative size of the dispersion for V_{\perp} compared to U_{\perp} is similar to the relative sizes of σ_V and σ_U . Of particular note is that, throughout the age of the thin disc, the dispersion of V_{\perp} is smaller than the dispersion of U_{\perp} . For the population of anomalously fast white dwarfs (mass bin $1.15 - 1.25 M_{\odot}$ and photometric age bin 0.5–1.5 Gyr)

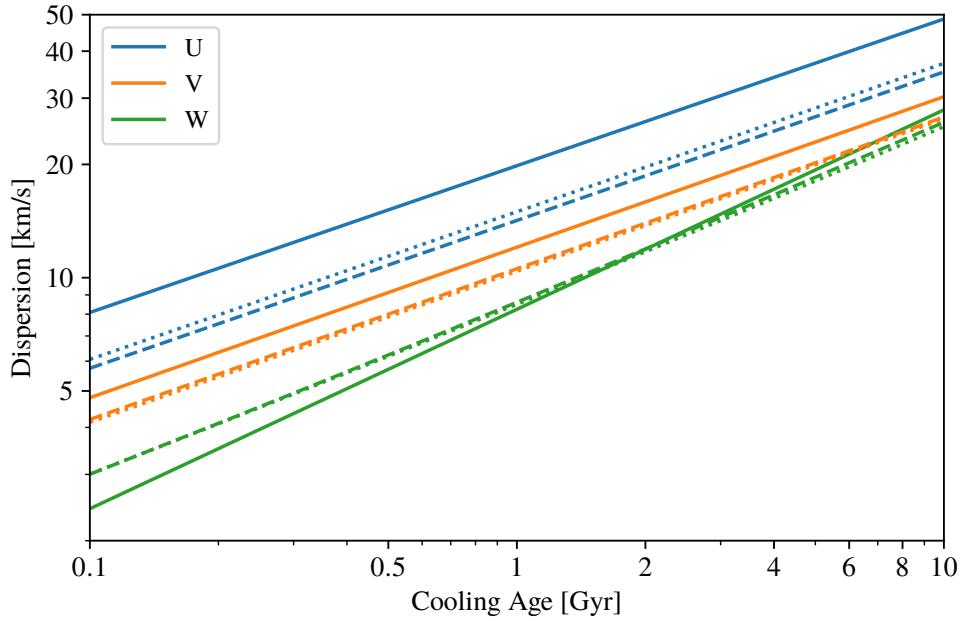


Figure 5.3: AVRs of projected and un-projected velocities. AVRs are shown for U , V , and W in blue, orange, and green, respectively. The solid curves show the AVRs for the full three-dimensional space velocities determined by Holmberg *et al.* [133] for main-sequence stars. The dotted curves show the expected AVRs for the projected velocities assuming that the un-projected velocities follow the AVRs of Holmberg *et al.* [133] and that the sky distribution is isotropic. The dashed curves also show the expected AVRs for the Holmberg *et al.* [133] projected velocities but assuming that the sky distribution is that observed for $1.15 - 1.25 M_{\odot}$ white dwarfs with $0.5 - 2.5$ Gyr cooling ages.

5.4. Discussion

with negative V_{\perp} , however, the dispersion of V_{\perp} is actually larger than the dispersion of U_{\perp} . A population in equilibrium in the disc should have a dispersion ratio for V relative to U of $\sigma_V/\sigma_U = 2/3$ [50], and the AVRs measured by Holmberg *et al.* [133] for thin disc main-sequence stars have a dispersion ratio nearly equal to this equilibrium ratio (though actually slightly smaller). The dispersion of V_{\perp} being larger than the dispersion of U_{\perp} in the anomalous population thus indicates that this population is not in equilibrium in the disc and thus does not originate from the local thin (or thick) disc.

5.4 Discussion

The presence of a population of anomalously fast-moving young, ultramassive white dwarfs on the Q branch was first discovered in *Gaia* DR2 observations by Cheng *et al.* [43], who argued this population must experience an extra cooling delay on the Q branch of $\sim 6 - 8$ Gyr that is not accounted for in current white dwarf cooling models. This proposed extra cooling delay scenario explained the large transverse velocities observed for the fast-moving population as a consequence of these white dwarfs being much older than their photometric ages indicate. The fast-movers in this scenario were assumed to originate from the local disc, where observed AVRs indicate that objects become more dispersed over time. As a check of the assumption of a local disc origin of the fast-mover population in the extra cooling delay scenario, Cheng *et al.* [43] ran a test in which they modelled the velocity distribution of the fast-movers on the Q branch as a Gaussian and determined kinematic relations expected for a disc in equilibrium, notably finding a dispersion ratio σ_V/σ_U of approximately 2/3. Like Cheng *et al.* [43], we also find a population of anomalously fast-moving young, ultramassive white dwarfs coincident with the Q branch. However, the empirical distributions of the transverse velocity components in Fig. 5.2 for *Gaia* EDR3 reveal kinematic features that are inconsistent with a local disc origin, in particular the distribution for V_{\perp} and the dispersion ratio for V_{\perp} and U_{\perp} .

5.4. Discussion

Fig. 5.2 suggests that the distribution of V_{\perp} for the anomalously fast ultramassive white dwarfs is actually too broad to be explained entirely by a cooling delay in (thin) disc white dwarfs. If white dwarfs with a true age of 10 Gyr produced through single stellar evolution in the Galactic thin disc follow the AVRs determined by Holmberg *et al.* [133], then the expected dispersion for the projected velocity component V_{\perp} in the tangent plane averaged over the sky distribution of the $1.15 - 1.25 M_{\odot}$ white dwarfs with photometric ages 0.5 – 2.5 Gyr would be approximately 26 km s^{-1} (see Fig. 5.3). From Fig. 5.2, we see that the anomalous $1.15 - 1.25 M_{\odot}$ white dwarfs with 0.5 – 1.5 Gyr photometric ages have a dispersion for V_{\perp} of $\sim 40 \text{ km s}^{-1}$ when V_{\perp} is negative (i.e. among white dwarfs that are moving in the opposite direction of Galactic rotation). The Galactic thin disc is not old enough to produce white dwarfs with such a high dispersion given the observed AVRs [133]. While stars in the older thick disc have slightly higher dispersion [138], the expected dispersion of $\sim 33 \text{ km s}^{-1}$ for V_{\perp} is still too small to explain this dispersion through a cooling delay in a thick disc population.

Furthermore, as noted in Section 5.3.2 and seen in Fig. 5.2, we find a population of highly dispersed ultramassive white dwarfs coincident with the Q branch for which the dispersion in V_{\perp} is larger than the dispersion in U_{\perp} . This is inconsistent with the dispersion ratio of a disc in equilibrium, for which the dispersion in V_{\perp} should be smaller than the dispersion in U_{\perp} with a dispersion ratio similar to that of the un-projected velocity components (see Fig. 5.3). Thus the empirical dispersion ratio for V_{\perp} and U_{\perp} suggests that these white dwarfs do not originate from the Milky Way disc.

The true origin of this population, however, is unclear. These white dwarfs may have originated from elsewhere in the Galaxy such as from the halo. The halo is older than the Galactic disc [165] and halo stars are typically more dispersed [151]. This scenario still requires a mechanism by which photometrically young white dwarfs acquire velocities faster than their photometric age would indicate. This mechanism could be an extra cooling delay as proposed by Cheng *et al.* [43] of nearly 10 Gyr, so that only white dwarfs with photometric ages of 0.5 – 1.5 Gyr are affected. Alternatively,

5.4. Discussion

these objects could originate from a dynamically distinct, hot yet young population.

Another possible explanation of the anomalous fast-movers is that they are the merger remnant of the inner binary of a hierarchical triple system whose outer companion was ejected. Observations indicate that triple systems are common, constituting about 10% of multiples with F and G dwarf primaries [166–169] and much larger fractions for higher mass B- and O-type primaries [169, 170], and population synthesis simulations indicate that stellar interactions occur in the majority of triple systems [171, 172]. The evolution of triple systems is significantly more complicated than that of isolated binaries due to the combination of stellar evolution, stellar interactions, and three-body dynamics (see e.g. Toonen *et al.* [173] for a review). Triple systems tend to be hierarchical, with an inner binary and a distant companion [174]. Such systems have a large variety of possible evolution channels [171, 172, 175] and can result in mergers and ejections [175–181]. The secular gravitational effects between the outer companion and inner binary can drive the inner binary to high eccentricity in Zeipel-Lidov-Kozai oscillations [182–186], which can enhance mergers of compact objects [172, 187–200]. For example, dynamical effects of the outer companion can enhance the merger rate of inner compact double white dwarf binaries relative to isolated binary systems [188]. Destabilized stellar triple systems in the field can eject a component with runaway speeds of tens of km s^{-1} [175].

Unstable triple systems in star clusters can also result in the ejection of the tertiary perturber and escape from the cluster [176]. From their N -body simulations of open clusters, Mardling and Aarseth [176] describe a particular event in which an initially stable triple system became unstable and the tertiary companion was ejected; the resulting binary and single star both escaped the cluster with respective terminal velocities of 20.6 and 65.7 km s^{-1} . Star clusters are known to harbour exotica such as blue stragglers, which can be formed through the dynamical evolution of stellar triple systems [201] and have been found in both open clusters [202–204] and globular clusters [205, 206]. Furthermore, cluster members can display different kinematic

5.5. Conclusions

relations than what is expected for a disc in equilibrium, and many nearby clusters are moving at speeds of tens of km s^{-1} relative to the local standard of rest [100, 207–209]. The kinematically anomalous population we observe could thus conceivably originate from old open clusters or globular clusters that are close to the Galactic disc. While open clusters are typically younger than the members of the anomalous population, some open clusters are old enough to be a source of these white dwarfs, and those old open clusters are also typically rich [100, 207, 208]. As open clusters are only loosely gravitationally bound, stars can also readily escape from them; for example, observations indicate that the Pleiades has lost $\sim 20\%$ of its mass over the past 100 Myr [122]. Though globular clusters are more tightly bound than open clusters and thus less prone to the escape of cluster members, they are also typically much more populated [100] and only a small fraction of escapees is needed to explain the anomalous population that we observe in this work.

In order to explain the clustering in photometric age in this scenario, the fast-moving population must be produced through a singular event rather than some continuous process. This could be cluster evaporation in the case of an open cluster origin, or it could be some transient event that disrupted a globular cluster. This transient event may relate to the burst of star formation in the local Galactic disc that peaked around 2 – 3 Gyr ago and continued until ~ 1 Gyr ago [106]. A travel time of around 1 Gyr from the location of origin after the disruption of the cluster would then produce the observed clustering in age of the fast-moving population.

5.5 Conclusions

As a follow-up to our work in Chapter 4 on the distribution of photometric cooling ages of massive white dwarfs in *Gaia* EDR3, we analysed the kinematics of the transverse motion of $0.95 - 1.25 M_{\odot}$ *Gaia* EDR3 white dwarfs. We find that the population of anomalously fast ultramassive white dwarfs on the Q branch reported by Cheng *et al.* [43] for *Gaia* DR2 is still present in *Gaia* EDR3. These white dwarfs appear photometrically young according

5.5. Conclusions

to single stellar evolution models, but are moving faster than expected for stars of that age according to AVR_s observed for the local thin disc. Using white dwarf cooling models for which the photometric age distributions are consistent with the expectation from the star formation history (see Chapter 4) to infer the masses and ages of our sample of white dwarfs and sorting our full sample into mass and age bins, we find that this population of fast-movers is concentrated to masses of $1.15 - 1.25 M_{\odot}$ and photometric cooling ages of $0.5 - 1.5$ Gyr.

Our analysis of the distributions of the individual components of the transverse velocity reveals that among the white dwarfs in the mass range $1.15 - 1.25 M_{\odot}$ and age range $0.5 - 1.5$ Gyr, there is a population of white dwarfs lagging the local Galactic rotation that is too dispersed in V_{\perp} to be explained solely by a cooling delay in white dwarfs born in the local thin disc. According to the AVR_s determined from observations of main-sequence stars in the local thin disc [133], it would take longer than the age of the thin disc (10 Gyr) for disc heating to produce a dispersion large enough to explain this anomalous population of white dwarfs. Furthermore, we find this population to be more dispersed in V_{\perp} than in U_{\perp} , which suggests that this population does not originate from the local disc. Some potential explanations of this population include a halo origin, in conjunction with an extra cooling delay or from a dynamically distinct population, or that they are produced through triple system dynamics such as the merger of the inner binary of a hierarchical triple system whose tertiary companion was ejected resulting in a large velocity for the binary that ultimately merges to form the white dwarf. However, the precise origin of this population of young ultramassive white dwarfs is an open question for future work.

Chapter 6

The cooling of old white dwarfs in 47 Tucanae

6.1 Introduction

Globular clusters are useful environments for studying white dwarf cooling because they provide populations of stars with many well controlled parameters, such as distance, interstellar reddening, age, and metallicity (e.g. [210–212]). Though metallicity is not directly a concern for observations of white dwarfs, since element sedimentation causes metals to sink below the surface envelope of the white dwarf by the early stages of white dwarf cooling, it affects the age and composition of the white dwarfs through earlier stages of stellar evolution. For a coeval population of stars with the same initial metallicity, white dwarfs of a similar mass are formed through single stellar evolution at an approximately constant rate.

Unlike open clusters, for which the parameters mentioned above are also well controlled, globular clusters are very old and typically much more rich (i.e. well-populated). These additional properties of globular clusters enable white dwarf cooling to be studied to much later times and with the greater statistical power of a larger sample size compared to what could be achieved with open clusters. The globular cluster 47 Tuc, in particular, has an especially rich white dwarf population and is known to be very old, with an age of ~ 10 Gyr [213]. Furthermore, the distance to 47 Tuc has been well determined [214].

47 Tuc is one of the most widely studied globular clusters, and white dwarfs in 47 Tuc have been used to study a variety of aspects of white dwarf

6.1. Introduction

cooling. For example, Obertas *et al.* [30] studied the onset of convective coupling and core crystallisation at late times in the white dwarf cooling process by comparing white dwarf evolution models to deep HST observations of white dwarfs in 47 Tuc. As another example, HST observations of younger white dwarfs in 47 Tuc were used by Goldsbury *et al.* [215] to study white dwarf cooling by the emission of neutrinos and to constrain neutrino physics. In both of these cases, the typical envelope thickness of white dwarfs in 47 Tuc was an important ancillary parameter in studying the relevant aspects of white dwarf cooling due to the relation between the envelope thickness and the cooling rate.

The envelope thickness is particularly important for the cooling rate at late cooling times when the envelope becomes convectively coupled to the core of the white dwarf. Convective coupling of the envelope to the core occurs when the convective layer that eventually develops in the outer layers of a white dwarf [67] breaks through to the degenerate interior, which enables energy to be transported from the core to the surface more quickly than if the envelope had remained radiative (see e.g. [68] for a review). This results in the white dwarf initially appearing more luminous at the onset of convective coupling than it otherwise would have, which manifests as a bump in the cooling curve showing the evolution of luminosity as a function of cooling time. Tassoul *et al.* [216] argued that the size of this bump depends sensitively on the thickness of the hydrogen envelope.

In this chapter, we analyse the cooling of white dwarfs in 47 Tuc to late cooling times where convective coupling becomes important for the cooling rate. At such late cooling times, the white dwarfs have cores that have long since finished transitioning from a gaseous to liquid state of matter and, as shown by Obertas *et al.* [30] in the case of 47 Tuc, are even in the process of crystallising from a liquid to solid state. Though from a theoretical standpoint a white dwarf core can only reasonably be approximated as an ideal gas at very early cooling times, the usual implementation of element diffusion in stellar evolution software like Modules for Experiments in Stellar Astrophysics (MESA; [217–222]) implicitly approximates the ions as an ideal gas in order to make the diffusion equations tractable. We per-

form stellar evolution simulations using MESA to create white dwarf cooling models for different envelope thicknesses, white dwarf masses, and diffusion scenarios. Three different treatments of element diffusion are considered: i) the standard MESA treatment where the ions are approximated as an ideal gas, ii) a custom modified treatment where diffusion is suppressed compared to the standard treatment to account for non-ideal gas effects, and iii) the case of no diffusion. We compare these cooling models to the deep HST data considered by Obertas *et al.* [30] using an unbinned likelihood analysis procedure similar to that of Goldsbury *et al.* [215]. This work provides both the best-fitting values of some parameters that are important for modelling white dwarf cooling in 47 Tuc, in particular the typical envelope thickness and mass of white dwarfs, and a test of the standard MESA treatment of diffusion to cooling times late enough that the core begins to crystallise.

6.2 Data

Our data consists of archival *Hubble Space Telescope* (HST) deep observations of the outer field of 47 Tuc imaged by the Advanced Camera for Surveys (ACS) using the Wide Field Camera (WFC) instrument, as described by Kalirai *et al.* [223]. These data were collected over 121 orbits during the time period extending from January 2010 to October 2010 as part of the HST Cycle 17 proposal GO-11677 (PI: H. Richer). The ACS/WFC observations were done using the broadband filters F606W and F814W. The deep exposures had a total integrated exposure time of 163.7 ks across 117 exposures in F606W and 172.8 ks across 125 exposures in F814W. These observations were centred at sky coordinates of $\alpha = 00^{\text{h}} 22^{\text{m}} 39^{\text{s}}$ and $\delta = -72^{\circ} 04' 04''$ in the international celestial reference frame (ICRS) at the reference epoch J2000, where α is the right ascension and δ is the declination. This corresponds to a distance of about 6.7 arcminutes (8.8 pc) from the cluster centre, which is located at $\alpha = 00^{\text{h}} 24^{\text{m}} 05.71^{\text{s}}$ and $\delta = -72^{\circ} 04' 52.7''$ [224].

For each filter, images from the various exposures were combined into a single final, stacked image. Photometric, astrometric, and morphological measurements were then performed on the final, stacked images for the two

6.2. Data

filters using iterative point-spread function (PSF) fitting techniques. Morphological information from the PSF fitting is stored in the **SHARP** diagnostic parameter, which provides a measure of how much broader the source’s profile is compared to the PSF profile it was fitted to. The **SHARP** parameter provides a way to distinguish stars from other contaminant sources such as galaxies and cosmic rays.

The PSF fitting produced a catalogue of sources that contains the F606W and F814W magnitudes, position, chi goodness-of-fit statistic, and **SHARP** statistic determined for each source, with the magnitudes reported in the Vega magnitude system. The full image processing and PSF fitting procedures are described in further detail in Kalirai *et al.* [223], and the resultant catalogue is publicly available through the Mikulski Archive for Space Telescopes (MAST) as a High-Level Science Product (HLSP)¹⁶. In addition to this catalogue, the results of the artificial stars tests documented in Kalirai *et al.* [223, 225] are also publicly available as part of the same HLSP collection¹⁶. These artificial stars tests are used to characterise the photometric uncertainties and completeness of the data in the final catalogue, and are discussed in detail in Section 6.3.

Photometric observations of 47 Tuc white dwarfs are contaminated at the faint end of the white dwarf cooling sequence by the Small Magellanic Cloud (SMC). Though the main body of the SMC lies more than 2° away from 47 Tuc¹⁷, a diffuse population of SMC stars persists out to very large radii and is present in the background of our 47 Tuc observations. This background SMC population overlaps with the faint end of the 47 Tuc white dwarf cooling sequence in the CMD of our data, which is the region of the cooling sequence we are most interested in. Fortunately, the SMC is moving with respect to 47 Tuc fast enough that the two populations can be mostly separated in proper motion space, and thus most of the SMC contaminants can be removed from our data. The reduction procedures to determine the

¹⁶The final stacked images, source catalogue, and artificial stars data are available at <https://archive.stsci.edu/prepds/deep47tuc/>.

¹⁷The SMC is located at the sky coordinates $\alpha = 00^{\text{h}} 52^{\text{m}} 45^{\text{s}}$, $\delta = -72^\circ 49' 43''$ [226]. This is an angular distance of $\sim 2.28^\circ$ from the centre of 47 Tuc and $\sim 2.39^\circ$ from the centre of the ACS/WFC observing field.

6.3. Artificial Stars Tests

proper motions accompanying the ACS/WFC photometric observations are described in detail in Richer *et al.* [227]. Note that we are interested in objects much fainter than those considered in Richer *et al.* [227] and that the uncertainty in the proper motions increases with magnitude. Our proper motion data thus has larger uncertainties and appears more dispersed overall than that of Richer *et al.* [227]. Since we are interested in the white dwarf cooling sequence down to very faint magnitudes, it is important to include the proper motions for the fainter white dwarfs, even though it results in a sample with larger proper motion uncertainties.

We identify white dwarfs associated with 47 Tuc in the HST data by making cuts in **SHARP**, proper motion, and the CMD. The cut in **SHARP** enables us to clean the data by removing contaminants such as cosmic rays and galaxies whose photometric profiles do not match the expected PSF profile for stars. The cut in proper motion is used to select sources likely to be members of 47 Tuc, further cleaning the data by removing most of the SMC and field stars. For this data cleaning, we perform a **SHARP** cut that selects objects with $|\text{SHARP}| < 0.5$ and a proper motion cut that selects objects with a total proper motion $< 2.5 \text{ mas yr}^{-1}$ relative to the mean proper motion of 47 Tuc. After cleaning the data using these cuts in **SHARP** and proper motion, we then perform a cut in the CMD to select our white dwarf sample, and the boundaries of this cut in the CMD define the data space for the unbinned likelihood analysis. The data cleaning procedure, including the choice of which cuts to make, is further explained in detail in Section 6.4. The unbinned likelihood analysis, including the CMD data space selection, is discussed in Section 6.7.

6.3 Artificial Stars Tests

The artificial stars tests were performed by adding artificial sources to the final stacked images and running the new images through the same PSF fitting procedure as was used for the real data. Each artificial source was given a unique ID in order to track whether the artificial source was detected by the PSF fitting procedure and, if the source was detected, to compare

6.3. Artificial Stars Tests

the output values of its magnitude and position to the known true value. A detailed description of the artificial stars tests and the corresponding artificial sources catalogue is described in detail in Kalirai *et al.* [223].

The catalogue of artificial sources produced by these artificial stars tests was used to construct a photometric error distribution function in a procedure similar to what is described in Goldsbury *et al.* [215]. Let $F606W_{\text{in}}$ and $F814W_{\text{in}}$ denote the input magnitudes of a source, and let $F606W_{\text{out}}$ and $F814W_{\text{out}}$ denote the output magnitudes determined by the PSF fitting procedure. The photometric error distribution function

$$E = E(\Delta F606W, \Delta F814W; F606W_{\text{in}}, F814W_{\text{in}}) \quad (6.1)$$

gives the joint probability density, normalised to the number of input stars, of the magnitude errors $\Delta F606W$ and $\Delta F814W$ as a function parameterised by the input magnitudes, where the errors are quantified as the differences between the output and input values of the magnitudes, $\Delta F606W = F606W_{\text{out}} - F606W_{\text{in}}$ and $\Delta F814W = F814W_{\text{out}} - F814W_{\text{in}}$. The normalisation of the error distribution function as a function of the input magnitudes is simply the completeness,

$$C(F606W_{\text{in}}, F814W_{\text{in}}) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} E d(\Delta F606W) d(\Delta F814W), \quad (6.2)$$

which in general is less than unity because not all of the sources that are actually present are recovered by the photometric reduction procedure. The completeness quantifies the probability of detecting a source and can in principle take values in the range of 0 to 1, though in practice the observations become unusable if the completeness becomes too poor.

In general, the error distribution function and completeness also depend on the position from the centre of the cluster. However, for the data considered in this work, the position dependence is negligible. Furthermore, it should be noted that since all of the cooling models that will be considered in our analysis lie along approximately the same curve in colour-magnitude (and likewise magnitude-magnitude) space before accounting for photomet-

ric errors, the photometric error distribution only needs to be constructed at the combination of ($F606W_{in}$, $F814W_{in}$) values that lie along this sequence.

6.4 Data Cleaning Procedures

6.4.1 Overview

To clean our deep HST ACS/WFC data, we want to remove sources in the catalogue that (1) are not stars or (2) are not members of 47 Tuc. In the latter case, we are particularly concerned with removing SMC stars that contaminate the faint end of the 47 Tuc white dwarf cooling sequence in the CMD. A cut in **SHARP** allows us to remove objects that are not stars, while a cut in proper motion allows us to remove objects that are unlikely to be 47 Tuc cluster members. We calibrate our data cleaning procedure using 47 Tuc main-sequence stars to choose what cuts to make in **SHARP** and proper motion and to quantify the effect of these cuts on the completeness of our white dwarf sample. We furthermore quantify the number of SMC stars expected to survive the cleaning procedure and contaminate our white dwarf sample. Any changes to the completeness arising from the cleaning procedure must be accounted for and applied to the error distribution function from the artificial stars tests.

6.4.2 SHARP

If an object identified by the DAOPHOT II program ALLSTAR is a star, then it should have a **SHARP** value near zero. Values of **SHARP** much larger than zero indicate the object is probably a galaxy or unrecognized double, while objects with **SHARP** much less than zero are probably cosmic rays or image defects such as bad pixels [228]. A cut in the **SHARP** parameter can thus be used to remove objects that are not stars. Galaxies in particular are a common contaminant of the white dwarf cooling sequence at very faint magnitudes [223].

To determine what a reasonable range of **SHARP** values is for stars in our sample, we analyze the distribution of **SHARP** values for 47 Tuc main-

sequence stars as a function of magnitude. The 47 Tuc white dwarfs should have a distribution in **SHARP** similar to that of the 47 Tuc main-sequence stars (at comparable magnitudes), so we use our analysis of the latter to choose the threshold value for our **SHARP** cut. The aim is to make a cut in **SHARP** that is generous enough to not reduce the completeness of the white dwarf data but strict enough to remove as many objects that are not stars as possible.

The 47 Tuc main-sequence stars are selected using a cut in the CMD. The CMD boundary used to select the 47 Tuc main-sequence stars is shown in Fig. 6.1 by the green lines (labelled “47 Tuc MS”). Figure 6.1 also shows the boundaries used to define CMD-selected populations of 47 Tuc white dwarfs (“47 Tuc WD”, blue lines) and SMC stars (“SMC”, orange lines), which are used later in the analysis in Section 6.4.3. The CMD boundary for the 47 Tuc white dwarfs is furthermore the same boundary ultimately used to define the data space in the unbinned likelihood analysis. In this section, these boundaries of the 47 Tuc white dwarfs and the SMC in the CMD simply serve as visual references to note the locations of these populations. From left to right across Fig. 6.1, the CMD boundaries correspond to 47 Tuc white dwarfs, SMC stars, and 47 Tuc main-sequence stars. The CMD boundaries of these three populations are the same in both sub-figures of Fig. 6.1, though the full span of each bounding region can only be seen in Fig. 6.1a. All three boundaries span the same range of F606W magnitude values along the y-axis, from 22 to 29, so the calibration analysis in both this section and Section 6.4.3 can be done as a function of F606W.

We construct the empirical **SHARP** distributions for the CMD-selected main-sequence stars both before and after performing a proper motion cut to select objects within 2.5 mas yr^{-1} of the mean proper motion of 47 Tuc¹⁸. This is the same proper motion cut that we use in the final cleaning procedure for our white dwarf data. The **SHARP** distributions for the proper-

¹⁸In contrast, the proper motion cut to select SMC stars used to calibrate the proper motion cleaning procedure (discussed in Section 6.4.3) selects objects within 0.75 mas yr^{-1} of the mean proper motion of the SMC, a much tighter cut in proper motion than what is used for 47 Tuc.

6.4. Data Cleaning Procedures

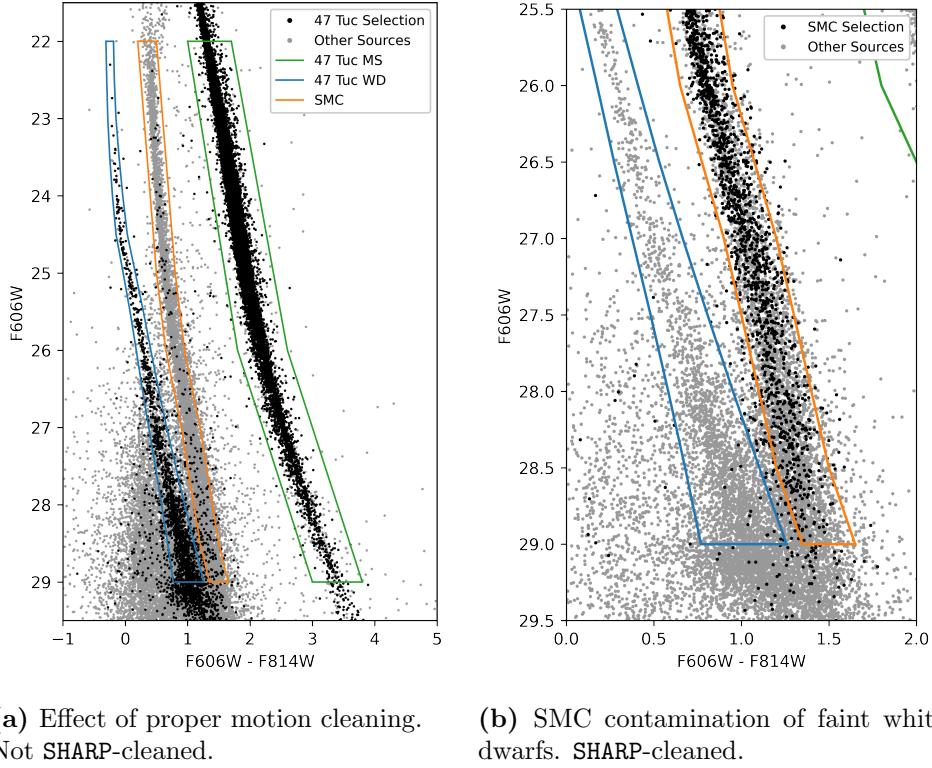


Figure 6.1: CMDs showing population boundaries and effects of cleaning procedures. The boundaries defining the CMD-selected populations are shown as solid lines. From left to right, these populations are 47 Tuc white dwarfs (blue), SMC stars (orange), and 47 Tuc main-sequence stars (green). The left panel shows the effect of proper motion cleaning, while the right panel shows the effect of SHARP cleaning and the SMC contamination at the faint end of the 47 Tuc white dwarf data space. The data point colours indicate different proper motion cuts for the left and right panels. *Left:* Data without SHARP cleaning. Sources selected by the 47 Tuc proper motion cut are shown in black, while rejected sources are shown in grey. *Right:* SHARP-cleaned data with focus on faint white dwarfs. Sources selected by the SMC proper motion cut are shown in black, while other sources are shown in grey. The SMC proper motion cut is used for calibration purposes in Section 6.4.3.2; its purpose is to select a very pure sample of SMC stars rather than as many SMC members as possible.

6.4. Data Cleaning Procedures

motion-cleaned main-sequence stars are of most interest to us because the proper motion cleaning produces a more pure sample of 47 Tuc members. However, comparing the **SHARP** distributions both before and after proper motion cleaning is also useful as it gives us information about the typical **SHARP** values of the outliers removed by the proper motion cut, especially those sources that are likely not stars. The reduction in completeness caused by the proper motion cut is analysed in Section 6.4.3. While that is an important consideration for our analysis of the white dwarf cooling, a reduction in completeness of the main-sequence sample due to this proper motion cut is not a concern for our analysis of the **SHARP** distribution, as it will only affect the amplitude, not the shape, of the distribution.

The effect of the proper motion cleaning on the CMD of the data is shown in Fig. 6.1a, where it can be seen that most of the SMC stars are removed by the cut in proper motion. While some SMC stars survive the proper motion cut, the SMC is located far enough away from the 47 Tuc main-sequence in the CMD that the SMC stars do not contaminate the main-sequence sample selected using the CMD cut. In Fig. 6.1a, the black points show the sources that were selected as likely 47 Tuc members by the proper motion cut, while the grey points show sources that were rejected by this cut. Note that the black and grey colour-coding of the data points in Fig. 6.1b has a different meaning than in Fig. 6.1a. In Fig. 6.1b, the colour-coding of the data points indicates which objects are selected (black) or rejected (grey) by a proper motion cut to select likely SMC members, which is used in Section 6.4.3.2 to analyse the SMC contamination in the 47 Tuc white dwarf data space and is described in detail in that section. It should also be noted that both the CMD and proper motion selections of SMC stars are only used for the purpose of calibrating the proper motion data cleaning procedure in Section 6.4.3 and do not need to be complete for this purpose, so these selections prioritize the purity of the SMC sample over the completeness of the sample. This results in many SMC stars being excluded from the SMC selections, particularly for the proper motion selection of SMC stars shown in Fig. 6.1b. These SMC selections are not relevant for the **SHARP** cleaning procedure discussed in the current section, so the discussion here is kept

6.4. Data Cleaning Procedures

brief, but more details can be found in Section 6.4.3.2.

The empirical number distributions of the **SHARP** parameter for the 47 Tuc main-sequence stars sub-divided into various F606W magnitude bins are shown in Fig. 6.2. The distributions before proper motion cleaning are shown in Fig. 6.2a (left column), while the distributions after proper motion cleaning are shown in Fig. 6.2b (right column). Each row corresponds to a different magnitude bin. From top to bottom, the F606W magnitude bins shown in Fig. 6.2 are 22 – 25, 25 – 26.5, 26.5 – 28, and 28 – 29. The ranges for these bins were chosen by first constructing the **SHARP** distributions for evenly spaced magnitude bins of 0.5 width and then grouping together adjacent bins for which the morphology of the distributions was similar. This grouping was done to facilitate visualisation. All of the **SHARP** distributions, plotted as histograms in Fig. 6.2, were constructed using the same **SHARP** bin width of 0.05. For each magnitude bin, we also calculated the following sample statistics: mean, median, standard deviation, and skewness. These sample statistics, along with the total number N of sources in the bin, are reported in Fig. 6.2 and summarised in Table 6.1.

Written explicitly, the sample mean for a particular magnitude bin is

$$\overline{\text{SHARP}} = \frac{1}{N} \sum_{i=1}^N \text{SHARP}_i, \quad (6.3)$$

where i is an index that labels the sources in the sample and SHARP_i is the value of **SHARP** for a particular source. The sample median is denoted as $\widetilde{\text{SHARP}}$, with a tilde instead of an overline.

The sample standard deviation of the **SHARP** values in each magnitude bin is

$$\sigma_{\text{SHARP}} = \sqrt{\frac{1}{N} \sum_{i=1}^N (\text{SHARP}_i - \overline{\text{SHARP}})^2}, \quad (6.4)$$

which is the square root of the sample variance. If the **SHARP** values are normally distributed, then Eq. (6.4) corresponds to the maximum likelihood estimate of the standard deviation parameter of the underlying distribution.

For the sample skewness statistic, we use the Fisher-Pearson coefficient

6.4. Data Cleaning Procedures

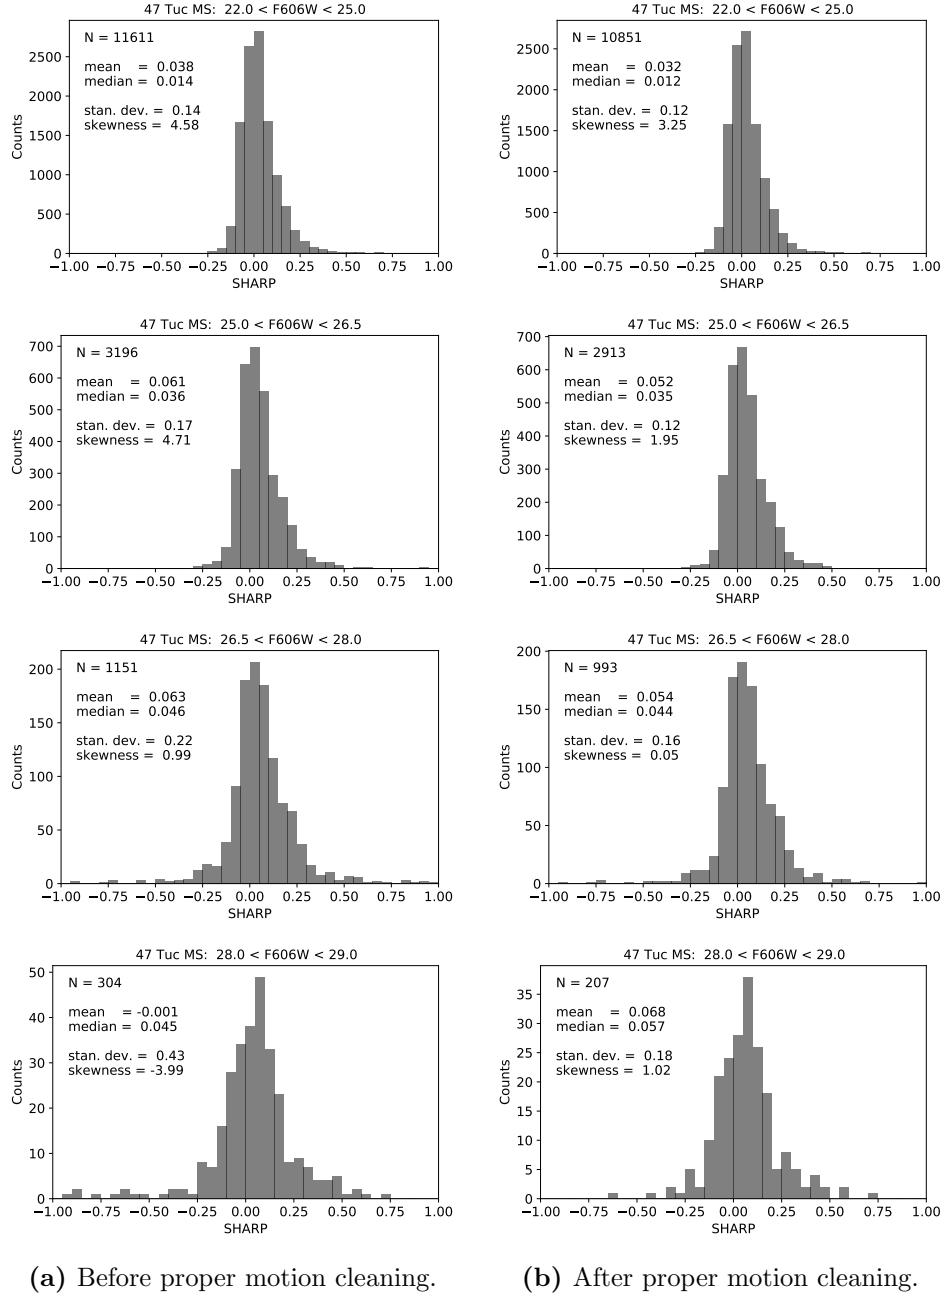


Figure 6.2: Distribution of SHARP for 47 Tuc main-sequence stars by F606W magnitude bin.

6.4. Data Cleaning Procedures

	F606W	N	$\overline{\text{SHARP}}$	$\widetilde{\text{SHARP}}$	σ_{SHARP}	g_1
Before proper motion cleaning	22.0 – 25.0	11611	0.038	0.014	0.14	4.58
	25.0 – 26.5	3196	0.061	0.036	0.17	4.71
	26.5 – 28.0	1151	0.063	0.046	0.22	0.99
	28.0 – 29.0	304	–0.001	0.045	0.43	–3.99
After proper motion cleaning	22.0 – 25.0	10851	0.032	0.012	0.12	3.25
	25.0 – 26.5	2913	0.052	0.035	0.12	1.95
	26.5 – 28.0	993	0.054	0.044	0.16	0.05
	28.0 – 29.0	207	0.068	0.057	0.18	1.02

Table 6.1: Statistics for the SHARP distributions of CMD-selected 47 Tuc main-sequence stars. For each F606W magnitude bin, the reported quantities are the number of sources (N), the mean ($\overline{\text{SHARP}}$), the median ($\widetilde{\text{SHARP}}$), the standard deviation (σ_{SHARP}), and the skewness (g_1).

of skewness

$$g_1 = \frac{m_3}{m_2^{3/2}}, \quad (6.5)$$

which is calculated from the biased sample second and third central moments, respectively m_2 and m_3 , where

$$m_n = \frac{1}{N} \sum_{i=1}^N (\text{SHARP}_i - \overline{\text{SHARP}})^n \quad (6.6)$$

is the biased sample n th central moment. The sample sizes are large enough that using the adjusted Fisher-Pearson standardized moment coefficient, $G_1 = g_1 \sqrt{N(N-1)}/(N-2)$, instead of g_1 to correct for bias makes negligible difference to the results. To the level of significance reported, the only difference is in the faintest bin, where $G_1 = -4.01$ (instead of $g_1 = -3.99$) before proper motion cleaning and $G_1 = 1.03$ (instead of $g_1 = 1.02$) after proper motion cleaning.

The mean and median are both measures of central tendency, with the median being more robust to outliers. Both the sample mean and sample median are close to zero for all magnitude bins, as expected. This applies for the SHARP distributions both before and after proper motion cleaning. For

6.4. Data Cleaning Procedures

the standard deviation, which is a measure of the spread of **SHARP** values, we see that proper motion cleaning makes more of a difference. Before proper motion cleaning, σ_{SHARP} increases more dramatically with magnitude; while after proper motion cleaning, the σ_{SHARP} values remain more similar even down to the faintest magnitude bin. σ_{SHARP} is also overall smaller after proper motion cleaning, though the value is similar before and after proper motion cleaning for the brightest stars. After proper motion cleaning, we see that a **SHARP** cut of $|\text{SHARP}| < 0.5$ corresponds to a cut of $\gtrsim 3\sigma_{\text{SHARP}}$ for all magnitude bins.

Even before proper motion cleaning, we find that σ_{SHARP} for the two magnitude bins containing the brightest sources ($22.0 < \text{F606W} < 25.0$ and $25.0 < \text{F606W} < 26.5$) are small enough that $3\sigma_{\text{SHARP}} \lesssim 0.5$. The two magnitude bins containing the faintest sources ($26.5 < \text{F606W} < 28.0$ and $28.0 < \text{F606W} < 29.0$) have larger estimated values of the standard deviation, but they also have much larger fractions of outliers that are likely not actually stars (particularly the faintest bin, $28.0 < \text{F606W} < 29.0$, which has the largest estimated standard deviation). These outliers can be seen in the tails of the distributions in Fig. 6.2a, and comparison of those distributions with the distributions for the same magnitude bins in Fig. 6.2b shows that most of these outliers are removed by the proper motion cut. The sample standard deviation is not very robust to outliers, so for these bins in particular, and before proper motion cleaning overall, the outliers lead to an overestimate of the standard deviation of the underlying **SHARP** distribution for sources that are actually single stars. The values of σ_{SHARP} after proper motion cleaning should be a better estimate, and we indeed find these values to be smaller than the values of σ_{SHARP} before proper motion cleaning.

While the **SHARP** values are approximately normally-distributed, it can be seen from Fig. 6.2 that the **SHARP** distributions have some asymmetry. The skewness statistic is a measure of this asymmetry. Before proper motion cleaning, the **SHARP** distributions for the brightest two magnitude bins have a similar value of skewness, both being positively skewed with a longer tail for positive **SHARP** values. This positive skewness may be in part due to the presence of unresolved binaries, which would have large positive **SHARP**

values [228] and would not be removed by proper motion cleaning. The negative skewness for the faintest bin ($28.0 < F606W < 29.0$) before proper motion cleaning, on the other hand, can be attributed to the long tail of likely non-stellar outliers with large negative **SHARP** values. After proper motion cleaning, we see that the skewness in the faintest bin becomes positive, like the skewness in the other magnitude bins, and the absolute value of the skewness decreases in all of the bins.

For all magnitudes, we find that most of the objects are contained within the range $-0.5 < \text{SHARP} < 0.5$. The small number of objects outside of this **SHARP** range are outliers of the **SHARP** distributions that are unlikely to be stars. We thus choose a cut in the **SHARP** parameter of $|\text{SHARP}| < 0.5$ for the cleaning procedure of our main white dwarf data. The result of this **SHARP** cleaning is shown in Fig. 6.1b, which focuses on faint magnitudes where the SMC sequence begins to overlap with the 47 Tuc white dwarf cooling sequence. The improvement in data quality achieved by **SHARP** cleaning can be seen by comparing Fig. 6.1b to Fig. 6.1a, the latter of which shows the data before any **SHARP** cleaning. The black and grey colour-coding of the data points in the two sub-figures sorts the data based on different proper motion cuts in each sub-figure, but these colours do not indicate anything about **SHARP** cleaning in either case. Neither the proper-motion-selected nor proper-motion-rejected data points shown in Fig. 6.1a have been **SHARP**-cleaned, while both the proper-motion-selected and proper-motion-rejected data points shown in Fig. 6.1b have been **SHARP**-cleaned so that only points with $|\text{SHARP}| < 0.5$ are shown in Fig. 6.1b.

6.4.3 Proper Motion

A cut in proper motion allows us to remove most of the SMC stars (and other contaminants like galaxies) from our white dwarf sample. This is particularly useful for cleaning our white dwarf sample at the faint end of the cooling sequence where it intersects with the SMC sequence in the CMD (see Fig. 6.1). However, the 47 Tuc and SMC populations overlap in the tails of their proper motion distributions, so the SMC contaminants

cannot be completely removed with a proper motion cut, at least not for the faintest stars of interest to us where the larger proper motion errors lead to more diffuse proper motion distributions for both the 47 Tuc and SMC populations. Furthermore, a cut in proper motion reduces the completeness of our sample. We quantify both of these effects in this section.

We first quantify the residual SMC contamination using both CMD-selected and proper-motion-selected SMC stars. Then we quantify the reduction in completeness using CMD-selected 47 Tuc main-sequence stars. For both of these procedures, we first clean our data by applying the **SHARP** parameter cut determined in Section 6.4.2, i.e. $|\text{SHARP}| < 0.5$, to the initial catalogue from the deep HST ACS/WFC observations. Based on the analysis of Section 6.4.2, this cut does not reduce the completeness of our 47 Tuc sample, but it removes non-stellar contaminants and facilitates identification of the CMD populations, particularly at the faint end of the 47 Tuc white dwarf and SMC sequences.

6.4.3.1 Proper Motion Distribution Model

To inform our choice of proper motion cuts in our analysis of the proper motion cleaning procedure, we want to know the mean proper motion of the SMC relative to the mean proper motion of 47 Tuc. We also want to know the spread of proper motion values for each of these two populations, which is quantified by the standard deviation for a population with normally-distributed proper motions. Let μ_α and μ_δ be the components of the tangent plane projection of the proper motion vector, where μ_α is the component in the direction of increasing right ascension and μ_δ is the component in the direction of increasing declination. While the mean proper motion of 47 Tuc has already been subtracted from our proper motion data, this mean proper motion is still included in the model discussed in this section to make the dependence on the mean motion of 47 Tuc explicit.

We model the distribution of proper motions for our data as a three-component Gaussian mixture model. In this model, the joint probability density function of μ_α and μ_δ is taken to be the linear superposition of three

6.4. Data Cleaning Procedures

bivariate normal distributions, each with its own mean $(\bar{\mu}_{\alpha,i}, \bar{\mu}_{\delta,i})$, standard deviation σ_i , and amplitude A_i , where i is an index labelling the constituent distributions. One of these Gaussian components accounts for the 47 Tuc population ($i = 1$), another one accounts for the SMC population ($i = 2$), and the final one accounts for outliers and background contaminants like field stars ($i = 3$). Written explicitly, the joint probability density function of μ_α and μ_δ is

$$f_{\mu_\alpha, \mu_\delta}(\mu_\alpha, \mu_\delta; \theta) = \sum_{i=1}^3 A_i f_{\mu_\alpha, \mu_\delta, i}(\mu_\alpha, \mu_\delta; \bar{\mu}_{\alpha,i}, \bar{\mu}_{\delta,i}, \sigma_i), \quad (6.7)$$

where θ denotes the full set of parameters that characterise the distribution and the probability density distribution of a single population labelled with index i is

$$\begin{aligned} f_{\mu_\alpha, \mu_\delta, i}(\mu_\alpha, \mu_\delta; \bar{\mu}_{\alpha,i}, \bar{\mu}_{\delta,i}, \sigma_i) \\ = \frac{1}{2\pi\sigma_i} \exp \left[-\frac{(\mu_\alpha - \bar{\mu}_{\alpha,i})^2}{2\sigma_i^2} - \frac{(\mu_\delta - \bar{\mu}_{\delta,i})^2}{2\sigma_i^2} \right]. \end{aligned} \quad (6.8)$$

It has been assumed that the two proper motion components μ_α and μ_δ are uncorrelated and have the same standard deviation for a given population. In general, these assumptions need not be true and could be relaxed in the model, but the current model given by Eq. (6.7) is sufficient for our purposes. The model is also further simplified by eliminating the dependence of $f_{\mu_\alpha, \mu_\delta}$ on a few parameters in Eq. (6.7) as follows.

The proper motion data are already given relative to the mean proper motion of 47 Tuc, so we keep the mean proper motion of 47 Tuc fixed in this model with the value $(\bar{\mu}_{\alpha,1}, \bar{\mu}_{\delta,1}) = (0, 0)$. As the total probability density function must be normalised to unity and each of the constituent Gaussian distributions is normalised to unity, the sum of amplitudes must be equal to unity, i.e. $\sum_i A_i = 1$. This relation eliminates dependence of the total probability density function on one of the amplitudes. We choose to eliminate A_3 by setting $A_3 = 1 - A_1 - A_2$. In summary, three parameters

6.4. Data Cleaning Procedures

are eliminated from Eq. (6.7) by enforcing the relations

$$\bar{\mu}_{\alpha,1} = 0, \quad (6.9)$$

$$\bar{\mu}_{\delta,1} = 0, \quad (6.10)$$

$$A_3 = 1 - A_1 - A_2. \quad (6.11)$$

This reduces the number of parameters that $f_{\mu_\alpha, \mu_\delta}$ depends on in Eq. (6.7) from 12 to 9. The set of remaining parameters is

$$\theta = \{A_1, \sigma_1, A_2, \sigma_2, \bar{\mu}_{\alpha,2}, \bar{\mu}_{\delta,2}, \sigma_3, \bar{\mu}_{\alpha,3}, \bar{\mu}_{\delta,3}\}. \quad (6.12)$$

The best-fitting values of the remaining 9 parameters, including the mean proper motion components of the SMC, are then determined using the maximum likelihood estimate. Let $d = \{d_j\}$ be the set of observed proper motion data points, where $d_j = (\mu_{\alpha j}, \mu_{\delta j})$ is a single data point in proper motion space and j is an index that labels the data points. The likelihood $\mathcal{L}(\theta)$ is the probability (density) of the observed data given the parameters and model,

$$\mathcal{L}(\theta) = p(d|\theta) \quad (6.13)$$

$$= \prod_j p(d_j|\theta) \quad (6.14)$$

$$= \prod_j f_{\mu_\alpha, \mu_\delta}(\mu_{\alpha j}, \mu_{\delta j}; \theta). \quad (6.15)$$

The natural logarithm of the likelihood for our model of the proper motion distribution is thus

$$\ln \mathcal{L}(\theta) = \sum_j \ln f_{\mu_\alpha, \mu_\delta}(\mu_{\alpha j}, \mu_{\delta j}; \theta). \quad (6.16)$$

The maximum likelihood estimate $\hat{\theta}$ of the model parameters is the set of parameter values that maximises $\mathcal{L}(\theta)$, i.e. that maximises the probability of the observed data. In practice, it is more computationally feasible to minimise the negative log-likelihood, $-\ln \mathcal{L}(\theta)$, and doing so is equivalent

to maximising $\mathcal{L}(\theta)$. We minimise the negative of Eq. (6.16) numerically to get $\hat{\theta}$, which is the set of best-fitting parameters for our model.

The three-component Gaussian mixture model was fit to the proper motions for different F606W magnitude bins spanning the range 22 – 29 in increments of 0.5. The maximum likelihood estimates of the distribution parameters for each of the magnitude bins are given in Table 6.2. The standard deviations of both the 47 Tuc population and the SMC population increase with magnitude due to the increasing error in the proper motion measurements with magnitude. Note that the errors can be taken to be Gaussian-distributed and the convolution of two Gaussians is another Gaussian, so the proper motion errors are naturally accounted for in the Gaussian mixture model as an adjustment to the standard deviations of the populations. These increasing proper motion errors do not have a notable effect on the location of the SMC in proper motion space, which was found to be similar across all of the magnitude bins. The coordinates of the SMC in proper motion space are taken to be the average of the coordinates found in each bin. These coordinates are $(\bar{\mu}_{\alpha,2}, \bar{\mu}_{\delta,2}) = (4.76, 1.59)$.

Though the best-fitting parameters are determined using the joint distribution of μ_α and μ_δ , it is also instructive to see how the total proper motion is distributed for a particular population in polar coordinates after marginalising over the polar angle. This marginal distribution is derived below, and some of its important features are discussed.

Consider a population labelled by index i with proper motion distribution given by Eq. (6.8). Define the polar coordinate variables μ and ϕ relative to the mean proper motion of this population, such that

$$\mu_\alpha - \bar{\mu}_{\alpha,i} = \mu \cos \phi, \quad (6.17)$$

$$\mu_\delta - \bar{\mu}_{\delta,i} = \mu \sin \phi. \quad (6.18)$$

The determinant of the Jacobian matrix $\mathbf{J}(\mu, \phi)$ for the transformation from the Cartesian to polar coordinates is $\det |\mathbf{J}(\mu, \phi)| = \mu$.

F606W	A_1	σ_1	A_2	σ_2	$\bar{\mu}_{\alpha,2}$	$\bar{\mu}_{\delta,2}$	A_3	σ_3	$\bar{\mu}_{\alpha,3}$	$\bar{\mu}_{\delta,3}$
22.0 – 22.5	0.82	0.62	0.09	0.44	4.75	1.53	0.10	3.34	-1.07	0.39
22.5 – 23.0	0.77	0.60	0.10	0.47	4.74	1.49	0.13	2.76	-0.76	0.00
23.0 – 23.5	0.78	0.60	0.10	0.43	4.67	1.44	0.12	3.04	-0.82	-0.05
23.5 – 24.0	0.75	0.61	0.12	0.53	4.72	1.52	0.13	2.51	-0.93	-0.15
24.0 – 24.5	0.75	0.62	0.14	0.48	4.75	1.48	0.11	2.99	-0.57	0.06
24.5 – 25.0	0.70	0.64	0.17	0.52	4.71	1.51	0.13	3.27	-0.29	0.26
25.0 – 25.5	0.63	0.64	0.22	0.56	4.73	1.56	0.15	3.55	0.83	0.73
25.5 – 26.0	0.56	0.69	0.28	0.62	4.73	1.61	0.16	4.28	0.65	0.40
26.0 – 26.5	0.48	0.74	0.37	0.81	4.74	1.69	0.15	5.82	1.02	0.68
26.5 – 27.0	0.35	0.82	0.43	0.93	4.75	1.67	0.22	6.27	2.06	1.15
27.0 – 27.5	0.27	0.90	0.49	1.15	4.86	1.61	0.24	6.83	1.18	0.40
27.5 – 28.0	0.21	1.08	0.49	1.37	4.84	1.78	0.30	7.94	1.33	0.49
28.0 – 28.5	0.20	1.36	0.44	1.55	4.71	1.67	0.35	8.32	1.44	0.57
28.5 – 29.0	0.22	1.63	0.35	1.79	4.88	1.68	0.43	8.38	0.92	0.46

Table 6.2: Results of fitting proper motion distribution by F606W magnitude bin. The subscript indices of the parameters denote which population that parameter describes in a three-component Gaussian mixture model: 1 denotes 47 Tuc, 2 denotes the SMC, and 3 denotes the background. The average proper motion coordinates of the SMC across all magnitude bins are $\bar{\mu}_{\alpha,2} = 4.76$ and $\bar{\mu}_{\delta,2} = 1.59$.

6.4. Data Cleaning Procedures

Accounting for the relevant factor of $\det |\mathbf{J}(\mu, \phi)|$, the joint probability density function of μ and ϕ is thus given by the relation

$$f_{\mu, \phi, i}(\mu, \phi; \sigma_i) \quad (6.19)$$

$$= \mu f_{\mu_\alpha, \mu_\delta, i}(\bar{\mu}_{\alpha, i} + \mu \cos \phi, \bar{\mu}_{\delta, i} + \mu \sin \phi; \bar{\mu}_{\alpha, i}, \bar{\mu}_{\delta, i}, \sigma_1) \\ = \frac{\mu}{2\pi\sigma_i^2} \exp\left(-\frac{\mu^2}{2\sigma_i^2}\right). \quad (6.20)$$

Marginalising $f_{\mu, \phi, i}(\mu, \phi; \sigma_i)$ over ϕ then gives the probability density function of μ ,

$$f_{\mu, i}(\mu; \sigma_i) = \int_0^{2\pi} d\phi f_{\mu, \phi, i}(\mu, \phi; \sigma_i) \quad (6.21)$$

$$= \frac{\mu}{\sigma_i^2} \exp\left(-\frac{\mu^2}{2\sigma_i^2}\right). \quad (6.22)$$

This function goes to zero as μ goes to zero, i.e. $\lim_{\mu \rightarrow 0} f_{\mu, i}(\mu; \sigma_i) = 0$. Also note that the probability density of μ is maximized when $\mu = \sigma_i$; this is in contrast to the joint probability density of μ_α and μ_δ , which is maximised at the coordinates $(\bar{\mu}_{\alpha, i}, \bar{\mu}_{\delta, i})$ and thus $\mu = 0$.

If proper motion values are given relative to the mean motion of 47 Tuc, then the probability density distribution of μ values for the 47 Tuc population goes to zero as $\mu \rightarrow 0$. This is straight-forwardly given by Eq. (6.22) with $i = 1$, as the mean proper motion of the 47 Tuc population coincides with the origin of the polar coordinate system in proper motion space. For the SMC population, the distribution of the total proper motion relative to the mean proper motion of 47 Tuc is more complicated, as the mean proper motion of the SMC does not coincide with the origin of the polar coordinate system. However, the factor of μ that appears in Eq. (6.19), which is the determinant of the Jacobian in transforming the proper motion distribution function from Cartesian coordinates to polar coordinates, also appears in the analogous distribution for the SMC when μ is defined with respect to $(\bar{\mu}_{\alpha, 1}, \bar{\mu}_{\delta, 1})$ instead of $(\bar{\mu}_{\alpha, 2}, \bar{\mu}_{\delta, 2})$, and this factor likewise results in the distribution of μ for SMC stars going to zero as $\mu \rightarrow 0$. The SMC population

is also far enough away from the 47 Tuc population in proper motion space that very few SMC stars are expected to be found at the centre of the 47 Tuc distribution in proper motion space in the first place.

6.4.3.2 SMC Contamination

While the **SHARP** cleaning makes the 47 Tuc white dwarf cooling sequence and the SMC sequence appear more distinct in the CMD (see Fig. 6.1), these sequences still intersect at faint magnitudes (for F606W greater than about 27). The proper motion cut to select likely 47 Tuc members, which will be referred to as simply the “47 Tuc proper motion cut”, also does not remove all of the SMC stars that overlap with the faint white dwarfs in the cooling sequence. This is illustrated by Fig. 6.1a, where it can be seen that some of the objects selected by the 47 Tuc proper motion cut (black points) lie along the SMC sequence (mostly composed of grey points). Though Fig. 6.1a has not been **SHARP**-cleaned, most of the black points along the SMC sequence in Fig. 6.1a persist after **SHARP** cleaning, so the problem of SMC contamination remains. We want to quantify the number of SMC stars that are expected to be in the final fully-cleaned white dwarf sample that we use for the unbinned likelihood analysis. The number of SMC contaminants in the white dwarf data space is in general a function of magnitude, so we determine this number for F606W magnitude bins, using bins of 0.5 width spanning the range 22 – 29.

To quantify the number of SMC stars expected to contaminate our final white dwarf sample, we define cuts in both the CMD and proper motion space that each independently select stars that are very likely to be SMC members. The ratio of the number of SMC stars in the 47 Tuc white dwarf CMD region vs the SMC CMD region should be the same regardless of what proper motion cut is used (as long as some SMC stars survive the proper motion cut). Thus, the number of SMC contaminants in the white dwarf CMD region after the 47 Tuc proper motion cut can be estimated by calculating this ratio using a very pure sample of SMC stars selected using a proper motion cut and multiplying this ratio by the number of stars in

6.4. Data Cleaning Procedures

the SMC CMD region after the 47 Tuc proper motion cut.

Another way of understanding this procedure is that we count the number of proper-motion-selected SMC stars in the 47 Tuc white dwarf CMD region, then re-scale this number using the ratio of the number of CMD-selected SMC stars that survive the 47 Tuc proper motion cut to the number of CMD-selected SMC stars that survive the SMC proper motion cut. If the cuts to select SMC stars yield pure SMC samples, then this procedure estimates the number of SMC stars expected to both survive the 47 Tuc proper motion cut and fall within the 47 Tuc white dwarf CMD region.

Let $N_{\text{ps},\text{cw}}$ be the number of stars that survive both the tight proper motion cut to select SMC stars and the CMD cut to select white dwarfs in 47 Tuc. Let $N_{\text{ps},\text{cs}}$ be the number of stars that survive the same SMC proper motion cut used to get $N_{\text{ps},\text{cw}}$ and that also survive the CMD cut to select SMC stars. Finally, let $N_{\text{pt},\text{cs}}$ be the number of stars that survive the 47 Tuc proper motion cut and also survive the same CMD cut to select SMC stars used to get $N_{\text{ps},\text{cs}}$. Then the number of SMC stars expected to survive both the 47 Tuc proper motion cut and the 47 Tuc white dwarf CMD cut is

$$N_{\text{contam}} = \frac{N_{\text{ps},\text{cw}} N_{\text{pt},\text{cs}}}{N_{\text{ps},\text{cs}}}. \quad (6.23)$$

This is the expected number of SMC contaminants in the proper-motion-cleaned white dwarf sample that we ultimately use in the unbinned likelihood analysis. A formal treatment of the derivation of Eq. (6.23) is given in Appendix A.1.

The goal of both the SMC proper motion cut and the SMC CMD cut is to get a pure sample of SMC stars, and the boundaries of these cuts are chosen with this goal in mind. However, it is still possible that some 47 Tuc stars could survive these cuts, particularly the SMC CMD cut where the SMC sequence and 47 Tuc white dwarf cooling sequence begin to overlap. These misclassified 47 Tuc stars would cause our count for the corresponding number used in the calculation of N_{contam} to be too large. This is most likely to be an issue in determining the number of stars that survive one 47 Tuc cut and one SMC cut, i.e. $N_{\text{ps},\text{cw}}$ or $N_{\text{pt},\text{cs}}$, and in particular $N_{\text{pt},\text{cs}}$. It is

6.4. Data Cleaning Procedures

least likely to be an issue in determining the number of stars that survive both SMC cuts, i.e. $N_{\text{ps,cs}}$, as these stars should be the purest sample of SMC stars. Note that both $N_{\text{ps,cw}}$ and $N_{\text{pt,cs}}$ appear in the numerator on the right-hand side of Eq. (6.23), while $N_{\text{ps,cs}}$ appears in the denominator. N_{contam} is thus more properly an upper limit on the number of SMC stars that contaminate the 47 Tuc white dwarf data space after the 47 Tuc proper motion cut. As the boundaries for the SMC cuts are specifically chosen to reduce the risk of misclassifying 47 Tuc stars, the true number of SMC contaminants in the proper-motion-cleaned white dwarf data space should be close to this upper limit. Furthermore, if this upper limit is found to be negligibly small compared to the total size of the white dwarf sample, then that is sufficient information to deem the possibility of SMC contamination in the white dwarf data space to be of no further concern.

The boundaries defining the CMD-selected SMC sample are shown in Fig. 6.1 as the middle boundary region (orange lines). The CMD boundaries for the SMC population select predominantly the red side of the SMC sequence at faint magnitudes in order to avoid including 47 Tuc white dwarfs in the CMD-selected SMC population. The priority here is to select a pure population of SMC stars, even if it results in the exclusion of some SMC members. This SMC sample does not need to be complete for our analysis. The boundary region defining the 47 Tuc white dwarf CMD selection is also shown in Fig. 6.1, as the left-most boundary region (blue lines). This CMD boundary region for 47 Tuc white dwarfs is the same as the white dwarf data space that will be used in the unbinned likelihood analysis. In Fig. 6.1b, objects in the SHARP-cleaned data that survive the SMC proper motion cut are shown as black points, while the other objects in that data are shown as grey points.

The proper motions for all sources in our SHARP-cleaned data in a frame relative to the mean motion of 47 Tuc are shown in Fig. 6.3. For objects that lie within one of the three CMD boundary regions shown in Fig. 6.1, the CMD-selected population to which each object belongs is indicated by colour. These CMD-selected populations are 47 Tuc white dwarfs (blue), 47 Tuc main-sequence stars (green), and SMC stars (orange). Objects that do

6.4. Data Cleaning Procedures

not correspond to any of these three CMD-selected populations are shown in grey. The boundary of the 47 Tuc proper motion cut is shown in both Fig. 6.3a and Fig. 6.3b as a solid black curve. The boundary of the tight SMC proper motion cut is shown in Fig. 6.3a as a dashed black curve.

Note that the 47 Tuc white dwarfs and main-sequence stars largely overlap in the proper motion plots of Fig. 6.3, as expected since they belong to the same dynamical population. Most of the white dwarfs are obscured in Fig. 6.3a and at bright magnitudes in Fig. 6.3b simply because the main-sequence stars have been plotted on top of them. However, some of the white dwarfs of most interest to us are still visible, especially those at faint magnitudes in Fig. 6.3b where the number density of the main-sequence stars is much lower than at brighter magnitudes. In Fig. 6.3a, it can also be seen that some of the objects selected by the 47 Tuc white dwarf CMD cut fall within the boundary of the SMC proper motion cut. These objects correspond to the black points in Fig. 6.1b that lie within the white dwarf CMD boundary. As the SMC proper motion cut is very tight about the mean motion of the SMC and far from the bulk of the 47 Tuc proper motion distribution, these objects are highly likely to be SMC stars that lie within the 47 Tuc white dwarf CMD boundary region.

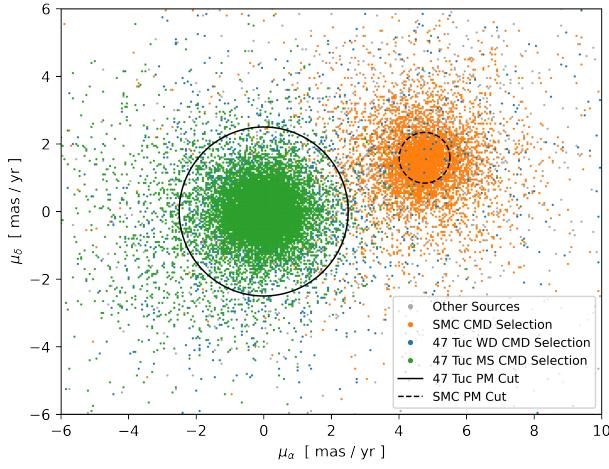
Fig. 6.3a shows the two-dimensional distribution of the components of proper motion in the directions of right ascension and declination, respectively μ_α and μ_δ . Note that a factor of $\cos\delta$ is included in the definition of μ_α (i.e. $\mu_\alpha = \dot{\alpha}\cos\delta$ and $\mu_\delta = \dot{\delta}$, where the overdot denotes a derivative with respect to time), making (μ_α, μ_δ) the tangent plane projection of the proper motion vector. Two distinct populations are clearly visible in Fig. 6.3a. The population on the left and centred at $(0, 0)$ corresponds to 47 Tuc, while the population on the right corresponds to the SMC.

The 47 Tuc proper motion cut selects objects with

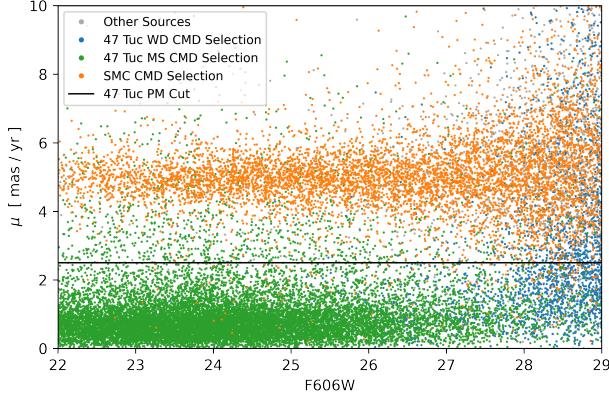
$$\sqrt{(\mu_\alpha - \bar{\mu}_{\alpha,t})^2 + (\mu_\delta - \bar{\mu}_{\delta,t})^2} < 2.5 \text{ mas yr}^{-1}, \quad (6.24)$$

where $(\bar{\mu}_{\alpha,t}, \bar{\mu}_{\delta,t})$ is the mean proper motion of 47 Tuc, which has the value $(\bar{\mu}_{\alpha,t}, \bar{\mu}_{\delta,t}) = (0, 0)$ in the reference frame of Fig. 6.3. This corresponds to

6.4. Data Cleaning Procedures



(a) Components of proper motion vector in tangent plane.



(b) Total proper motions as a function of F606W magnitude.

Figure 6.3: Proper motions of SHARP-cleaned data relative to mean motion of 47 Tuc. The CMD-selected population to which each source belongs is indicated by colour: 47 Tuc white dwarfs (blue), 47 Tuc main-sequence stars (green), SMC stars (orange), and other sources with $22 < F606W < 29$ but not in a CMD boundary region (grey). The boundary of the proper motion cut to select 47 Tuc members is shown as a solid black curve indicating a circle of radius 2.5 mas yr^{-1} centred on the mean 47 Tuc proper motion coordinates $(0, 0)$. The boundary of the tight SMC proper motion cut used in the proper motion calibration procedure is shown in sub-figure (a) as a dashed black curve of radius 0.75 mas yr^{-1} centred on the mean motion of the SMC at $(4.76, 1.59)$.

6.4. Data Cleaning Procedures

selecting the points in Fig. 6.3a that lie within the circle of radius 2.5 centred on $(0, 0)$, which is shown as a solid black line.

The SMC proper motion cut selects objects with

$$\sqrt{(\mu_\alpha - \bar{\mu}_{\alpha,s})^2 + (\mu_\delta - \bar{\mu}_{\delta,s})^2} < 0.75 \text{ mas yr}^{-1}, \quad (6.25)$$

where $(\bar{\mu}_{\alpha,s}, \bar{\mu}_{\delta,s})$ is the mean proper motion of the SMC, which we take to be $(\bar{\mu}_{\alpha,s}, \bar{\mu}_{\delta,s}) = (4.76, 1.59)$ in the reference frame of Fig. 6.3. This corresponds to selecting the points in Fig. 6.3a that lie within the circle of radius 0.75 centred on $(4.76, 1.59)$, which is shown as a dashed black line.

The value that we use for the mean motion of the SMC was determined by fitting a three-component Gaussian mixture model to the proper motion data, where the first of these Gaussian components accounts for the 47 Tuc population, the second one accounts for the SMC population, and the third one accounts for outliers and background contaminants like field stars. The best-fitting parameters of this model for the proper motion distribution were determined using the maximum likelihood estimate. This model and fitting procedure are described in detail in Section 6.4.3.1.

Fig. 6.3b shows the total proper motion in the tangent plane as a function of F606W magnitude. This total proper motion is defined in terms of μ_α and μ_δ as

$$\mu = \sqrt{\mu_\alpha^2 + \mu_\delta^2}. \quad (6.26)$$

Since the mean proper motion of 47 Tuc is located at proper motion coordinates $(0, 0)$, this is the total proper motion relative to the 47 Tuc mean. As in Fig. 6.3a, distinct 47 Tuc and SMC populations can also be seen in Fig. 6.3b. Most of the 47 Tuc stars have $\mu < 2.5 \text{ mas yr}^{-1}$, while the total proper motions of the SMC stars are clustered near $\mu \sim 5 \text{ mas yr}^{-1}$. Note that the density of stars at $\mu = 0$ is approximately zero due to the factor of μ that appears through the Jacobian in transforming the proper motion distribution function from Cartesian coordinates to polar coordinates, as explained in Section 6.4.3.1.

In terms of the total proper motion given by Eq. (6.26), the 47 Tuc

6.4. Data Cleaning Procedures

proper motion cut selects objects with $\mu < 2.5 \text{ mas yr}^{-1}$, which corresponds to all objects that lie below the solid black line in Fig. 6.3b. The SMC proper motion cut is more complicated in terms of μ . The distance between the mean proper motion of the SMC and the mean proper motion of 47 Tuc in proper motion space is $\bar{\mu}_s = \sqrt{\bar{\mu}_{\alpha,s}^2 + \bar{\mu}_{\delta,s}^2}$, which has the value $\bar{\mu}_s = 5.02 \text{ mas yr}^{-1}$. So the smallest value of μ along the boundary of the SMC proper motion cut is $\bar{\mu}_s - 0.75 \text{ mas yr}^{-1} = 4.27 \text{ mas yr}^{-1}$, and the largest value of μ along the boundary of the SMC proper motion cut is $\bar{\mu}_s + 0.75 \text{ mas yr}^{-1} = 5.77 \text{ mas yr}^{-1}$. Thus, all objects selected by the SMC proper motion cut have total proper motion in the range $4.27 \text{ mas yr}^{-1} < \mu < 5.77 \text{ mas yr}^{-1}$; however, not all objects with μ in this range are actually selected by the SMC proper motion cut.

The increasing uncertainty of the proper motion measurements with increasing magnitude causes fainter stars in both 47 Tuc and the SMC to appear more dispersed, which manifests in Fig. 6.3b as the increasingly large spread of μ values along the y-axis as the magnitude increases. This leads to the proper motion distributions of the two populations overlapping more as the magnitude increases. The number of SMC contaminants in the 47 Tuc white dwarf data space is thus expected to increase with magnitude, especially for $F606W \gtrsim 27$ where the proper motions start to become noticeably more dispersed and the CMD sequences of the SMC stars and 47 Tuc white dwarfs begin to intersect.

The value of N_{contam} that we calculate in each $F606W$ magnitude bin is given in Table 6.3. For reference, the total number of objects, N_{WD} , found in that bin for our proper-motion-cleaned white dwarf data space (i.e. after both the 47 Tuc proper motion cut and the 47 Tuc white dwarf CMD cut) is also given in Table 6.3. The estimate of the true number of white dwarfs in that bin is $N_{\text{WD}} - N_{\text{contam}}$. Since N_{contam} is really an upper limit on the number of SMC contaminants in the 47 Tuc white dwarf data space (which should also be close to the actual number of contaminants), the quantity $N_{\text{WD}} - N_{\text{contam}}$ is correspondingly really a lower limit on the true number of white dwarfs. Note that the interpretation of N_{contam} as an upper limit makes it sensible to report non-integer values for N_{contam} .

F606W	N_{WD}	N_{contam}	f_{CR}	Error (f_{CR})
22.0 – 22.5	1	0.0	0.9413	0.0062
22.5 – 23.0	1	0.0	0.9373	0.0058
23.0 – 23.5	5	0.0	0.9349	0.0055
23.5 – 24.0	10	0.0	0.9359	0.0050
24.0 – 24.5	28	0.0	0.9431	0.0051
24.5 – 25.0	27	0.0	0.9328	0.0060
25.0 – 25.5	44	0.0	0.9293	0.0069
25.5 – 26.0	48	0.0	0.9114	0.0088
26.0 – 26.5	75	0.0	0.9093	0.0106
26.5 – 27.0	78	0.0	0.9087	0.0131
27.0 – 27.5	117	0.1	0.8849	0.0167
27.5 – 28.0	150	0.0	0.8161	0.0240
28.0 – 28.5	320	0.6	0.7419	0.0351
28.5 – 29.0	424	4.5	0.7016	0.0411

Table 6.3: Results of calibrating proper motion cleaning procedure by F606W magnitude bin. The number of objects in the proper-motion-cleaned 47 Tuc white dwarf data space (N_{WD}) is the total number of objects in a given magnitude bin that survive both the 47 Tuc proper motion cut and the 47 Tuc white dwarf CMD cut. The number of contaminants (N_{contam}) is the estimated number of SMC stars that survive the same cuts used to calculate N_{WD} . The completeness reduction factor (f_{CR}) is the fraction of CMD-selected 47 Tuc main-sequence stars that survive the 47 Tuc proper motion cut. The error for f_{CR} is reported in the final column, following its value. All of these quantities were calculated using the SHARP-cleaned data.

6.4. Data Cleaning Procedures

From the values tabulated in Table 6.3, it can be seen that N_{contam} is negligible compared to N_{WD} for our choice of 47 Tuc proper motion cut. Most magnitude bins contain no contaminants, and even for the faintest magnitudes, most of the bins have $N_{\text{contam}} < 1$. The largest number of contaminants is found in the faintest magnitude bin, $28.5 < \text{F606W} < 29.0$, and N_{contam} is still negligible compared to N_{WD} for this bin. As our cleaning procedure successfully removes all but a negligible number of SMC stars from the data, we do not need to apply a correction in our unbinned likelihood analysis to account for SMC contaminants in the white dwarf data space. However, the proper motion cleaning also removes some objects that are actually 47 Tuc white dwarfs. This reduces the completeness of our white dwarf sample below what is found from the artificial stars tests. This completeness reduction effect is analysed below in Section 6.4.3.3, and the result of that analysis is also included in Table 6.3.

6.4.3.3 Completeness Reduction

The reduction in completeness from the proper motion cut to select likely 47 Tuc members is quantified using 47 Tuc main-sequence stars. These main-sequence stars are identified in the CMD using the same bounding region as was used for the SHARP calibration, which is fully shown in Fig. 6.1a as the right-most bounding region (in green). The proper motion cut used in our ultimate cleaning procedure for the white dwarf data, which selects objects with a total proper motion $< 2.5 \text{ mas yr}^{-1}$ relative to the mean proper motion of 47 Tuc, was then applied to the CMD-selected 47 Tuc main-sequence sample. The fraction of main-sequence stars remaining after the proper motion cut quantifies the completeness reduction due to the proper motion cleaning procedure. This survival fraction corresponds to the magnitude-dependent fraction of green points that lie below the solid black line in Fig. 6.3b and tends to decreases with magnitude as the proper motion uncertainties increase. The completeness reduction was determined as a function of F606W magnitude by sorting the main-sequence sample into bins of 0.5 magnitude width over the magnitude range 22 to 29 and

Parameter	Value
$F606W_0$	27.00 ± 0.17
$f_{CR,0}$	0.9158 ± 0.0059
a_1	-0.0060 ± 0.0017
a_2	-0.1297 ± 0.0247

Table 6.4: Best-fitting parameter values for piece-wise linear model of f_{CR} as a function of $F606W$.

calculating the fraction remaining in each bin after the proper motion cut. These are the same magnitude bins used to calculate N_{contam} . The results are tabulated in Table 6.3.

Let N_{cm} be the number of CMD-selected 47 Tuc main-sequence stars without any proper motion cut applied, and let $N_{\text{pt,cm}}$ be the number of CMD-selected 47 Tuc main-sequence stars that survive the 47 Tuc proper motion cut. We define a completeness reduction factor f_{CR} , which is given by the survival fraction of the CMD-selected main-sequence stars

$$f_{CR} = \frac{N_{\text{pt,cm}}}{N_{\text{cm}}}. \quad (6.27)$$

The standard error in f_{CR} is taken to be binomially-distributed and thus given by

$$\text{Error}(f_{CR}) = \sqrt{\frac{f_{CR}(1-f_{CR})}{N_{\text{cm}}}}. \quad (6.28)$$

The errors calculated in this way are reported in the final column of Table 6.3, following the corresponding values of f_{CR} .

To avoid numerical artefacts due to binning when using f_{CR} in the unbinned likelihood analysis, we modelled f_{CR} as a piece-wise linear function of $F606W$ with parameter values determined by fitting this function to the reference values given in Table 6.3. The analytic function used for f_{CR} consists of two linear segments, with the switch occurring at the $F606W$ value $F606W_0$ and corresponding f_{CR} value $f_{CR,0}$. We let $F606W_0$ and $f_{CR,0}$ be free parameters of the fit, along with the slope a_1 of the seg-

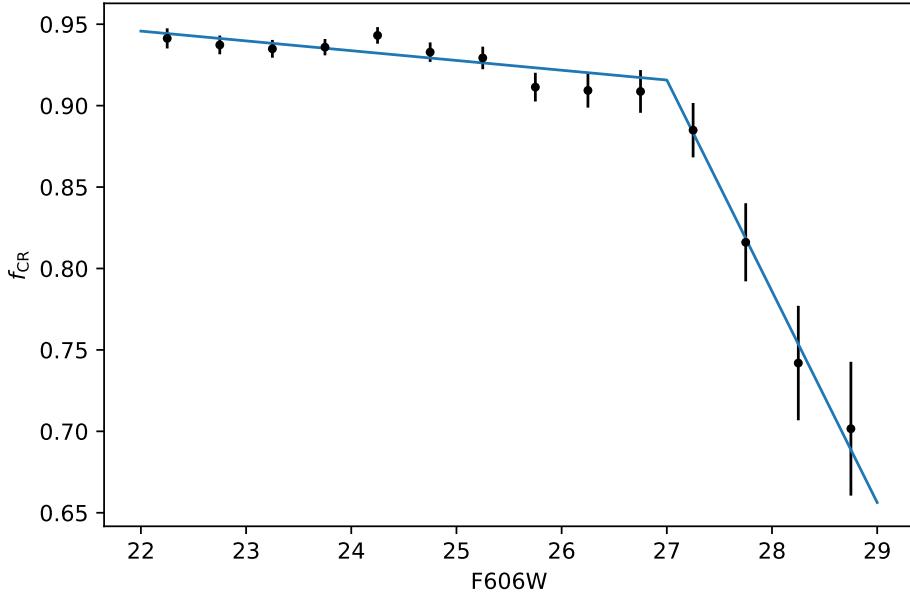


Figure 6.4: Completeness reduction factor f_{CR} due to proper motion cleaning as a function of F606W magnitude. The analytic function (blue curve) is shown for the best-fitting parameters determined by fitting the binned values of f_{CR} (black points) calculated in the calibration of the cleaning procedure.

ment where $F606W < F606W_0$ and the slope a_2 of the other segment where $F606W \geq F606W_0$. The best-fitting parameter values are given in Table 6.4, and the corresponding best-fitting function is plotted in Fig. 6.4, along with the values it was fitted to. The reduced chi-squared χ_ν value (with $\nu = 10$ degrees of freedom) is $\chi_\nu = 0.976$, indicating a good fit.

The completeness reduction due to proper motion cleaning is accounted for in our unbinned likelihood analysis after applying the photometric error distribution function, Eq. (6.1), to the cooling model. Note that the error distribution is taken to depend on the input magnitudes, whereas the completeness reduction factor should be treated as a function of the output magnitudes after accounting for photometric errors (as f_{CR} was calculated directly from observations). After applying the error distribution to the

cooling model, the resultant distribution function (which is a function of the output magnitudes) is then multiplied by f_{CR} to account for the completeness reduction in the unbinned likelihood analysis.

6.5 Models

We created white dwarf cooling models using the stellar evolution software MESA [217–221]¹⁹. Using MESA, we ran a suite of white dwarf cooling simulations for different parameter values, varying the white dwarf mass, the treatment of diffusion, and the thickness of the H envelope. This suite of simulations was generated from an initial model of a young, hot white dwarf that was created by simulating the pre-white dwarf evolution of a progenitor star. The simulation that produced this initial model is described in Section 6.5.1. The white dwarf cooling models are described in Section 6.5.2.

6.5.1 Creation of Initial Model

The initial model for the white dwarf cooling simulations was created by simulating the evolution of a $0.9 M_{\odot}$ progenitor star with parameters appropriate for 47 Tuc from the pre-main sequence until the birth of the white dwarf. This simulation was done using MESA version 10398 (mesa-r10398) and was created from the MESA `test_suite` example `1M_pre_ms_to_wd`. We modified the parameters of the `1M_pre_ms_to_wd` inlist by changing the initial mass, the initial composition parameters, and the wind parameters. We set the following initial parameters

```
initial_mass = 0.9d0
initial_z = 4.0d-3
initial_y = 0.256d0
```

and set `Zbase` (for use with Type 2 opacities) to be the same as `initial_z`.

We used a Reimers mass loss scheme [229] on the RGB and a Blocker mass loss scheme [230] on the AGB. Previous work has shown that stars in

¹⁹There is also a more recent sixth instrument paper [222]; however, it is only relevant for MESA versions released later than the oldest version used in this work.

47 Tuc lose most of their mass on the AGB [231], rather than the RGB, so we set the scaling parameter values

```
Reimers_scaling_factor = 0.1d0
Blocker_scaling_factor = 0.7d0
```

for our prescriptions of mass loss via stellar winds. As we are not primarily concerned with the details of stellar evolution before the white dwarf stage, the particular choice of scaling factors is not a concern for this work; the values simply need to be reasonable for 47 Tuc and produce a white dwarf with a thick H envelope from which our set of white dwarf models can be created.

A custom stopping condition was used to ensure that the simulation ended shortly after the star became a white dwarf. The simulation ended when the two conditions `log_Teff > 4.5` and `log_L < 2` were both met, with the code for this implemented through the `run_star_extras` module. This stopped the simulation before the luminosity of the white dwarf had dropped into the luminosity range of interest for studying white dwarf cooling in 47 Tuc, and thus the output of this simulation can be used as a starting point to generate white dwarf cooling models that span the entire luminosity range of interest. In the end, the simulation of pre-main sequence to white dwarf evolution produced a model of a newly born white dwarf with a mass of $0.5388 M_{\odot}$ and a thick H envelope.

6.5.2 White Dwarf Cooling Models

Starting from the model created by the simulation described in Section 6.5.1, we generated additional initial models with different masses and envelope thicknesses from which to begin the white dwarf cooling simulations. We simulated the evolution of white dwarfs with these different parameter values using different treatments of diffusion as described below. Both the procedures to modify the initial model and the main white dwarf cooling simulations were performed using MESA version 15140 (mesa-r15140).

From the initial model of a $0.5338 M_{\odot}$ white dwarf, we created less massive white dwarf models with masses of 0.5092, 0.5166, 0.5240, and

0.5314 M_{\odot} using the `relax_mass_scale` control provided by MESA and running a brief simulation in which the model was allowed to evolve for a few short time steps to adjust to the change. Note that a mass of 0.5240 M_{\odot} was chosen as one of the target mass values because this was the mass of the model used in Obertas *et al.* [30], which was produced by a stellar evolution simulation of a 0.974 M_{\odot} progenitor for 47 Tuc and was found by Obertas *et al.* [30] to reasonably replicate the cooling curve of the old white dwarfs in the same data as used in this work. The procedure to reduce the mass of the initial model simply re-scaled the profile of the 0.5338 M_{\odot} input model to the target mass, so like the 0.5338 M_{\odot} model, the less massive models produced in this way also have thick H envelopes.

We found that the `relax_mass_scale` procedure was only able to successfully re-scale models to a target mass that was lower than the mass of the initial model. In order to extend the mass grid to masses above 0.5338 M_{\odot} , we thus created another, heavier initial model from which the simulations for more massive white dwarfs could be generated. We created a 0.5644 M_{\odot} initial white dwarf model using a simulation based on the `test_suite` example `make_co_wd` with the relevant parameters modified to make the simulation appropriate for 47 Tuc. From the new 0.5644 M_{\odot} initial white dwarf model, we used the `relax_mass_scale` procedure described above to create models with white dwarf masses of 0.5462 and 0.5536 M_{\odot} . We also checked to confirm that relaxing the heavier initial white dwarf model down to lower masses produced the same cooling models as starting from the 0.5388 M_{\odot} initial model. In total, we created white dwarf simulations for white dwarf masses spanning the range 0.5092 – 0.5535 M_{\odot} (inclusive) in increments of 0.0074 M_{\odot} .

For each of these masses, models with thinner envelopes were created by using the `relax_mass` control. This removes mass from the white dwarf via a wind, which takes the mass from the H envelope. For this procedure the new mass is set to be just slightly smaller than the mass of the input model, with the difference being the amount of mass to remove from the H envelope. As even the thickest envelopes only have a mass on the order of 10^{-4} times the total mass M_{WD} of the white dwarf, the change to M_{WD}

due to reducing the mass of the envelope in this way is negligible in terms of the effect that varying M_{WD} can have on the cooling curves. For white dwarfs with thick envelopes, the envelope thickness also decreases over time at early times in the white dwarf’s evolution due to residual H burning near the boundary of the He layer and H envelope. This residual H burning does not change the total mass of the white dwarf, but it makes the envelope thickness in general a function of time. To define a parameter quantifying the thickness of the H envelope, we select a reference cooling time of 10 Myr into the white dwarf cooling simulations at which to define the envelope thickness parameter, and for this parameter we use the relative mass of the H envelope $q_H = M_H/M_{\text{WD}}$, where M_H is the mass of the H envelope.

For each value of M_{WD} , white dwarf evolution was simulated for a range of q_H values using three different treatments of diffusion: standard diffusion as implemented by MESA, a custom modified treatment of diffusion, and no diffusion. The diffusion equations solved by MESA implicitly assume an ideal gas approximation. As we are interested in white dwarf evolution at cooling times well after the phase transition from gas to liquid has occurred in the white dwarf core, the ideal gas approximation should not be assumed to hold throughout the duration of the white dwarf cooling simulations. We thus consider a modified treatment of diffusion that corrects for non-ideal gas effects. We do this by modifying the parameter `SIG_factor` in MESA’s diffusion solver routine, which by default is set equal to the inlist control parameter `diffusion_SIG_factor` at every location in the star. In our modified diffusion routine, this parameter is instead set to be

$$\text{SIG_factor} = \text{diffusion_SIG_factor} \times \begin{cases} 0.3 & t \leq t_{\text{on}} \\ \frac{1}{[1 + (\Gamma_k/c_3)^{c_4}]} & t > t_{\text{on}} \end{cases} \quad (6.29)$$

at each position cell indexed by k with plasma coupling parameter Γ_k , where $t_{\text{on}} = 5 \times 10^3$ yrs and c_3 and c_4 are parameters that can be set in the inlist using custom `x_ctrl` inlist control parameters. The case of $t \leq t_{\text{on}}$ is simply to give the model time to adjust before implementing the main modified diffusion code. The initial relaxation period $t \leq 10^4$ yrs, which includes all of

the period $t \leq t_{\text{on}}$ and the very early part of the main simulation, is removed from each cooling model before performing the analysis, so continuity at $t = t_{\text{on}}$ is not a concern. The modification to diffusion given by Eq. (6.29) uses a very general functional form. We consider a fiducial case where the parameters are set to be $c_3 = 0.0625$ and $c_4 = 1$. We obtained this function form and estimated the value of c_3 (to set as the fiducial value) by performing a small suite of molecular dynamics simulations of hydrogen-helium plasmas with different values of Γ . We implement our modified diffusion routine via the `other_diffusion` subroutine in `run_stars_extras`. The treatment of element diffusion in MESA and our modified treatment of diffusion are explained in more detail in Appendix A.2.

The white dwarf cooling simulations were all based on the `test_suite` example `wd_cool_0.6M` provided by MESA with the modifications described above. We also set `use_Skye` to be false²⁰ and set `Zbase` to be 0.004, the same value used for the simulation that produced the initial 47 Tuc white dwarf model. These simulations were run for 12.5 Gyr of cooling time or until T_{eff} reached a lower limit of 2,000 K, whichever condition was reached first. These stopping conditions enabled the white dwarf models to become sufficiently faint to span the entire magnitude range of our data space while remaining within the parameter regime limits of MESA.

The cooling models produced by these simulations describe how luminosity changes over time, and the cooling curves showing this relation for the different parameter combinations that were simulated are plotted in Fig. 6.5. The different curves within a given sub-figure show the effect of varying the envelope thickness for a fixed white dwarf mass and diffusion treatment. The different rows of sub-figures show the effect of varying the white dwarf mass, with the mass increasing from the top to bottom row; only a few select masses are shown for the sake of visualisation. The different columns correspond to the different diffusion scenarios: no diffusion (left), modified diffusion (middle), and standard diffusion (right).

²⁰The Skye equation of state was experimental in mesa-r15140, the version of MESA used for our cooling simulations.

6.5. Models

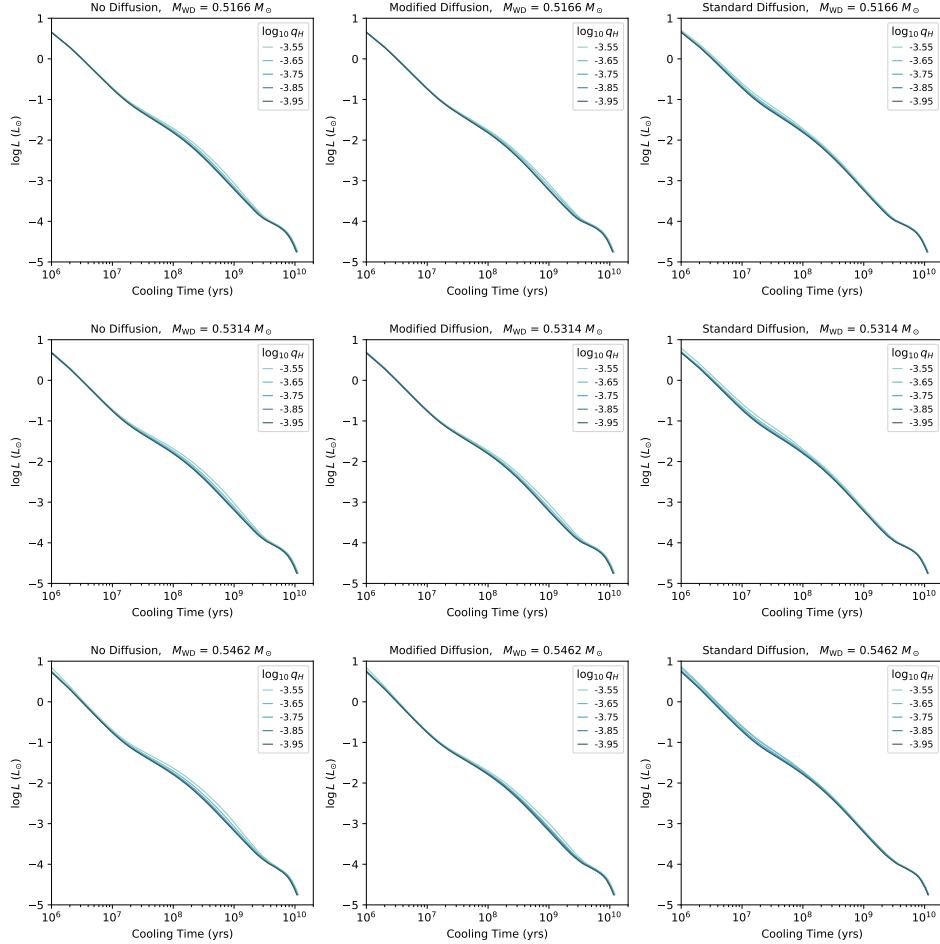


Figure 6.5: Theoretical cooling curves for different model parameters varied in MESA simulations. Each subplot shows cooling curves for white dwarfs with the same mass (M_{WD}) and diffusion treatment but different H envelope thickness at a reference cooling time of 10 Myr. The envelope thickness is indicated in the legend by the parameter q_H , which gives the ratio of the mass of H in the white dwarf to the total white dwarf mass at the reference time. Each row corresponds to a fixed white dwarf mass, which from top to bottom are $0.5166 M_{\odot}$, $0.5314 M_{\odot}$, and $0.5462 M_{\odot}$. Each column corresponds to a particular diffusion scenario, which from left to right are no diffusion, our custom modified diffusion, and standard MESA diffusion.

In the standard diffusion scenario, the faster diffusion, in particular at the boundary of the H envelope, causes the white dwarfs to undergo stronger residual H burning at early times and finish this residual burning at earlier times than in the other diffusion scenarios considered. In the modified treatment of diffusion, it takes longer for the elements to diffuse, so the residual H burning is delayed, which results in a cooling curve more similar to the scenario of no diffusion at early times.

The final bump in the cooling curve that occurs after a few billion years of cooling time in Fig. 6.5 (in all of the plots shown) corresponds to the onset of convective coupling of the envelope, which also approximately coincides with the onset of core crystallisation for white dwarfs in 47 Tuc [30]. This feature is sensitive to the envelope thickness and occurs at a luminosity where there are a large number of white dwarfs in the deep ACS data set.

6.6 Birthrate

The white dwarf birthrate for our sample is one of the parameters that is determined in the unbinned likelihood analysis described in Section 6.7. Before performing that analysis, however, we first determine a prior for the birthrate using observations of red giant stars from *Gaia* EDR3 data in the HST footprint. Since the members of a star cluster population are approximately the same age and the stars evolve quickly through the evolutionary stages between the end of the main sequence stage and start of the white dwarf stage, the rate of stars leaving the main sequence should be approximately the same as the white dwarf birthrate, and this rate can be measured using stars on the RGB.

The HST data cannot be used for this birthrate calculation because the red giant stars are so bright that they saturate the deep observations, leading to high incompleteness. Instead, we must use another dataset and ensure that we select stars over the same field of view as was used for the HST white dwarf observations. *Gaia* EDR3 observations are essentially complete for $17 < G < 12$, which spans nearly the entire RGB of 47 Tuc, and have

6.6. Birthrate

the most precise astrometry ever measured²¹, making *Gaia* EDR3 a good dataset for this purpose.

We first retrieved all *Gaia* EDR3 sources within a radius²² far enough from the centre of 47 Tuc that the selection included the entire HST field of view. The *Gaia* data is publicly available and was retrieved through Vizier. We then selected the *Gaia* EDR3 sources within the field boundaries of the HST ACS/WFC observations of our white dwarf data by performing a cut in position space to get a sample of *Gaia* EDR3 sources in the HST footprint. To get the field boundary for the full HST observation, we merged the boundaries of all of the orbits. The observational plan grouped the orbits into 24 visits, and the boundaries for the 24 visits (for which the observing regions of the constituent orbits²³ have been merged) are publicly available through MAST. We merged the boundaries provided by MAST for these visits to get the boundary for the final stacked image that combined the observations from all of the orbits. The final merged boundary is shown as a red curve in Fig. 6.6, where it is plotted over the *Gaia* EDR3 data (shown as the black points). Figure 6.6a shows the *Gaia* EDR3 data before selecting sources within the HST field boundary; this plot is approximately centred on the centre of 47 Tuc and shows the location of the HST observations relative to the cluster centre. Figure 6.6b shows only the *Gaia* EDR3 sample after selecting sources within the HST field boundary, with the plot approximately centred on the centre of the HST field boundary.

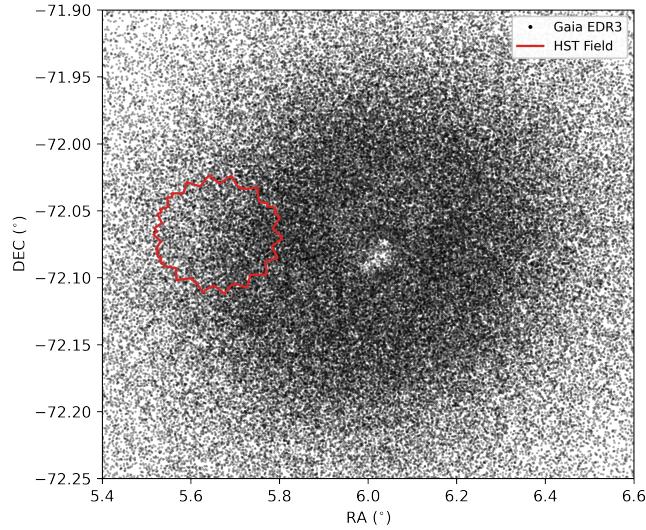
The dependence of the number density and completeness on the radial distance from the cluster centre is clearly shown in Fig. 6.6a. The number density in general increases as the distance to the cluster centre decreases; however, overcrowding near the very centre of the cluster notably reduces the completeness in this region, leading to an absence of stars observed at the centre where the cluster is most densely populated. The dependence

²¹Though the more recent *Gaia* DR3 is now available, the astrometric measurements are the same as *Gaia* EDR3.

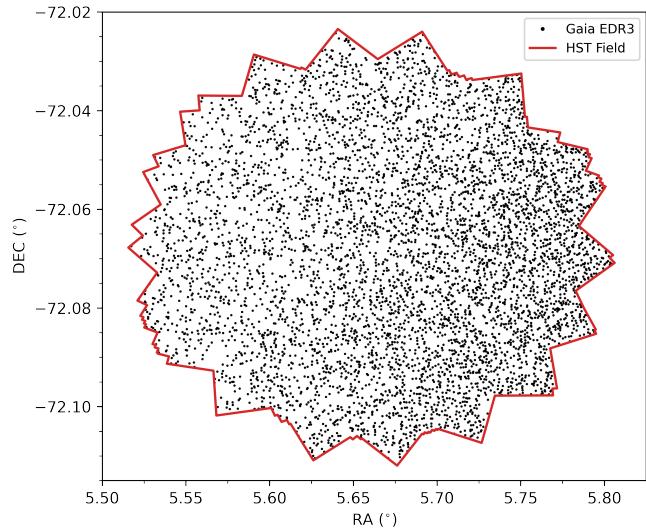
²²A value of 5° was used for this radius, but the exact value is not important as long as it is large enough to include the entire HST field of view.

²³Each visit consists of a group of five orbits (except for Visit 24, which consists of six orbits). See Kalirai *et al.* [223] for details.

6.6. Birthrate



(a) Zoomed-out view before selection.



(b) Zoomed-in view after selection.

Figure 6.6: Field boundaries (red curves) for the HST ACS/WFC deep observations overlaid on Gaia EDR3 observations (black points) of 47 Tuc. The Gaia EDR3 sources that are located within these boundaries are used for the birthrate calculations.

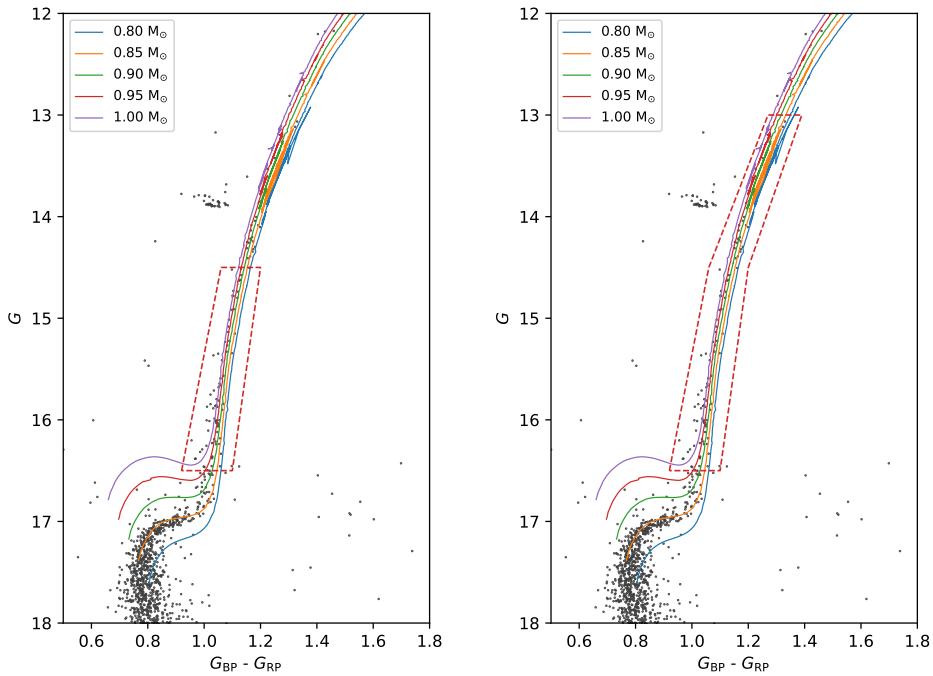
6.6. Birthrate

of the completeness and thus photometric error distribution on the radial coordinate from the cluster centre would be a concern if the region observed to gather the white dwarf data were closer to the centre; however, the field of view for the HST observations considered in this chapter is far enough away from the cluster centre that the photometric error distribution does not depend appreciably on radius. The dependence of the number density on radius can also be seen in Fig. 6.6b, where the number density of the data points increases with increasing values of right ascension, which correspond to positions closer to the cluster centre within the HST field region shown in that plot. Since the photometric error distribution does not depend on radius, the dependence of the number density on radius does not need to be accounted for in the unbinned likelihood analysis.

From the sample of Gaia EDR3 sources within the HST field boundaries, we then selected RGB stars using a cut in the CMD. The main boundary regions for selecting RGB stars is shown in Fig. 6.7 by the dashed red curves. Figure 6.7 also shows stellar evolution models (solid curves) for various initial masses plotted over the Gaia EDR3 data (black points). The models are shown in Fig. 6.7 from the end of the main sequence stage (taken to be when the mass fraction of hydrogen at the centre of the star drops below 10^{-4}) until the tip of the RGB. The birthrate is calculated by dividing the number N_{RGB} of stars observed in the RGB boundary region by the time t_{RGB} it takes a star to traverse this region according to the models, so (the prior value for) the birthrate is given by $\dot{N}_0 = N_{\text{RGB}} / t_{\text{RGB}}$. To access the uncertainty in the birthrate, we calculated the birthrate using models of different initial masses and using different RGB boundary definitions consisting of sub-regions of the main boundary regions shown in Fig. 6.7.

The stellar evolution models used to calculate the birthrate (and shown in Fig. 6.7) were created by running MESA simulations analogous to the simulation described in Section 6.5.1 (which created the initial white dwarf model from which the white dwarf cooling simulations were generated) but with different initial masses. Like the simulation of Section 6.5.1, these simulations were based on the MESA `test_suite` example `1M_pre_ms_to_wd` with composition appropriate for 47 Tuc (i.e. with an `initial_z` of 4×10^{-4}

6.6. Birthrate



(a) Selection excluding red-giant bump. (b) Selection including red-giant bump.

Figure 6.7: CMD selections of RGB stars for birthrate calculations. The black points correspond to Gaia EDR3 data in the HST footprint, while the dashed red curves indicate the boundaries to select the RGB stars. From right to left, the solid curves correspond to stellar evolution models with initial masses of 0.80 , 0.85 , 0.90 , 0.95 , and $1.00 M_{\odot}$.

6.6. Birthrate

and `initial_y` of 0.256) and were run using mesa-r10398. The magnitudes (and colour) in the *Gaia* bandpass filters were calculated for these MESA evolution models using bolometric corrections. We used the bolometric corrections²⁴ of Chen *et al.* [232] from PARSEC [233] calculated using synthetic spectra from a mix of the ATLAS9 ODFNEW [234] and PHOENIX BT-Settl [235] spectral libraries (combined with the spectra from COMARCS [236, 237] for cool giants and Chen *et al.* [238] for very hot stars) with the transmission curves of the *Gaia* EDR3 bandpass filters [239, 240] provided on the ESA/Gaia website²⁵. We used the same values for distance modulus and colour excess as are used in the unbinned likelihood analysis for the main white dwarf data, i.e. $\mu = 13.24$ and $E(B - V) = 0.04$, and we likewise used the same extinction curves, i.e. those of Cardelli *et al.* [38] and O'Donnell [39] with a total extinction²⁶ of $A_V = 0.124$ and a relative visibility of $R_V = 3.1$. The procedure for applying the bolometric corrections is essentially the same as what is described in Section 6.7 for the white dwarf models; we simply used different filters and got the corresponding bolometric corrections from a different source for the RGB models used here compared to the white dwarf models used in the main analysis later.

The main models used for the birthrate calculations were those with initial masses of 0.85, 0.90, and $0.95 M_\odot$. The models shown in Fig. 6.7 indicate that stars with an initial mass of $\sim 0.85 M_\odot$ are just beginning to leave the main sequence, which is in keeping with expectations from other work (e.g. [241–244]). The $0.80 M_\odot$ model is only shown for reference to illustrate the mass-dependence of the main sequence turnoff; it was not used for any of the birthrate calculations because the models indicate that stars with mass this low in 47 Tuc have not yet left the main sequence. The $1.00 M_\odot$ model was only used when calculating birthrates for truncated boundary regions that do not include its subgiant stage (i.e. boundary selection with a maximum value of $G \leq 16.4$). The models with initial mass $\geq 0.85 M_\odot$ align well with

²⁴ Accessed through <http://stev.oapd.inaf.it/cgi-bin/cmd> using version 3.4. This particular version is available at http://stev.oapd.inaf.it/cgi-bin/cmd_3.4.

²⁵ <https://www.cosmos.esa.int/web/gaia/edr3-passbands>

²⁶ The total extinction is related to the colour excess through the relative visibility by the equation $A_V = R_V E(B - V)$, and a value of $R_V = 3.1$ is typical for the Milky Way.

6.6. Birthrate

the observed RGB sequence aside from the discrepancy between the location of the red-giant bump in the data and the corresponding feature in the models. The red-giant bump is the accumulation of stars located at $G \sim 14$ on the RGB of the data in Fig. 6.7 and corresponds to a temporary decrease in luminosity (and thus a “bump” in the luminosity function). The occurrence of this feature is related to the H-burning shell passing through the composition gradient left over by the convective envelope where it reached its maximum depth [245]. The location of the red-giant bump is thus sensitive to the details of mixing processes beyond convective boundaries, which current stellar evolution models struggle to accurately predict [246]. The discrepancy between observations and model predictions of the luminosity of the red-giant bump in globular clusters has been identified repeatedly in the literature, with a variety of different potential explanations suggested (e.g. [246–252]), and it remains an open issue. However, more accurately modelling the location of this feature is not needed for our purposes as long as the RGB boundary region is chosen appropriately.

The boundary region shown in Fig. 6.7a selects stars at the early stage of evolution along the RGB and excludes the red-giant bump feature of both the data and the models. This boundary region was chosen to maximise the number of stars selected while not needing to be concerned about the misalignment of the red-giant bump between the models and the data. For each model used (as described above), we calculated $N_{\text{RGB}} / t_{\text{RGB}}$ for this region as well as sub-regions where the maximum G value (at which the bottom horizontal boundary line in Fig. 6.7a was drawn) was reduced in increments of 0.1 from 16.5 to 16.0. This varies the boundary close to the subgiant branch, which stars evolve through more slowly than the RGB stage. The spread in values from these calculations was taken to be the 3σ range of the Gaussian prior for the birthrate (i.e. the range from $\dot{N}_0 - 3\sigma_{\dot{N}}$ to $\dot{N}_0 + 3\sigma_{\dot{N}}$), yielding a prior birthrate of $\dot{N}_0 = 2.21 \times 10^{-7} \text{ yr}^{-1}$ with $\sigma_{\dot{N}} = 0.04 \times 10^{-7} \text{ yr}^{-1}$.

As a check, we considered an extended RGB boundary region shown in Fig. 6.7b that includes the red-giant bump of both the data and the

6.6. Birthrate

models²⁷ and performed a similar set of calculations in which the boundary regions of both Fig. 6.7a and Fig. 6.7b were truncated at the maximum G end in increments of 0.25 from 16.5 to 16.0, with $N_{\text{RGB}} / t_{\text{RGB}}$ calculated in each of these cases for each relevant model. Considering the spread of values for just the extended regions gives a much larger birthrate estimate of $\dot{N} = (2.43 \pm 0.03) \times 10^{-7} \text{ yr}^{-1}$, but we do not consider this estimate reliable due to the discrepancy between the models and data in this region. Considering both the shorter and extended RGB boundary regions gives a wider spread of values and thus a larger value of $\sigma_{\dot{N}}$ in addition to a larger value of \dot{N}_0 compared to considering just the shorter regions, yielding an estimate of $\dot{N} = (2.31 \pm 0.07) \times 10^{-7} \text{ yr}^{-1}$.

This tendency to shift the estimate of the birthrate to larger values when the RGB boundary region is extended to include the red-giant bump is worth keeping in mind when considering the results of this chapter. However, including the estimates from the extended RGB regions likely overestimates the uncertainty in the birthrate. It is important to have a tight prior on the birthrate for the unbinned likelihood analysis, otherwise the birthrate will tend to simply be adjusted so that the number of white dwarfs predicted by the likelihood function matches the total number of white dwarfs observed in the data space, which effectively prioritises fitting the normalisation constant rather than the morphology of the cooling curve and is weighted towards the faintest end of the cooling sequence where there are the most white dwarfs but also the poorest completeness. We consider the birthrate estimate given by just the shorter RGB regions that exclude the red-giant bump to be the most accurate and thus use the result $\dot{N} = (2.21 \pm 0.04) \times 10^{-7} \text{ yr}^{-1}$ given by those calculations as the prior for the birthrate in the unbinned likelihood analysis of Section 6.7.

²⁷Using a boundary region that includes the red-giant bump in the data but excludes it in the models gives a larger birthrate estimate than regions that either include or exclude the red-giant bump in both the data and models, but that is not a reasonable selection.

6.7 Unbinned Likelihood Analysis

We perform an unbinned likelihood analysis similar to that of Goldsbury *et al.* [215]. As the data used in our analysis consists of observations far from the centre of 47 Tuc, where the density profile depends very little on the distance R from the cluster centre, we take the density profile to be uniform. The distribution of photometric errors is also approximately independent of R for our data, so the number density distribution function (and thus likelihood) does not depend on R for our analysis.

For the analysis performed in this chapter, the number density distribution function f for the magnitudes $m_1 = \text{F606W}$ and $m_2 = \text{F814W}$ is given by the expression

$$f(m_1, m_2; \theta) = \dot{N} f_{\text{CR}}(m_1) \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f_M(m'_1, m'_2; \theta_M) \times E(m_1 - m'_1, m_2 - m'_2; m'_1, m'_2) dm'_1 dm'_2, \quad (6.30)$$

where θ denotes the set of parameters that the full model depends on, θ_M is the set of all model parameters except the birthrate \dot{N} , E is the photometric error distribution function, and the quantity $\dot{N} f_M$ is the theoretical number density distribution function given by the model before accounting for photometric errors. The function f_M gives the rate of change of the cooling time with respect to magnitude and depends only on the subset of parameters θ_M . The completeness correction factor f_{CR} has been inserted into the expression for $f(m_1, m_2; \theta)$ given by Goldsbury *et al.* [215] to account for the additional cleaning procedure performed in our work. Note that f_{CR} is a function of the magnitude m_1 after accounting for photometric errors. The primed magnitudes m'_1 and m'_2 are the magnitude values before accounting for photometric errors, while the magnitudes m_1 and m_2 without a prime symbol are the magnitude values after accounting for photometric errors.

In the notation of Goldsbury *et al.* [215], the expression giving f_M would be written as $f_M = \frac{dt}{dm'_1 dm'_2}$ for a cooling time dt over a cell of volume $dm'_1 dm'_2$ in magnitude-magnitude space. Since m'_1 and m'_2 are dependent

6.7. Unbinned Likelihood Analysis

variables, the expression for f_M is written more formally as

$$f_M(m'_1, m'_2; \theta_M) = \frac{dt}{dm'_1} \delta [m'_2 - m_{2,\text{mod}}(m'_1; \theta_M)], \quad (6.31)$$

where $m_{2,\text{mod}}(m'_1; \theta_M)$ is the function that relates m'_2 to m'_1 for a particular cooling model (parameterised by θ_M). The set of parameters represented by θ_M consists of the parameters of the theoretical cooling models (i.e. M_{WD} , $\log_{10} q_H$, and diffusion scenario), as well as the distance d to 47 Tuc (from Earth) and the colour excess $E(B - V)$ due to reddening. The full set of parameters represented by θ additionally includes the white dwarf birthrate, such that $\theta = \{\dot{N}, \theta_M\}$.

Note that the different diffusion scenarios can equivalently be viewed as different models (in a model comparison problem) or as a single model of the modified diffusion parameterised by a diffusion parameter (in a parameter estimation problem), with the cases of standard diffusion and no diffusion being extremum cases of this model. The parameters d and $E(B - V)$ are needed to move the models from theory space to data space, as they are used to determine the magnitudes m'_1 and m'_2 predicted by a model (before accounting for photometric errors). Since the publication of Goldsbury *et al.* [215], Gaia observations have been used to determine the distance to 47 Tuc to much higher accuracy, $d = 4.45 \pm 0.13$ kpc [214], than was known at the time the work of Goldsbury *et al.* [215] was done, so there is less concern about its value in our work. Likewise, varying $E(B - V)$ over the very limited range of allowable values for 47 Tuc, $E(B - V) = 0.04 \pm 0.02$ [210], would have little effect on our results while increasing the computational cost of the analysis (if added as an additional axis of the parameter grid). We thus keep d and $E(B - V)$ fixed at values of $d = 4.45$ kpc [214] and $E(B - V) = 0.04$ [210] in our analysis. The predicted magnitudes are calculated from the relevant theory-space model variables using bolometric corrections and these d and $E(B - V)$ values.

The bolometric correction for a particular filter labelled by the index i is the difference between the bolometric magnitude and absolute magnitude

6.7. Unbinned Likelihood Analysis

for that filter,

$$\text{BC}_i = M_{\text{bol}} - M_i, \quad (6.32)$$

where M_{bol} is the bolometric magnitude and M_i is the absolute magnitude in the i th filter. For a fixed metallicity, BC_i is a function of the effective temperature T_{eff} and surface gravity g , which the white dwarf cooling models give as a function of cooling age. The bolometric magnitude is given by

$$M_{\text{bol}} = M_{\text{bol},\odot} - 2.5 \log \left(\frac{L}{L_{\odot}} \right), \quad (6.33)$$

where $M_{\text{bol},\odot}$ is the absolute bolometric magnitude for the nominal solar luminosity L_{\odot} and L is the luminosity of the white dwarf. According to the IAU 2015 resolution [69], the nominal solar luminosity is $L_{\odot} = 3.828 \times 10^{26}$ W, corresponding to $M_{\text{bol},\odot} = 4.74$ mag.

The apparent magnitude predicted by a model in the i th filter is thus

$$m'_i = M_{\text{bol}} - \text{BC}_i + \mu + A_i, \quad (6.34)$$

where μ is the distance modulus, given by

$$\mu = 5 \log_{10} \left(\frac{d}{10 \text{ pc}} \right), \quad (6.35)$$

and A_i is the extinction in the i th filter due to interstellar reddening, which depends on $E(B - V)$ through the extinction law. We take the distance to be $d = 4.45$ kpc [214], which gives a distance modulus of $\mu = 13.24$, and we take the colour excess to be $E(B - V) = 0.04$ [210]. The extinctions in the filters F606W and F814W were determined using the extinction law²⁸ of Cardelli *et al.* [38] and O'Donnell [39] with a total V -band extinction of $A_V = 0.124$ and relative visibility of $R_V = 3.1$, which is typical of the Milky Way. The relative visibility relates the total extinction to the colour excess through the expression $A_V = R_V E(B - V)$. The resultant values for the

²⁸The extinction values were retrieved using http://stev.oapd.inaf.it/cgi-bin/cmd_3.7, which uses this extinction law.

6.7. Unbinned Likelihood Analysis

extinctions A_i are

$$A_{\text{F606W}} = 0.59696 A_V, \quad (6.36)$$

$$A_{\text{F814W}} = 0.90328 A_V. \quad (6.37)$$

We use bolometric corrections calculated by the procedure described in Holberg and Bergeron [253], which is an extension of the earlier work of Bergeron *et al.* [254], for pure-hydrogen models and the relevant HST filters²⁹. These bolometric corrections for DA white dwarfs were calculated using the models of Blouin *et al.* [255] at low temperatures $T_{\text{eff}} < 5,000$ K, the models of Bédard *et al.* [109] at high temperatures $T_{\text{eff}} > 30,000$ K, and the models of Tremblay *et al.* [256] at intermediate temperatures. They also incorporate the Lyman alpha profile calculations of Kowalski and Saumon [257].

With the procedure established for calculating the distribution function given by Eq. (6.30), the unbinned likelihood can then be calculated from the distribution function. The unbinned likelihood \mathcal{L} is given by the expression³⁰

$$\ln \mathcal{L}(\theta) = \sum_i \ln f(m_{1i}, m_{2i}; \theta) - N_{\text{pred}}(\theta), \quad (6.38)$$

where $\ln \mathcal{L}$ is the natural logarithm of the likelihood, i is an index that enumerates the data points (i.e. the white dwarfs observed within the boundaries defining the data space), $f(m_{1i}, m_{2i}; \theta)$ is the number density distribution function evaluated at the magnitude values m_{1i} and m_{2i} of the i th data point, and N_{pred} is the total number of white dwarfs predicted by the model to be in the data space. N_{pred} is calculated by integrating $f(m_1, m_2; \theta)$ over

²⁹Tables with bolometric and absolute magnitudes for a number of photometric systems calculated using the same atmosphere models are available at <https://www.astro.umontreal.ca/~bergeron/CoolingModels/>. Absolute magnitudes in the HST filters relevant for this work were provided by Pierre Bergeron upon request.

³⁰As was done in Goldsbury *et al.* [215], an additive constant that does not depend on the model parameters has been dropped from the expression for $\ln \mathcal{L}$. See Goldsbury *et al.* [215] for a full derivation of this expression.

6.7. Unbinned Likelihood Analysis

the whole data space,

$$N_{\text{pred}}(\theta) = \iint_{\text{data space}} f(m_1, m_2; \theta) dm_1 dm_2. \quad (6.39)$$

The data space used for this analysis is a truncated version of the white dwarf data space shown in Fig. 6.1 that ends at a horizontal lower limit of $F606W = 28.5$ instead of extending to the limit of $F606W = 29.0$ used in the cleaning procedure. This truncated data space region is used to ensure that the completeness (as a function of $F606W$ magnitude) remains reasonable even after applying the completeness correction factor. To extend the analysis to larger magnitudes, there is a trade-off between the improved statistical power of having more objects vs the reduced completeness and increased spread of the error distribution at larger magnitudes. Multiple cut-offs for the data space between $F606W$ of 28.0 and 29.0 were tested, and it was found that a cut-off of 28.5 optimised this trade-off while still extending to a large enough magnitude to capture the relevant feature in the cumulative luminosity function associated with convective coupling.

The data space ultimately used in the unbinned likelihood analysis is shown in Fig. 6.8. The data space boundaries are indicated by the solid red curves enclosing most of the 47 Tuc white dwarf cooling sequence. For reference, a cooling model as a function of input magnitudes before accounting for the distribution of photometric errors is shown as a solid orange curve that passes through the middle of the data space. The data shown in Fig. 6.8 are the HST data after applying the full data cleaning procedure described in Section 6.4. It can be seen in Fig. 6.8 that the reference model aligns well with the observed white dwarf cooling sequence of 47 Tuc, falling approximately along the centre of this sequence. For the model parameter ranges considered in this work, there is little variation in the corresponding model curve in the CMD, so Fig. 6.8 looks similar for every model in our grid of cooling models. Note the data space boundaries are shown in colour-magnitude space in Fig. 6.8 for the ease of visualisation; however, the analysis was actually performed in magnitude-magnitude space.

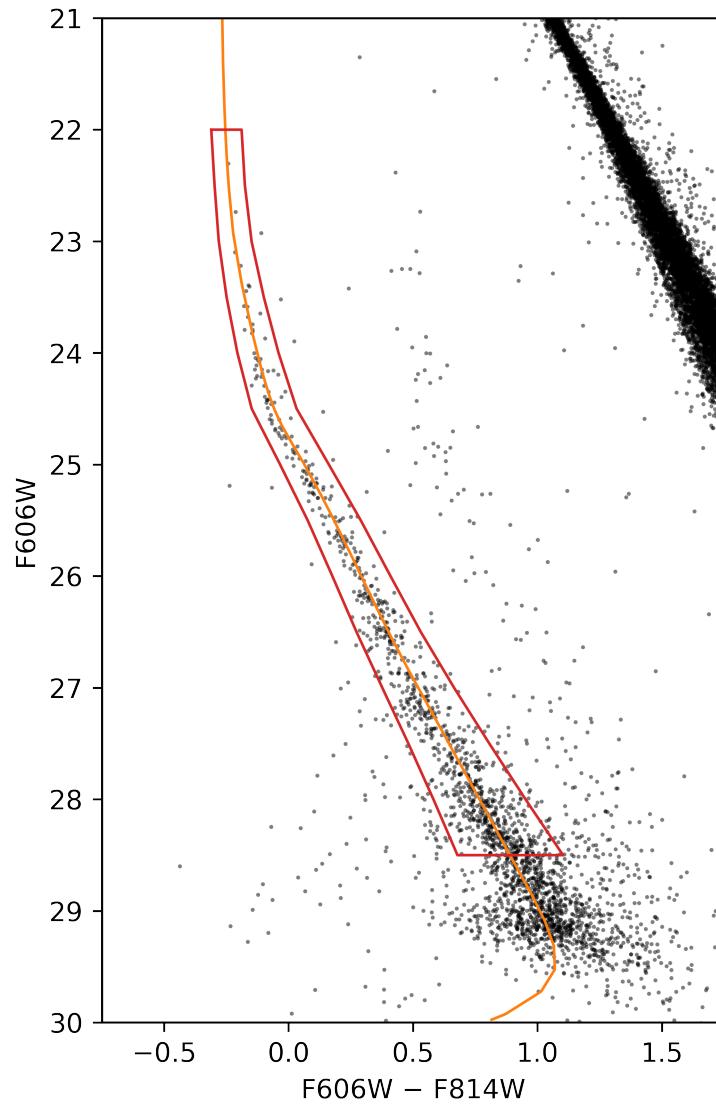


Figure 6.8: Data space used in unbinned likelihood analysis. The boundaries of the data space are indicated by the solid red curves. A reference cooling model in terms of input magnitudes (before accounting for photometric errors) is shown as the solid orange curve passing through the data space. The data space boundaries and model are shown overlaid on the cleaned HST data (points).

6.7. Unbinned Likelihood Analysis

To maximise the log-likelihood, we first evaluate the likelihood over a grid of values for the 3 parameters M_{WD} , diffusion, and q_H with the birthrate fixed to the value \dot{N}_0 calculated from stars leaving the main sequence. We then analytically solve for the birthrate that optimises the likelihood at each point on the 3-parameter grid and find the combination of the other parameters that give the maximum value of the likelihood after re-scaling the birthrate. The birthrate is treated in this way instead of as an additional axis of the parameter grid in order to reduce the memory usage of the procedure.

Let \dot{N}_0 be the fixed birthrate used to calculate the likelihoods on the 3-parameter grid described above, and let $\ln \mathcal{L}_0$ and $N_{\text{pred},0}$ be the corresponding (natural) log-likelihood and total predicted number of white dwarfs in the data space. The likelihood for this fixed birthrate is given by

$$\ln \mathcal{L}_0 = \sum_{i=1}^{N_{\text{obs}}} \ln f_{0,i} - N_{\text{pred},0}, \quad (6.40)$$

where $f_{0,i}$ is the distribution function evaluated at the magnitude values of the i th data point with a birthrate of \dot{N}_0 . Let a be a factor that re-scales the birthrate such that

$$\dot{N} = a \dot{N}_0. \quad (6.41)$$

We assign a Gaussian prior to the birthrate such that $\dot{N} \sim \mathcal{N}(\dot{N}_0, \sigma_{\dot{N}})$ and account for this prior in the expression for $\ln \mathcal{L}$, resulting in the expression

$$\ln \mathcal{L} = \ln \mathcal{L}_0 + N_{\text{obs}} \ln(a) - N_{\text{pred},0} (a - 1) - \frac{1}{2} \left(\frac{\dot{N}_0}{\sigma_{\dot{N}}} \right)^2 (a - 1)^2. \quad (6.42)$$

The values of \dot{N}_0 and $\sigma_{\dot{N}}$ that parameterise the prior distribution for \dot{N} were determined using the procedure described in Section 6.6. Note that if the priors for the other parameters represented by θ (aside from the birthrate) are taken to be uniform, then the logarithm of the joint posterior probability distribution for the parameters is equal to $\ln \mathcal{L}$ as given by Eq. (6.42) up to an additive constant (which includes the uniform prior distributions and an overall normalisation term).

6.7. Unbinned Likelihood Analysis

The expression for $\ln \mathcal{L}$ given by Eq. (6.42) can be optimised analytically with respect to the re-scale factor a without needing to evaluate the distribution function Eq. (6.30) for different birthrate values, which reduces the computational cost of the analysis. The extremum values for $\ln \mathcal{L}$ with respect to the birthrate re-scale factor occur when

$$a = \frac{1}{2} (1 - b) \pm \frac{1}{2} \sqrt{(1 - b)^2 + 4b \frac{N_{\text{obs}}}{N_{\text{pred},0}}}, \quad (6.43)$$

where the constant b has been defined to be

$$b \equiv N_{\text{pred},0} \left(\frac{\sigma_{\dot{N}}}{\dot{N}_0} \right)^2. \quad (6.44)$$

To find the combination of parameters that maximises the likelihood, i.e. the maximum likelihood estimates of the parameters, we first calculate $\ln \mathcal{L}_0$ at each point on the 3-parameter grid as given by Eq. (6.40). At each point on this parameter grid, we then calculate a as given by Eq. (6.43) and use the result to calculate $\ln \mathcal{L}$ as given by Eq. (6.42). Finally, we find the combination of parameters on the parameter grid that maximises this new $\ln \mathcal{L}$, and we get the corresponding maximum likelihood estimate of the birthrate from Eq. (6.41) using the optimal value of the re-scale factor for that parameter grid point.

The results of the unbinned likelihood analysis are presented in Section 6.8, with the results of the procedure to find the optimal parameters using the likelihood re-scaling technique described above presented in Section 6.8.1. The optimal model found by this procedure is compared to the data in Section 6.8.2 by comparing the predicted and empirical (inverse) cumulative luminosity functions.

6.8 Results

6.8.1 Likelihood Distribution

The distribution of likelihood values after locally optimising the birthrate at each point on the cooling model parameter grid is shown in Fig. 6.9. The quantity \mathcal{L} plotted in Fig. 6.9 is really the likelihood plus birthrate prior after re-scaling the birthrate, i.e. \mathcal{L} as given by Eq. (6.42), evaluated at each point on the parameter grid using the value of the birthrate re-scale factor that maximises \mathcal{L} for that combination of cooling model parameter values. To make the significance of the likelihood values clear, \mathcal{L} has been scaled by the value $\hat{\mathcal{L}}$ of its global maximum across the entire parameter grid, and the distribution of $\mathcal{L}/\hat{\mathcal{L}}$ values is shown as a filled contour plot with the levels corresponding to the σ -level values of the similarly scaled probability density for a two-dimensional (spherically symmetric) normal distribution. Letting p denote the probability density of a two-dimensional normal distribution and \hat{p} denote the maximum value of p , the value of the scaled probability density corresponding to the level $n_\sigma \sigma$ is simply given by $\hat{p}/p = \exp(-0.5 n_\sigma^2)$. More specifically, the contour levels corresponding to 1, 2, 3, 4, and 5 σ in Fig. 6.9 are drawn at values of $-0.5, -2.0, -4.5, -8.0$, and -12.5 (from darkest to lightest region).

Each plot in Fig. 6.9 corresponds to a slice in the distribution for a particular diffusion scenario: no diffusion (top), modified diffusion (middle), and standard diffusion (bottom). The likelihoods in these three plots are all scaled by the same value $\hat{\mathcal{L}}$, which is the global maximum across the different diffusion scenarios (as well as with respect to M_{WD} and $\log_{10} q_H$); this global normalisation makes these plots a meaningful comparison of the different diffusion scenarios. From Fig. 6.9, it can be seen that the cases of modified diffusion and standard diffusion give similarly likely results for similar M_{WD} and $\log_{10} q_H$ values, though with slightly smaller $\log_{10} q_H$ values favoured in the case of modified diffusion. The 2 σ level contains both the case of standard diffusion with $M_{WD} = 0.5314 M_\odot$ and $\log_{10} q_H$ between -3.65 and -3.55 and the case of modified diffusion with the same mass value $M_{WD} = 0.5314 M_\odot$ and $\log_{10} q_H$ between -3.75 and -3.60 . Even the 1 σ

6.8. Results

level contains both the case of standard diffusion (with $M_{\text{WD}} = 0.5314 M_{\odot}$ and $\log_{10} q_H$ between -3.60 and -3.55) and the case of modified diffusion (with $M_{\text{WD}} = 0.5314 M_{\odot}$ and $\log_{10} q_H = -3.65$). Similar to the case of modified diffusion, the case of no diffusion also favours slightly lower $\log_{10} q_H$ values than the case of standard diffusion, though the case of no diffusion is overall less likely than modified or standard diffusion, with the case of no diffusion not included within the 2σ level for any combination of M_{WD} or $\log_{10} q_H$ values.

It should be emphasised that the filled contours in Fig. 6.9 indicate ranges of probability density values, not regions of enclosed probabilities, so they are not credible regions. The distribution shown in Fig. 6.9 is also not quite the joint posterior distribution distribution; rather, it is a distribution of the local maximum of the posterior with respect to birthrate at each point on the parameter grid. Since the prior for the birthrate is a narrow Gaussian, this is expected to be very similar to the posterior distribution after marginalising over the birthrate. This marginal posterior is shown in Fig. 6.10, which confirms the expected similarly of these two distributions. Fig. 6.10 is analogous to Fig. 6.9, with contours drawn at the same levels, but for the marginal posterior density distribution p instead of \mathcal{L} . Like Fig. 6.9, the filled contours in Fig. 6.10 are also not credible regions, though the proper credible regions would be expected to be similar for a posterior distribution that is approximately a normal distribution in the remaining parameters.

Though the quantity plotted in Fig. 6.9 is not quite the posterior distribution, the location of its maximum value on the parameter grid is the same as the location at which the posterior probability density distribution is maximised (and the global maximum $\hat{\mathcal{L}}$ is equal to the maximum value of the full joint posterior distribution). We find that $\hat{\mathcal{L}}$ corresponds to the case of standard (fully on) diffusion with parameter values of $M_{\text{WD}} = 0.5314 M_{\odot}$ and $\log_{10} q_H = -3.55$ (corresponding to $q_H = 2.82 \times 10^{-4}$) on the parameter grid and a birthrate of $\dot{N} = 2.7 \times 10^{-7} \text{ yrs}^{-1}$ after re-scaling. The corresponding value of the log-likelihood is $\ln \hat{\mathcal{L}} = 5385.35$. These maximum likelihood estimate results are summarised in Table 6.5.

6.8. Results

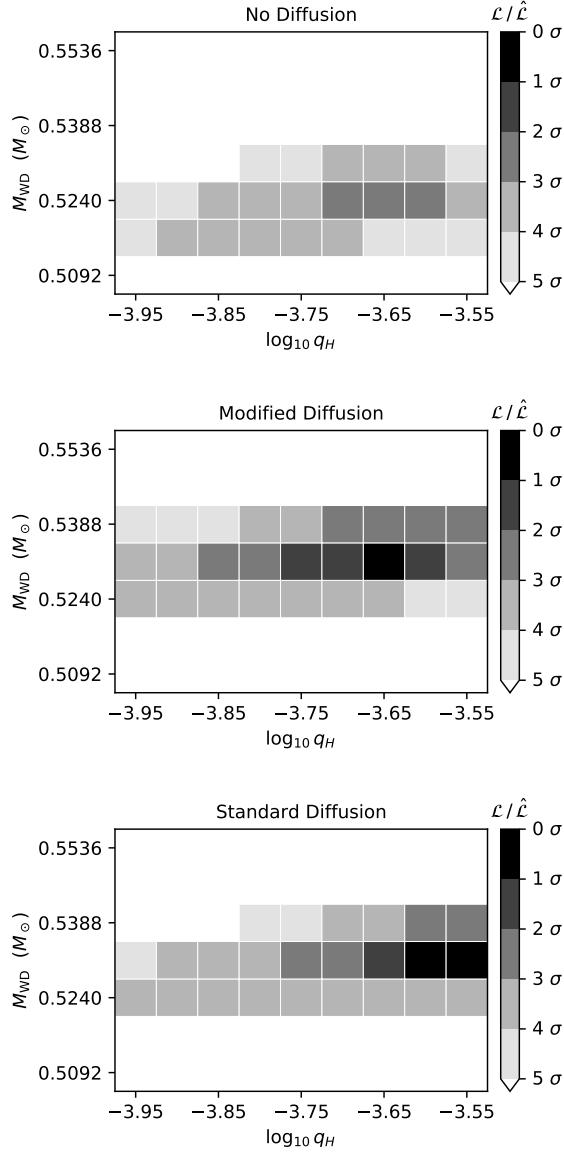


Figure 6.9: Likelihood (including birthrate prior) locally maximised with respect to birthrate at each location on the parameter grid. The likelihood has been scaled by its global maximum across the parameter grid (including different diffusion scenarios), with contours drawn at the values of the analogously scaled probability density of a two-dimensional normal distribution evaluated at 1, 2, 3, 4, and 5 σ (darkest to lightest).

6.8. Results

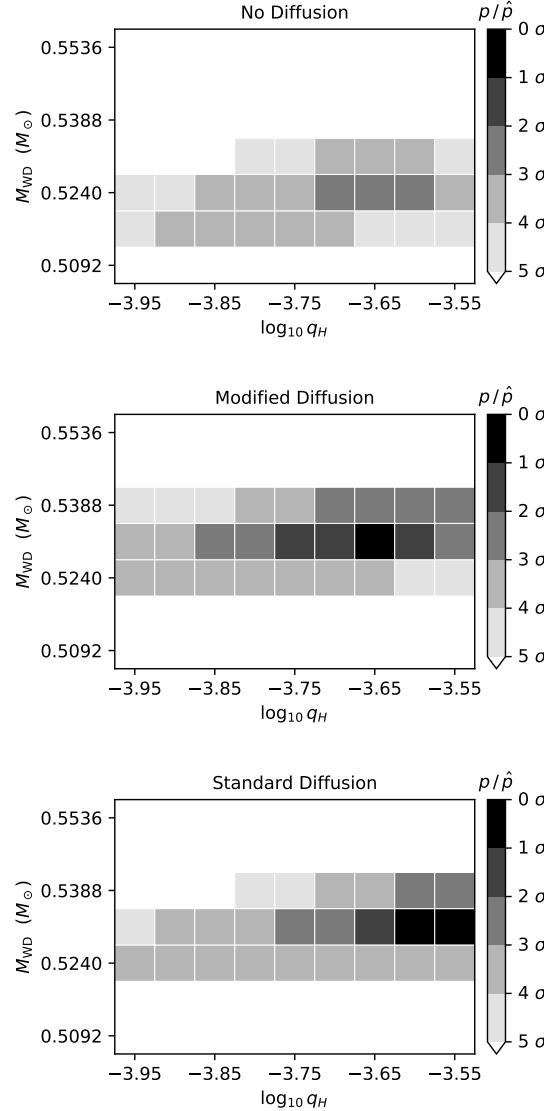


Figure 6.10: Joint posterior probability density distribution after marginalising over the birthrates. The distribution as a function of white dwarf mass (M_{WD}) and envelope thickness (q_H) is shown for each diffusion scenario. The probability density (p) has been scaled by its maximum value (\hat{p}) across all diffusion scenarios so that the plotted quantity is p/\hat{p} . The filled contours are drawn at level values corresponding to the σ levels for a two-dimensional normal distribution.

Parameter	Value
Diffusion scenario	Standard diffusion
M_{WD}	$0.5314 M_{\odot}$
$\log_{10} q_H$	-3.55
N	$2.27 \times 10^{-7} \text{ yrs}^{-1}$
$\ln \hat{\mathcal{L}}$	5385.35

Table 6.5: Maximum likelihood estimates of parameters and corresponding log-likelihood value from unbinned likelihood analysis. These are the parameter values and likelihood of the best-fitting model.

6.8.2 Best-Fitting Model

The inverse cumulative luminosity function, i.e. the inverse of the cumulative number distribution of white dwarfs in the data space as a function of magnitude, is shown in Figs. 6.11 and 6.12 for the best-fitting model from the unbinned likelihood analysis (red curve) in comparison to the empirical distribution for the HST data (black points) that the models were fitted to. The cumulative number as a function of magnitude for the model is given by integrating the number density distribution function,

$$N(m_1; \theta) = \int_{-\infty}^{m_1} dm_1^* \int_{-\infty}^{\infty} dm_2 f(m_1^*, m_2; \theta) \quad (6.45)$$

where N is the (predicted) number of white dwarfs in the data space with magnitude $\leq m_1$ in filter 1 (in this case F606W) and $f(m_1, m_2; \theta)$ is the number density distribution function given by Eq. (6.30) evaluated over the (m_1, m_2) coordinates of the data space and equal to zero outside of the data space. An equivalent expression for the cumulative number distribution as a function of m_2 can be written by swapping the indices 1 and 2 in Eq. (6.45).

Note that the cumulative number distribution is proportional to the white dwarf birthrate through the proportional dependence of $f(m_1, m_2)$ on the birthrate; a larger birthrate would shift the whole model distribution to the right in Figs. 6.11 and 6.12, while a smaller birthrate would shift it to the left. The other model parameters that were varied in the analysis affect

6.8. Results

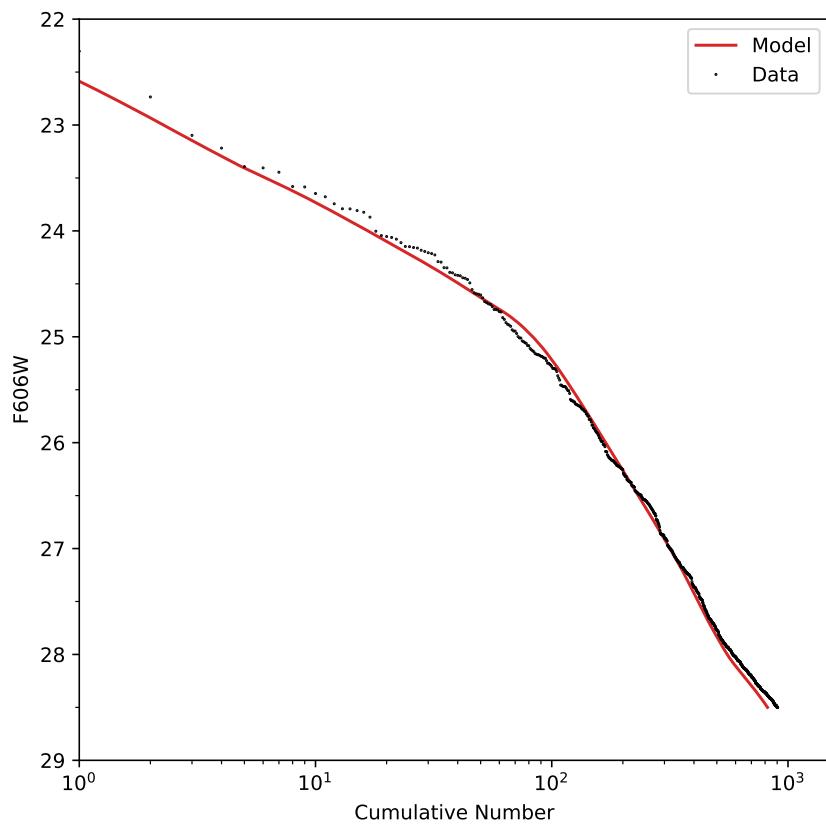


Figure 6.11: Inverse cumulative luminosity function for F606W magnitude of data (black points) compared to optimal model determined by the unbinned likelihood analysis (red curve).

6.8. Results

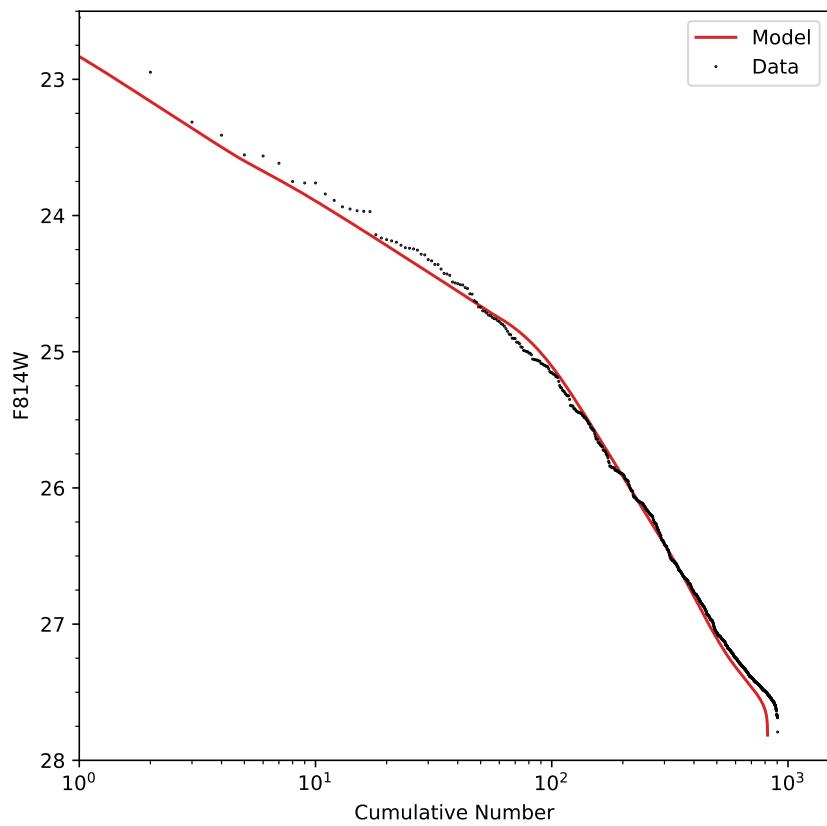


Figure 6.12: Inverse cumulative luminosity function for F814W magnitude of data (black points) compared to optimal model determined by the unbinned likelihood analysis (red curve).

6.8. Results

Magnitude Variable	<i>p</i> -value
F606W	0.077
F814W	0.041

Table 6.6: Results of KS tests comparing the one-dimensional marginal cumulative probability distribution functions predicted by the optimal model to the corresponding empirical distribution.

the morphology of the distribution. The cumulative number distribution for the data is given by simply counting a list of the data points ordered by magnitude. The incompleteness of the data is accounted for in the distribution functions of the models, so it does not need to be corrected for in plotting the distribution of the data in Figs. 6.11 and 6.12.

It can be seen from Figs. 6.11 and 6.12 that the optimal model from the unbinned likelihood analysis well reproduces the empirical distribution for most of the white dwarf data space, which is indicative of a good fit to the data. As a simple quantitative check of the goodness-of-fit of the optimal model to the data, we performed a one-sample Kolmogorov-Smirnov (KS) test for the cumulative marginal probability distribution with respect to each magnitude (after integrating over the other magnitude). Each KS test compares the cumulative probability density function predicted by the model for a given magnitude to the empirical cumulative fraction as a function of that magnitude. The *p*-value returned by the one-sample KS test gives the probability of drawing a sample from the model distribution for which the corresponding sample distribution differs from the model distribution by at least as much as the observed empirical distribution differs from the model distribution. Very small *p*-values indicate a low probability that the data was drawn from the model distribution and thus a poor fit, while larger *p*-values are expected if the data was drawn from the model distribution (i.e. if the model distribution and empirical distribution are equivalent) and thus indicate a good fit. The *p*-values from these KS tests are given in Table 6.6. The large *p*-values (well above a reasonable threshold of 10^{-4}) in both cases indicate a good fit.

6.9. Discussion

Note that the distributions compared by each KS test are cumulative fractions normalized to unity over the magnitude range, rather than the cumulative number distributions as shown in Figs. 6.11 and 6.12. These KS tests specifically assess the morphology of the (inverse) cumulative luminosity function with respect to a particular magnitude for the optimal combination of cooling model parameters (MWD , q_H) and diffusion scenario. The birthrate parameter simply re-scales the total number of white dwarfs predicted by the model (i.e. the normalization of the number density distribution), but the KS test is applied to the probability distribution function (normalized to unity). Thus, while the birthrate was important in determining the optimal cooling model using the unbinned likelihood analysis, it does not directly affect the morphology assessed by the KS test. Nonetheless, we note that the birthrate found for the optimal model only deviates from the prior value by

$$\dot{N} - \dot{N}_0 = 1.5 \sigma_{\dot{N}}, \quad (6.46)$$

suggesting the birthrate value is reasonable.

6.9 Discussion

Previous work studying the cooling of white dwarfs in 47 Tuc such as Goldsbury *et al.* [215] and Obertas *et al.* [30] used cooling models based on `test_suite` examples from older versions of MESA that had diffusion turned off for white dwarf evolution. Obertas *et al.* [30] used the same data as our work and considered old white dwarfs down to faint enough magnitudes in the cooling sequence to reveal the effect of core crystallisation in the (inverse) cumulative luminosity function, but the models shown in that work were not fit to the data. Goldsbury *et al.* [215] found models that fit the data well, but compared the models to data of younger white dwarfs than considered in this chapter. Since the focus of Goldsbury *et al.* [215] was neutrino cooling, which is important at early white dwarf cooling times (and negligible at late cooling times), the data space selections used in that

6.9. Discussion

work stopped at much brighter magnitudes³¹ than the data space used in our analysis. Since we are interested in the thickness of the H envelope as a key parameter in our work (rather than it being a nuisance parameter as it was in the work of Goldsbury *et al.* [215]) and the effect of this parameter is best analysed using very old white dwarfs, our data space needs to extend to fainter magnitudes. Truncating the data space used in our analysis so that it stops at brighter magnitudes (maximum F606W ≤ 28.25) results in the case of diffusion fully turned off giving a better fit, though the optimal q_H tends towards lower values in that case.

In this work, the parameters upon which the modified treatment of diffusion depend have been fixed at fiducial values, though the code to implement modified diffusion has been written with the flexibility to vary the values of these diffusion parameters. The theoretically ideal approach to comparing the different diffusion scenarios in a Bayesian framework would be to compare the diffusion models after marginalising over the diffusion parameters in the case of modified diffusion. In principle, these parameters could be varied to produce white dwarf cooling models over a grid of diffusion parameter values to extend the analysis done in this work. This would be worth doing for data that has greater sensitivity to the treatment of diffusion, which could potentially be produced by e.g. future JWST observations. However, our results indicate that it is already difficult to differentiate between standard diffusion and our fiducial modified diffusion (when the H envelope thickness is a free parameter) using the data considered in this work, so extending our analysis in this way for this data is unlikely to provide additional information.

³¹ Goldsbury *et al.* [215] used HST data in different filters than our data, so the magnitude values are not directly comparable, but they can be compared through the corresponding theoretical luminosity values.

6.10 Conclusions

In this chapter, we performed a detailed analysis of the cooling of white dwarfs to late cooling times in the globular cluster 47 Tuc using archival data of deep observations taken by HST. These deep photometric observations resolve the white dwarf cooling sequence of 47 Tuc to faint enough magnitudes that a bump associated with the onset of convective coupling and core crystallisation can be seen in the luminosity function. This was shown by Obertas *et al.* [30] using this same data, though a statistical analysis of model fits was not performed in that work. We built upon the work of Obertas *et al.* [30] by performing a detailed statistical analysis that accounted for different treatments of diffusion, H envelope thickness, white dwarf mass, and white dwarf birthrate using the unbinned likelihood.

A cleaning procedure consisting of cuts in proper motion and the photometry quality-of-fit parameter **SHARP** were performed to remove contamination from the SMC in the 47 Tuc white dwarf cooling sequence, and this cleaning procedure was carefully calibrated to account for any potential residual contamination and the reduction of completeness due to cleaning. This cleaning procedure and calibration was particularly important for using data at the very faint end of the cooling sequence where the phenomena of convective coupling of the envelope to the core and the crystallisation of the core begin to occur. The thickness of the bump in the luminosity function due to convective coupling in particular is sensitive to the H envelope thickness, and extending the analysis to the onset of core crystallisation provides a good test of the regime over which the standard MESA treatment of diffusion can be used. The inclusion of these very old, faint white dwarfs in the analysis is thus important to distinguish (cooling) models of different H envelope thickness and diffusion treatment. Though the data for these older, fainter white dwarfs has lower completeness than younger, brighter white dwarfs, there is a much larger number of the fainter white dwarfs in the data, so including them in the analysis furthermore provides improved statistical power.

The stellar evolution software MESA was used to produce a suite of

6.10. Conclusions

white dwarf cooling models for different treatments of diffusion, H envelope thicknesses, and white dwarf mass. Three different diffusion scenarios were considered: i) the standard MESA treatment of diffusion where the ions are approximated as an ideal gas, ii) a custom modified treatment of diffusion that accounts for non-ideal gas effects (that notably suppress diffusion at the H/He boundary compared to the ideal gas case), and iii) the case of no diffusion. In addition to the cooling model parameters, the full model for the number density distribution function used in the unbinned likelihood analysis was also sensitive to the white dwarf birthrate, for which a prior value was determined using Gaia EDR3 data of stars on the RGB. The optimal model found by the unbinned likelihood analysis corresponded to the case of standard diffusion with a H envelope thickness parameter of $\log_{10} q_H = -3.55$ (corresponding to $q_H = 2.82 \times 10^{-4}$), a white dwarf mass of $M_{\text{WD}} = 0.5314 M_{\odot}$, and a white dwarf birthrate of $\dot{N} = 2.27 \times 10^{-7} \text{ yrs}^{-1}$.

We find that the standard MESA treatment of diffusion, in which the ions are approximated as an ideal gas, produces white dwarf cooling models that well reproduce the cumulative white dwarf luminosity function to magnitudes faint enough to resolve features related to the onset of convective coupling and core crystallisation, with thicker H envelopes being favoured. There is some degeneracy between the treatment of diffusion and the H envelope thickness; the modified treatment of diffusion with somewhat thinner (though still generally considered thick) envelopes produced cooling models that were similarly likely to the best-fitting models produced using the standard MESA treatment of diffusion with thicker envelopes, and these scenarios could not be distinguished at a statistically significant level using the data considered in this work.

The analysis of the deep HST ACS/WFC data considered in this work is limited by the increasingly poor completeness with increasing magnitude, particularly at the faint end of the cooling sequence that is most important for our analysis. Improved data with better completeness at the faintest magnitudes of the cooling sequence, as could for example be obtained with a newer telescope like JWST, may enable the degeneracy between the detailed treatment of diffusion and the H envelope thickness to be better disentangled

6.10. Conclusions

in the future. In that case, a more thorough, extensive analysis of how the treatment of diffusion should be modified to account for non-ideal gas effects may be warranted and would be an interesting avenue for future research. The analysis presented in this chapter, however, has already pushed the limits of what can be learned about this from the data considered in this work.

Chapter 7

Axion constraints from white dwarfs in 47 Tucanae

7.1 Introduction

White dwarfs are a popular target of indirect searches for axions and axion-like particles. For axion models such as the benchmark DFSZ model [81, 82] that include a coupling between the axion and electrons, axions can be produced in the very dense, electron-degenerate interior of white dwarfs through axion bremsstrahlung from electrons. Axions can also in principle be produced in the interiors of white dwarfs through the axion Primakoff effect, but this is a sub-dominant process compared to axion bremsstrahlung for the density and temperature scales typical of white dwarfs. Furthermore, the typical densities and temperatures of young white dwarfs would lead to high rates of emission for axions produced via axion bremsstrahlung relative to the emission of neutrinos and photons, which are the standard channels through which white dwarfs lose energy over time.

For example, X-ray observations of a hot, magnetic white dwarf have been used to search for the spectroscopic signal of axions that were produced in the core of the white dwarf through axion bremsstrahlung (from electron-ion scattering) and then converted to photons in the strong magnetic field surrounding the white dwarf [53, 54]. As that scenario relies on the axion-photon interaction (to produce the observable photon signal) in addition to an axion-electron interaction (to produce the axions), that work constrained the product of the axion-electron coupling g_{aee} and the axion-photon coupling $g_{a\gamma\gamma}$, rather than just g_{aee} alone. Dessert *et al.* [54] found

7.1. Introduction

a bound of $|g_{aee}g_{a\gamma\gamma}| \lesssim 1.3 \times 10^{-25}$ GeV $^{-1}$ at 95% confidence for low axion masses $m_a \ll 10^{-5}$ eV.

More typically, axion searches targeting white dwarfs focus on how the additional energy loss due to axion emission modifies the cooling behaviour of white dwarfs, and in particular the observable consequences of the altered cooling rate on the white dwarf luminosity function of a population of white dwarfs. Due to the weakness of the interaction between axions and standard model particles, axions that are produced in the interior of a white dwarf typically leave the white dwarf without further interaction. The emission of axions provides an extra source of energy loss for a white dwarf in addition to the standard energy loss due to the radiation of photons from the white dwarf surface and, for young white dwarfs, the emission of neutrinos from the white dwarf interior. This causes the white dwarf to cool at a faster rate than predicted by standard cooling processes alone.

Evidence for an additional white dwarf cooling source compatible with axion emission has been reported for the empirical white dwarf luminosity functions [44–46] of white dwarfs in the Galactic thin and thick discs, as well as the halo, and it has been suggested that this hints at the existence of an axion that couples to electrons, with a mass m_a (and angular parameter β) value of $m_a \sin^2 \beta \sim 4 - 10$ meV for a DFSZ axion model [49]. This corresponds to an axion-electron coupling of $g_{aee} \sim (1.1 - 2.8) \times 10^{-13}$. The work of Isern *et al.* [49] expanded upon the earlier work of Isern *et al.* [47] and Miller Bertolami *et al.* [48], which likewise presented evidence for similar hints of axions based on the shape of some measurements of Galactic white dwarf luminosity functions. Using Galactic white dwarf luminosity function measurements, Miller Bertolami *et al.* [48] placed a bound of $m_a \sin^2 \beta \lesssim 10$ meV on the mass of a DFSZ axion at a 3σ confidence level, which corresponds to a 3σ bound on the axion-electron coupling of $g_{aee} \lesssim 2.8 \times 10^{-13}$.

Recent reviews summarising both hints of axions and bounds on axion parameters from various observations pertaining to stellar evolution are given by Di Luzio *et al.* [84] and Di Luzio *et al.* [85]. These reviews compare bounds reported at a 2σ confidence level, where the 2σ bound from Miller

7.1. Introduction

Bertolami *et al.* [48] is given as $g_{aee} \leq 2.1 \times 10^{-13}$. The bound on g_{aee} has been improved upon since the work of Miller Bertolami *et al.* [48] (and Isern *et al.* [49]) using observations of the tip of the red giant branch in globular clusters [92, 93], with the strongest 2σ bound being $g_{aee} \leq 1.5 \times 10^{-13}$ [92].

Globular clusters are good environments for studying stellar evolution because they provide coeval populations of stars with well-controlled values of parameters like distance, reddening, and birthrate. The globular cluster 47 Tuc in particular is a good environment for studying the cooling of white dwarfs, and parameters important for modelling the cooling of white dwarfs in 47 Tuc, such as the typical white dwarf mass and envelope thickness, were studied extensively in Chapter 6. If axions exist and couple to electrons with the interaction strength hinted at by Galactic white dwarf luminosity functions, there should be a measurable effect on the cooling of white dwarfs in 47 Tuc as well. The emission of axions affects the white dwarf cooling rate in a manner similar to the emission of neutrinos at early cooling times, and the effect of neutrino emission on the cooling of young white dwarfs in 47 Tuc was studied in detail by Goldsbury *et al.* [215]. Like in the case of neutrino emission, the emission of axions produced in the dense interior of a white dwarf provides an additional mechanism of energy loss that increases the cooling rate compared to what is expected from photon radiation alone, though the effect of axion emission could persist to later cooling ages than neutrino emission due to the particular temperature and density values at which these emission processes are optimised [65, 66].

In this chapter, a detailed analysis of the cooling of white dwarfs in 47 Tuc is performed in order to look for indirect evidence of axions through the effect of axion emission on the white dwarf cooling rate. This work follows a procedure similar to the procedure of Goldsbury *et al.* [215], but applied to axions instead of neutrinos. We use the same data and many of the same analysis techniques as described in Goldsbury *et al.* [215], but we use new white dwarf cooling models that account for energy loss due to the emission of axions. The analysis procedure that we use in this work is also similar to that of Chapter 6, where a separate set of data was analysed to study the cooling of older white dwarfs in 47 Tuc. In the current work, we additionally

make use of prior knowledge gained from the work of Chapter 6 about the typical mass and envelope thickness of white dwarfs in 47 Tuc, as well as the appropriate treatment of diffusion in modelling white dwarf cooling, in our analysis of the effect of axion emission on the cooling of younger white dwarfs in 47 Tuc. The use of this prior information enables stronger constraints to be placed on the relevant axion parameters.

7.2 Data

In this chapter, we use archival data from HST observations of 47 Tuc performed as part of the HST Cycle 20 proposal GO-12971 (PI: H. Richer). The data was collected over 10 orbits, one of which was rejected due to loss of guide stars. In each orbit, observations were done simultaneously using WFC3/UVIS and ACS/WFC, with the WFC3/UVIS exposures split between the F225W and F336W filters and the ACS/WFC exposures split between the F435W and F555W filters. Across all of the orbits, WFC3 observed the inner field at the centre of the cluster while ACS simultaneously observed the outer field in a ring about the WFC3 field of view. This data is described in more detail in Goldsbury *et al.* [215], along with a detailed description of artificial stars tests that were performed to determine the photometric errors of this data.

7.3 Models

7.3.1 MESA Simulations

Following the procedure of Chapter 6, we performed white dwarf cooling simulations using MESA to produce a suite of white dwarf cooling models over a grid of model parameter values. The simulations are analogous to those described in Section 6.5.2, except they additionally implemented energy loss due to the emission of axions, which introduces axion couplings as additional parameters that the models depend on. The parameters varied in these simulations were the white dwarf mass M_{WD} , the thickness of the

H envelope, and the axion coupling constants.

The standard MESA treatment of diffusion was used for all of the simulations in this chapter. This was the diffusion scenario that produced the optimal cooling model in Chapter 6. In principle, different diffusion scenarios could be considered in this chapter as was done in Chapter 6, but the results of Chapter 6 do not indicate a need to do so. Furthermore, as the extreme ends of the mass grid considered in Chapter 6 are strongly excluded by the results of that chapter, we restrict our work in this chapter to white dwarf masses spanning the smaller range $0.5240 M_{\odot} \leq M_{\text{WD}} \leq 0.5388 M_{\odot}$. Within this white dwarf mass range, we performed simulations for the same grid of M_{WD} values considered in Chapter 6 (i.e. M_{WD} values of 0.5240, 0.5314, and $0.5388 M_{\odot}$), starting from the same initial models as the simulations described in Section 6.5.

In our MESA simulations of white dwarf cooling, we implement energy loss due to axion emission via the `run_star_extras` module in the `other_neu` subroutine. We account for energy loss due to both axion (electron) bremsstrahlung, which depends on the axion-electron coupling g_{aee} , and the axion Primakoff effect, which depends on the axion-photon coupling $g_{a\gamma\gamma}$. We implement axion bremsstrahlung emission using the prescription given by Nakagawa *et al.* [65, 66]. To account for axion Primakoff emission, we use the fitting formula given by Friedland *et al.* [87], which has previously been implemented (and for which code is publicly available) in the MESA `test_suite` example `axion_cooling`. We specify the value of each coupling constant, g_{aee} and $g_{a\gamma\gamma}$, through MESA inlist `x_ctrl` parameters. Note that axion bremsstrahlung is the dominant axion emission mechanism affecting white dwarf cooling. Though axion Primakoff emission is a sub-dominant effect for white dwarf cooling, it was included in the simulations for completeness and to confirm that this effect is negligible enough to neglect any potential dependence of the cooling models on $g_{a\gamma\gamma}$ in our analysis.

We performed simulations for a set of g_{aee} and $g_{a\gamma\gamma}$ values chosen based on the benchmark DFSZ (type I) model. For this model, the coupling constants g_{aee} and $g_{a\gamma\gamma}$ are related to the axion mass m_a through the expres-

sions³² [51]

$$g_{a\gamma\gamma} = 1.44 \times 10^{-10} \left(\frac{m_a}{1 \text{ eV}} \right) \text{ GeV}^{-1}, \quad (7.1)$$

$$g_{aee} = 2.83 \times 10^{-11} \sin^2 \beta \left(\frac{m_a}{1 \text{ eV}} \right), \quad (7.2)$$

where the expression for g_{aee} also depends on the angular parameter β . We performed simulations for a grid of m_a values ranging from 0.0 to 8.0 meV in increments of 0.5 meV for the limiting case that $\sin \beta = 1$. In this case, the cooling models only depend on a single axion parameter, the axion mass m_a . More generally, in parameter regimes where energy loss due to the axion Primakoff effect is negligible, as is the case for the white dwarf cooling regime that we consider in this work, the cooling models only depend on the coupling constant g_{aee} (i.e. the dependence of the cooling models on $g_{a\gamma\gamma}$ is negligible). The work presented in this chapter in terms of m_a (for a DFSZ axion with $\sin \beta = 1$) can thus equivalently be expressed in terms g_{aee} , with the results in terms of g_{aee} applying more generally. Table 7.1 provides the conversion between m_a and the g_{aee} (and $g_{a\gamma\gamma}$) values used for the simulations.

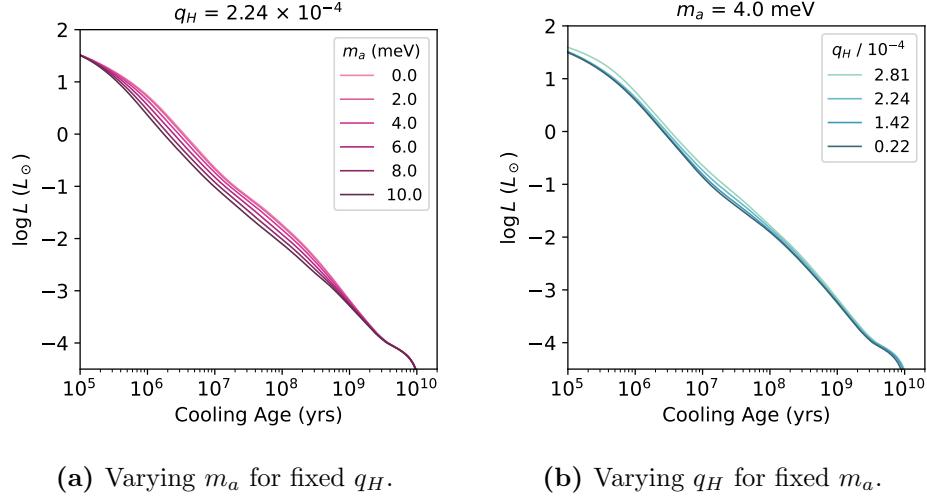
7.3.2 Cooling Curves

Fig. 7.1 shows how the cooling curves varying with axion mass (m_a) and H envelope thickness (q_H). The left panel (Fig. 7.1a) shows a series of cooling curves with different m_a values over the range 0 – 10 meV for a fixed envelope thickness of $q_H = 2.24 \times 10^{-4}$, while the right panel (Fig. 7.1b) shows a series of cooling curves with different q_H values for a fixed axion mass of $m_a = 4.0$. For both series of cooling curves, the white dwarf mass is fixed at $M_{\text{WD}} = 0.5388 M_{\odot}$ and the diffusion treatment is always standard MESA diffusion (i.e. diffusion is fully on). Note that cooling curves are only shown for a subset of m_a and q_H values in order to facilitate visualisation;

³²See in particular Table 2.1 of Raffelt [51] for a summary of the expressions for various axion coupling constants. Note that $\cos \beta$ has been replaced by $\sin \beta$ in the expression for g_{aee} in accordance with the definition used for β elsewhere in this dissertation. The choice of which definition of β to use is simply a matter of convention.

$m_a / (1 \text{ meV})$	$g_{aee} / (10^{-13})$	$g_{a\gamma\gamma} / (10^{-12} \text{ GeV}^{-1})$
0.0	0.00	0.00
0.5	0.14	0.07
1.0	0.28	0.14
1.5	0.42	0.22
2.0	0.57	0.29
2.5	0.71	0.36
3.0	0.85	0.43
3.5	0.99	0.51
4.0	1.13	0.58
4.5	1.28	0.65
5.0	1.42	0.72
5.5	1.56	0.79
6.0	1.70	0.87
6.5	1.84	0.94
7.0	1.98	1.01
7.5	2.12	1.08
8.0	2.27	1.15

Table 7.1: Conversion between axion mass values and the corresponding values used for the coupling constants in the MESA simulations that produced the white dwarf cooling models. These values apply for a DFSZ model with $\sin \beta = 1$. For these values of $g_{a\gamma\gamma}$, the effect of axion-photon interactions on white dwarf cooling is negligible; these values are simply reported for completeness.


 (a) Varying m_a for fixed q_H .

 (b) Varying q_H for fixed m_a .

Figure 7.1: Theoretical cooling curves for various axion mass values (left; a) and envelope thickness (right; b) when all other model parameters are fixed. In both cases, the white dwarf mass is $M_{\text{wd}} = 0.5388 M_{\odot}$ and diffusion is fully on. For the series of axion mass values shown in the left panel, the envelope thickness always has a fixed value of $q_H = 2.24 \times 10^{-4}$. For the series of q_H values shown in the right panel, the axion mass always has a fixed value of $m_a = 4.0 \text{ meV}$.

the full set of models used in the analysis is more finely spaced in m_a and q_H than what is shown in Fig. 7.1.

For the axion mass range $0 - 10.0 \text{ meV}$ considered in our analysis, the emission of axions produced in the core of the white dwarf alters the shape of the cooling curves for cooling ages between $\sim 1 \times 10^5 \text{ yrs}$ and $3 \times 10^9 \text{ yrs}$ and over the luminosity range $1.5 \lesssim \log L \lesssim -4$, where L is the luminosity in units of solar luminosity and \log denotes the base-10 logarithm. This can be seen in Fig. 7.1a, where the cooling curves differ over the luminosity (and cooling age) range where axion emission is an important energy loss mechanism. Larger values of the axion coupling (corresponding to large values of m_a) cause the luminosity of the white dwarf to decrease more quickly with cooling age. The cooling curves for all of the m_a values re-converge at late cooling times once energy loss due to axion emission becomes

negligible compared to energy loss due to photon radiation.

Over the cooling regime where axion emission is an important energy emission mechanism, the morphology of the cooling curve also depends on the H envelope thickness. This can be seen in Fig. 7.1b, where it can be seen that the shape of the cooling curves varies within the relevant luminosity range with varying q_H for a fixed axion mass value. The cooling curve depends on q_H over this regime because thicker H envelopes result in more residual nuclear burning at the boundary of the envelope, which in turn causes the white dwarf to cool more slowly and remain brighter for a longer period of time early in the cooling process.

7.4 Analysis

Our analysis uses the unbinned likelihood in a procedure very similar to the analysis in Goldsbury *et al.* [215]. Similar methods were also used in Section 6.7 and are discussed more in that section. Unlike in Section 6.7, however, the dependence of the likelihood function on the distance R from the cluster centre cannot be neglected.

The radial dependence is in particular important for analysing the WFC3 data because the WFC3 observations are centred on the cluster centre, where stars belonging to the cluster are most densely concentrated. Overcrowding in the observations near the cluster centre due to this dense concentration results in reduced completeness at small values of R . The completeness of the WFC3 observations is also reduced to a lesser extent at large values of R due to the geometry of the observations, with less overlap in the fields at larger values of R resulting in fewer fields that an object may be detected in. Furthermore, the radial density distribution (in projection) of a globular cluster in general, and 47 Tuc in particular, varies most rapidly with R at small values of R where the cluster is most concentrated; the distribution is less sensitive to R in the tails where the density varies more slowly with R . The ACS observations were taken far enough from the centre that radial dependence can be neglected for the ACS data.

For a particular data set, the number density distribution function ac-

counting for radial dependence is given by the expression [215]

$$f(m_1, m_2, R; \theta) = \dot{N} f_R(R; r_0, r_t) \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f_M(m'_1, m'_2; \theta_M) \times E(m_1 - m'_1, m_2 - m'_2; m'_1, m'_2, R) dm'_1 dm'_2, \quad (7.3)$$

where the photometric error distribution function E is parameterised by R (in addition to the input magnitudes m'_1 and m'_2), f_R is the projected radial (surface) density distribution after integrating over the azimuth angle, f_M is a function that quantifies the cooling rate, θ_M denotes the set of parameters that the cooling models depends on (after being moved from theory space to data space), and $\theta = \{\theta_M, \dot{N}, r_0, r_t\}$ denotes the set of all parameters that the full model depends on, which includes the white dwarf birthrate \dot{N} and any parameters that parameterise the radial density distribution. In the case that radial dependence can be neglected, Eq. (7.3) reduces to Eq. (6.30), the expression used for the analysis in Section 6.7.

The radial density distribution f_R is, stated properly, a probability density function that we take to be normalised over the range of R values defining the relevant data space. We use the same radial density distribution³³ as Goldsbury *et al.* [215], which is a King-Michie model [258, 259] with a King radius of $r_0 = 32''$ and tidal radius of $r_t = 3800''$. The parametric dependence of f_R is written explicitly in Eq. (7.3) to make the model dependence on these parameters in general clear; however, these parameters are held fixed in the analysis performed in this chapter. We follow the notation of Goldsbury *et al.* [260] and Goldsbury *et al.* [215] whereby a lowercase r denotes a three-dimensional radial distance from the cluster centre and an uppercase R denotes the corresponding two-dimensional radius in projection. The two-dimensional nature of f_R is made more transparent by expressing it as

$$f_R(R) = \frac{2\pi R \Sigma(R)}{\int 2\pi R \Sigma(R) dR}, \quad (7.4)$$

³³Note that the function $f_R(R)$ in our work is called $\rho(R)$ in Goldsbury *et al.* [215]. We use the notation f_R to make the nature of this quantity as a two-dimensional probability density function more clear.

where Σ is the projected surface density given by an Abel transform of the three-dimensional density distribution, as described in Goldsbury *et al.* [260], and the integration to normalise f_R is performed over the R limits defining the data space. The procedure for evaluating Σ as a function of R is described in Goldsbury *et al.* [260], which entails numerically solving the system of equations describing the King-Michie model and then performing an Abel transform on the result.

The magnitudes represented by m_1 and m_2 (and their primed equivalents) are $m_1 = \text{F225W}$ and $m_2 = \text{F336W}$ in the case of the WFC3 data, while they are $m_1 = \text{F435W}$ and $m_2 = \text{F555W}$ in the case of the ACS data. The function f_M quantifies the theoretical cooling rate predicted by the cooling model as a function of both magnitudes before accounting for photometric errors. Note that the quantity $\dot{N} f_M$ is the theoretical number density distribution of the magnitude values before accounting for photometric errors. The expression for f_M given in Chapter 6 and related discussion also applies here,

$$f_M(m'_1, m'_2; \theta_M) = \frac{dt}{dm'_1} \delta[m'_2 - m_{2,\text{mod}}(m'_1; \theta_M)], \quad (7.5)$$

though some of the parameters differ in the set of model parameters represented by θ_M . The theoretical cooling models before moving to data space are parameterised by the white dwarf mass M_{WD} , the H envelope thickness parameter $\log_{10} q_H$, and the axion mass m_a . Moving the theoretical models to data space depends on the distance modulus μ and colour excess $E(B - V)$, so the function f_M is parameterised by the parameter set $\theta_M = \{m_a, M_{\text{WD}}, \log_{10} q_H, \mu, E(B - V)\}$. Since converting the cooling rate to a number density depends on the white dwarf birthrate \dot{N} and the radial density distribution depends on the additional set of parameters $\theta_R = \{r_0, r_t\}$, the full set of parameters that the likelihood depends on is $\theta = \{m_a, M_{\text{WD}}, \log_{10} q_H, \dot{N}, \mu, E(B - V), r_0, r_t\}$.

As was done in Chapter 6, the values of μ and $E(B - V)$ are held fixed, with a value of $\mu = 13.24$ used for the distance modulus [214] and a value of $E(B - V) = 0.04$ used for the colour excess [210]. For this value of the colour

7.4. Analysis

Parameter	Value
r_0	32"
r_t	3800"
μ	13.24
$E(B - V)$	0.04

Table 7.2: Summary of model parameter values held fixed in analysis.

excess, the corresponding value of the total V-band extinction is $A_V = 0.124$, where the relative visibility has been taken to be $R_V = 3.1$, the typical value for the Milky Way. The values of all of the parameters that are held fixed are summarised in Table 7.2. With μ and $E(B - V)$ held fixed in addition to r_0 and r_t , as described above, the set of parameters θ_M that the cooling rate depends on in the analysis is reduced to $\theta_M = \{m_a, M_{WD}, \log_{10} q_H, \}$, and the full model (and thus the likelihood) depends only on the reduced set of parameters $\theta = \{m_a, M_{WD}, \log_{10} q_H, \dot{N}\}$.

The input magnitudes before accounting for photometric errors, denoted by a prime symbol in Eqs. (7.3) and (7.5), are calculated from the relevant theory-space cooling model variables using bolometric corrections via the procedure described in Section 6.7, but with extinction values appropriate for the magnitude filters of the data considered in this chapter. The relevant extinctions for the WFC3/UVIS filters are given by

$$A_{F225W} = 2.62940 A_V, \quad (7.6)$$

$$A_{F336W} = 1.67536 A_V, \quad (7.7)$$

and the relevant extinctions for the ACS/WFC filters are given by

$$A_{F435W} = 1.33879 A_V, \quad (7.8)$$

$$A_{F555W} = 1.03065 A_V. \quad (7.9)$$

Note that these extinctions were determined using the same extinction law³⁴,

³⁴The extinction values were retrieved using http://stev.oapd.inaf.it/cgi-bin/cmd_3.7, which uses this extinction law.

7.4. Analysis

that of Cardelli *et al.* [38] and O'Donnell [39], as was used in Section 6.7.

The natural logarithm of the unbinned likelihood is calculated from the distribution function f given by Eq. (7.3) using the expression

$$\ln \mathcal{L}(\theta) = \sum_i \ln f_i - N_{\text{pred}}, \quad (7.10)$$

where f_i is the distribution function evaluated at the data coordinate $d_i = (m_{1i}, m_{2i}, R_i)$ observed for the i th data point,

$$f_i = f(m_{1i}, m_{2i}, R_i; \theta), \quad (7.11)$$

and N_{pred} is the model prediction for the total number of white dwarfs in the data space,

$$N_{\text{pred}} = \iint_{\text{data space}} f(m_1, m_2, R; \theta) \, dR \, dm_1 \, dm_2. \quad (7.12)$$

The data space boundaries as they appear in a CMD for each data set considered in our analysis are shown in Fig. 7.2. Note that the analysis is actually performed in magnitude-magnitude space. The data space is simply shown in colour-magnitude space for better visualisation, and it is straightforward to move between colour-magnitude and magnitude-magnitude space. The WFC3/UVIS data and corresponding data space selection are shown in Fig. 7.2a, while the ACS/WFC data and corresponding data space selection are shown in Fig. 7.2b. These plots are focused on the white dwarf cooling sequence of 47 Tuc, with the data space boundaries shown as red curves enclosing a large part of the cooling sequence. For each data set, the evolutionary track predicted by the model for the relevant magnitudes before accounting for photometric errors is shown by the orange curve that lies along the 47 Tuc cooling sequence and passes through the data space.

The particular reference model shown in Fig. 7.2 corresponds to parameter values of $M_{\text{WD}} = 0.5388 M_{\odot}$, $m_a = 0$ meV, and $\log_{10} q_H = -3.60$, though it should be noted that the evolutionary track in the CMD shows

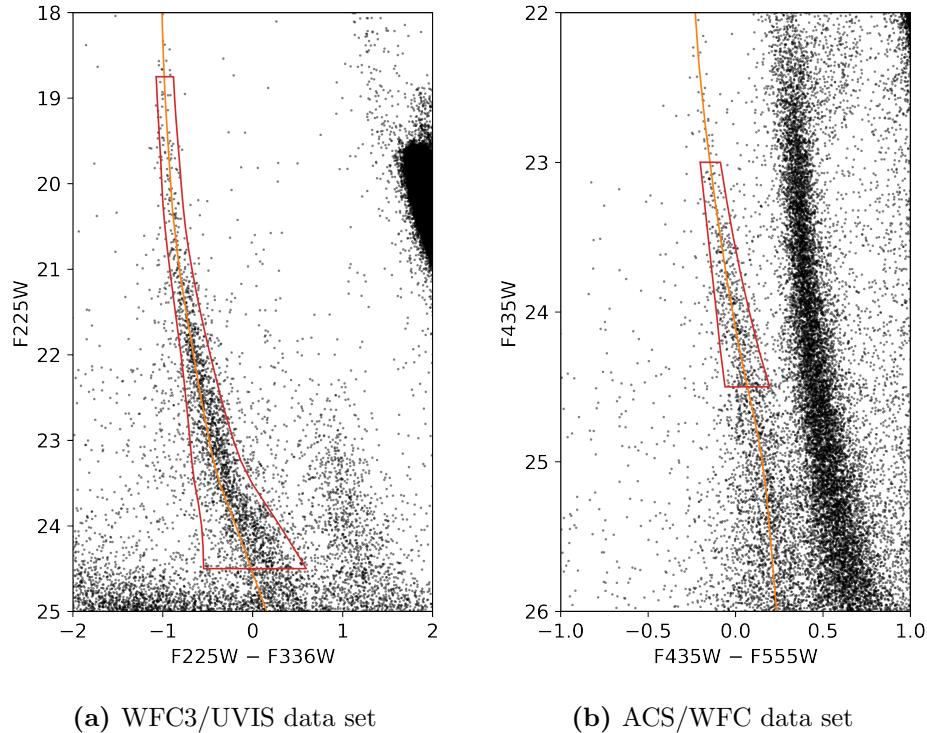


Figure 7.2: CMDs showing the data space selections used in the unbinned likelihood analysis for the WFC3/UVIS data (left; a) and ACS/WFC data (right; b). The plots are focused on the white dwarf cooling sequence, with the data shown as black points. The boundaries of the white dwarf data space selection for each data set are indicated by the red curves. The evolutionary track predicted by the model in each case before accounting for photometric errors is shown as an orange curve.

7.4. Analysis

very little variation over the range of parameter values considered in our analysis. It can be seen visually from Fig. 7.2 that the theoretical cooling tracks align reasonably well with the observed data in the CMD. For the WFC3 data, the model lies left of the centre line of the empirical cooling sequence because the photometric errors are not distributed equally in the two filters and the error distribution has not yet been applied to the model shown in Fig. 7.2. For the ACS data, the photometric error distribution is more symmetric in the two filters, and it can correspondingly be seen in Fig. 7.2 that the model lies more closely along the centre line of the empirical white dwarf sequence of the ACS data.

The WFC3 data space boundaries are chosen to be approximately 3σ in colour from the model cooling track, while the ACS data space boundaries are chosen to be approximately 2σ from the model track. In each case, the error in colour was determined as a function of magnitude m_1 (the y -axis magnitude in the relevant sub-figure of Fig. 7.2) from the photometric error distribution given by artificial stars tests. These data space boundaries are chosen such they enclose the 47 Tuc white dwarf sequence tightly enough in colour to avoid contamination from the SMC. In Fig. 7.2, the SMC corresponds to the sequence of stars located to the right of the 47 Tuc white dwarf sequence and running approximately parallel to it. The ACS data space is shorter than the WFC3 data space because the SMC sequence begins to intersect with the 47 Tuc white dwarf sequence at brighter magnitudes (corresponding to earlier cooling times) in the ACS data than in the WFC3 data, as can be seen in Fig. 7.2.

In the case of the WFC3 data (for which the variable R is included in the analysis), the data space definition also includes a cut in the radial distance from the cluster centre of $R \leq 4,000$ pixels³⁵. The WFC3/UVIS pixel scale is 0.04 arcseconds per pixel, so this corresponds to a cut of $R \leq 160$ arcsec in physical units. This cut removes objects located at the outermost radii of the total WFC3 field of view where there is not full coverage about the entire circumference of a circle of that radius; thus, this cut circumvents

³⁵This cut was also used in the analysis of Goldsbury *et al.* [215], though it is not explicitly mentioned in that paper.

7.4. Analysis

any potential concerns about a reduction in completeness due to the field geometry at these outermost radii.

The birthrate of white dwarfs for a particular data set is specific to the field of view for those observations, so the WFC3 and ACS data are expected to have different birthrates. A combined analysis of both of these data sets thus depends on two birthrate parameters, which will be denoted as \dot{N}_{WFC3} for WFC3 and \dot{N}_{ACS} for ACS. This is accounted for in a combined analysis of the WFC3 and ACS data sets as follows.

The theoretical cooling models predict separate distribution functions for each set of data: $f_{\text{WFC3}}(\text{F225W}, \text{F336W}, R; \theta)$ with $\theta = \{\dot{N}_{\text{WFC3}}, \theta_M\}$ for the WFC3 data and $f_{\text{ACS}}(\text{F435W}, \text{F555W}; \theta)$ with $\theta = \{\dot{N}_{\text{ACS}}, \theta_M\}$ for the ACS data. Likelihoods can thus be calculated for each data set using Eq. (7.10) for the appropriate distribution function. Letting D_{WFC3} and D_{ACS} denote the set of WFC3 data points and ACS data points, respectively, in the relevant data space for each field, i.e.

$$D_{\text{WFC3}} \equiv \{\text{F225W}_i, \text{F336W}_i, R_i\} \quad \text{for } i \in \text{WFC3 data points}, \quad (7.13)$$

$$D_{\text{ACS}} \equiv \{\text{F435W}_i, \text{F555W}_i\} \quad \text{for } i \in \text{ACS data points}, \quad (7.14)$$

the corresponding likelihoods $\mathcal{L}_{\text{WFC3}}$ and \mathcal{L}_{ACS} for each of these data sets separately are defined as

$$\mathcal{L}_{\text{WFC3}}(\dot{N}_{\text{WFC3}}, \theta_M) \equiv p(D_{\text{WFC3}} | \dot{N}_{\text{WFC3}}, \theta_M), \quad (7.15)$$

$$\mathcal{L}_{\text{ACS}}(\dot{N}_{\text{ACS}}, \theta_M) \equiv p(D_{\text{ACS}} | \dot{N}_{\text{ACS}}, \theta_M), \quad (7.16)$$

where p denotes a probability density function. The likelihood $\mathcal{L}_{\text{comb}}$ of both data sets combined is defined as

$$\mathcal{L}_{\text{comb}}(\theta) \equiv p(D_{\text{WFC3}}, D_{\text{ACS}} | \dot{N}_{\text{WFC3}}, \dot{N}_{\text{ACS}}, \theta_M), \quad (7.17)$$

where $\theta = \{\dot{N}_{\text{WFC3}}, \dot{N}_{\text{ACS}}, \theta_M\}$ is the set of all parameters considered in the combined analysis.

7.4. Analysis

Since each data point is statistically independent, it follows that

$$\begin{aligned} p(D_{\text{WFC3}}, D_{\text{ACS}} | \dot{N}_{\text{WFC3}}, \dot{N}_{\text{ACS}}, \theta_M) \\ = p(D_{\text{WFC3}} | \dot{N}_{\text{WFC3}}, \theta_M) p(D_{\text{ACS}} | \dot{N}_{\text{ACS}}, \theta_M). \end{aligned} \quad (7.18)$$

The log-likelihood of the combined data is thus simply given by the sum of the log-likelihoods of the separate data sets, i.e.

$$\ln \mathcal{L}_{\text{comb}}(\theta) = \ln \mathcal{L}_{\text{WFC3}}(\dot{N}_{\text{WFC3}}, \theta_M) + \ln \mathcal{L}_{\text{ACS}}(\dot{N}_{\text{ACS}}, \theta_M). \quad (7.19)$$

The joint posterior for the combined analysis is then given by

$$p(\theta | D_{\text{WFC3}}, D_{\text{ACS}}) = \frac{p(\theta) \mathcal{L}_{\text{comb}}(\theta)}{\int p(\theta) \mathcal{L}_{\text{comb}}(\theta) d\theta}, \quad (7.20)$$

where the joint prior distribution $p(\theta)$ for all of the parameters can be expressed as

$$p(\theta) = p(\dot{N}_{\text{WFC3}}) p(\dot{N}_{\text{ACS}}) p(M_{\text{WD}}, \log_{10} q_H) p(m_a) \quad (7.21)$$

and the integral in the denominator of Eq. (7.20) is performed over the entire parameter space.

For the main results reported in Section 7.5 below, uniform priors were used for both of the birthrate parameters. The option of using Gaussian birthrate priors with the same values as used by Goldsbury *et al.* [215] is considered in Appendix B. However, it is found in Appendix B (for each data set individually) and shown below in Section 7.5 (for the combined analysis) that the birthrate priors of Goldsbury *et al.* [215] both overestimate \dot{N}_{WFC3} and underestimate \dot{N}_{ACS} by a similar amount. This suggests that there may be a relevant effect that was not accounted for in determining (or applying) those birthrate priors, such as the effect of cluster relaxation [261, 262] causing stars to leave the WFC3 field and enter the ACS field over time. Though using Gaussian priors for the birthrates could potentially give tighter bounds on the parameters, uniform priors are used for the birthrates

to avoid the risk of imposing an incorrect prior value.

The results of Chapter 6 provided priors for M_{WD} and $\log_{10} q_H$. The analysis presented in Section 7.5 used a joint prior $p(M_{\text{WD}}, \log_{10} q_H)$ for M_{WD} and $\log_{10} q_H$ given by the joint posterior distribution from Chapter 6 for the standard diffusion scenario after marginalising over the birthrate. This corresponds to the distribution plotted in the bottom panel of Fig. 6.10, but locally normalised to that particular diffusion scenario. Using uniform priors for M_{WD} and $\log_{10} q_H$ instead yields similar best-fitting parameter values (when uniform birthrate priors are used) but a more extended posterior distribution (with some degeneracy between M_{WD} and $\log_{10} q_H$ for the WFC3 data in particular), so using the Chapter 6 joint prior for M_{WD} and $\log_{10} q_H$ provides tighter parameter constraints. This is shown in detail in Appendix B. In all cases considered, in both Section 7.5 and Appendix B, the prior for m_a was taken to be uniform.

7.5 Results

The results for the combined fit of the cooling models to both the WFC3/UVIS and ACS/WFC data are presented in this section. The combined fit to both of these data sets gives stronger constraints than fitting either one of these data sets separately. Nevertheless, separate fits of the models to each of these data sets should give similar results to the combined fit if each of the data sets are reasonably well-fitted by the models, which is verified in Appendix B. In addition to the combined analysis presented in this section, which gives the main results of the work described in this chapter, we also independently fitted the cooling models to each of the two data sets and present those results in Appendix B. The results for the WFC3 data alone are given in Appendix B.2, while the results for the ACS data alone are given in Appendix B.3. The key results of the separate WFC3 and ACS analyses are compared with the results of the combined analysis in Appendix B.4.

7.5.1 Posterior Distributions

The joint posterior density distribution after marginalising over both of the birthrates is shown in Fig. 7.3 as filled contour plots. Each plot shows a slice of this distribution as a function of m_a and $\log_{10} q_H$ for a particular value of M_{WD} . In these plots, the posterior density p (marginalised over \dot{N}_{WFC3} and \dot{N}_{ACS}) has been scaled by its maximum value \hat{p} attained on the parameter grid, and the contour levels are drawn at p/\hat{p} values of -0.5, -2.0, -4.5, -8.0, and -12.5. These level values correspond to the analogously scaled probability density values of a normal distribution evaluated at 1, 2, 3, 4, and 5 σ , with $p/\hat{p} = \exp(-0.5 n_\sigma^2)$ for a level value corresponding to $n_\sigma \sigma$. Plots are only shown for the M_{WD} values from the parameter grid that are most significant for the posterior distribution: 0.5314 and $0.5388 M_\odot$. For the case of $M_{WD} = 0.5240 M_\odot$, which has been omitted from Fig. 7.3, the value of p/\hat{p} is less than the smallest displayed level, i.e. $p/\hat{p} < 5 \sigma$, for every point on the grid of m_a and $\log_{10} q_H$ values.

It should be emphasised that the filled contours in Fig. 7.3 indicate ranges of probability density values, not regions of enclosed probabilities, so they are not credible regions. The two-dimensional credible regions in the joint parameter space of m_a and $\log_{10} q_H$ are shown in Fig. 7.4. The credible regions shown in Fig. 7.4 are more specifically the highest posterior density credible regions calculated from the joint posterior distribution shown in Fig. 7.3 after marginalising over M_{WD} as well as \dot{N}_{WFC3} and \dot{N}_{ACS} . The filled contours in Fig. 7.4 demarcate regions of enclosed probability, and the total probability enclosed by each contour (including the regions coloured darker than that contour level) is indicated on the colour bar in terms of the number of standard deviations of a (spherically symmetric) two-dimensional normal distribution that encloses the same probability. The contour levels of 1, 2, 3, 4, and 5 σ respectively correspond to total enclosed probabilities of 39.35, 86.47, 98.89, 99.97, and 99.9996%.

Fig. 7.3 and Fig. 7.4 both show that smaller m_a and larger $\log_{10} q_H$ values are favoured, with the peak of the joint posterior distribution occurring for $m_a = 0$ meV and $\log_{10} q_H = -3.55$ on the parameter grid. Fig. 7.3 and

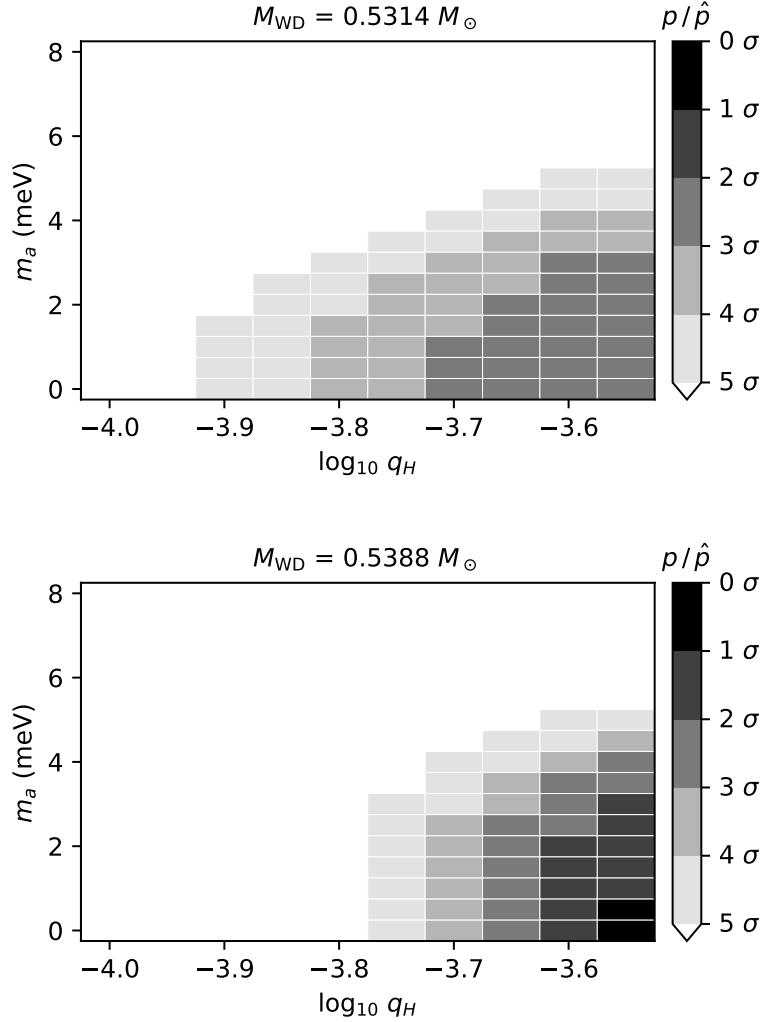


Figure 7.3: Joint posterior probability density distribution after marginalising over the birthrates. Slices of the distribution as a function of axion mass (m_a) and envelope thickness (q_H) are shown for fixed values of white dwarf mass (M_{WD}). The probability density (p) has been scaled by its maximum value (\hat{p}) so that the plotted quantity is p/\hat{p} . The filled contours are drawn at level values corresponding to the σ levels indicated on the legend for a two-dimensional normal distribution. The lowest mass case ($M_{\text{WD}} = 0.5240 M_{\odot}$) is not shown because it is excluded at 5σ .

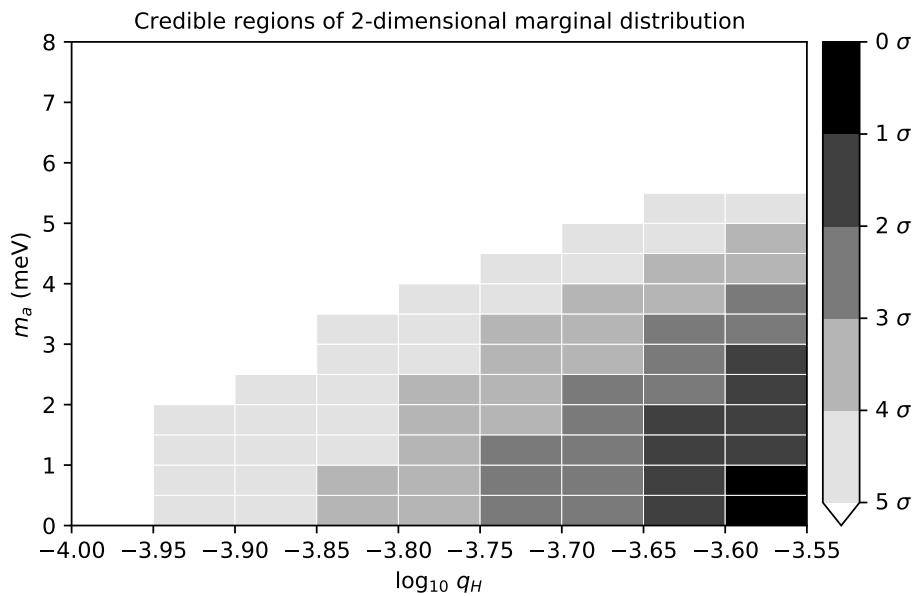


Figure 7.4: Two-dimensional joint credible regions of axion mass (m_a) and envelope thickness (q_H) after marginalising over the other parameters. The filled contours show regions of enclosed probability containing the same probability as enclosed by a two-dimensional normal distribution for the σ levels indicated by the legend.

7.5. Results

Fig. 7.4 also highlight the degeneracy between m_a and $\log_{10} q_H$: larger m_a values are more probable when the value of $\log_{10} q_H$ is also large, though small m_a values are still favoured even at the largest $\log_{10} q_H$ values. As can be seen in Fig. 7.3, the morphology of the marginal posterior distribution as a function of m_a and $\log_{10} q_H$ is similar regardless of M_{WD} , with the distribution simply becoming more tightly concentrated on the highest $\log_{10} q_H$ values for larger M_{WD} . The peak in the posterior distribution occurs for $M_{\text{WD}} = 0.5388 M_{\odot}$. This is the largest value of M_{WD} considered in the analysis, but larger values (on the extended parameter grid considered in Chapter 6) are strongly disfavoured by the results of Chapter 6 (which were used as the prior on M_{WD} and $\log_{10} q_H$ in this analysis).

The maximum of the posterior distribution also occurs at the (upper) edge of the $\log_{10} q_H$ grid, but this grid cannot be extended to larger $\log_{10} q_H$ values (with the current grid spacing) because this is a physical limit. In creating the suite of white dwarf cooling simulations used in this work, white dwarfs that started with thicker H envelopes simply burned away the extra H in the envelope through residual nuclear burning at very early cooling times and ended up with the same envelope thickness at the reference cooling age of 10 Myr at which $\log_{10} q_H$ is defined in this work. Furthermore, the cooling models with the same $\log_{10} q_H$ value (at this reference time) produced by these simulations were the same over the magnitude range of interest for this work, regardless of the initial thickness at very early times.

The one-dimensional posterior probability density distributions of each model parameter after marginalising over all other parameters are shown in Fig. 7.5. The posterior distribution for the WFC3 birthrate is shown in Fig. 7.5a, while the posterior distribution for the ACS birthrate is shown in Fig. 7.5b. For each birthrate distribution, the value of \dot{N} that maximises the posterior distribution is indicated by a dashed vertical line and the prior value \dot{N}_0 used by Goldsbury *et al.* [215] is indicated by a dotted vertical line for comparison. It can be seen in these plots that the posterior for \dot{N}_{WFC3} is notably smaller than the Goldsbury *et al.* [215] prior value, while the posterior for \dot{N}_{ACS} is larger than the Goldsbury *et al.* [215] prior value. This may be an indication that an appreciable number of white dwarfs are

7.5. Results

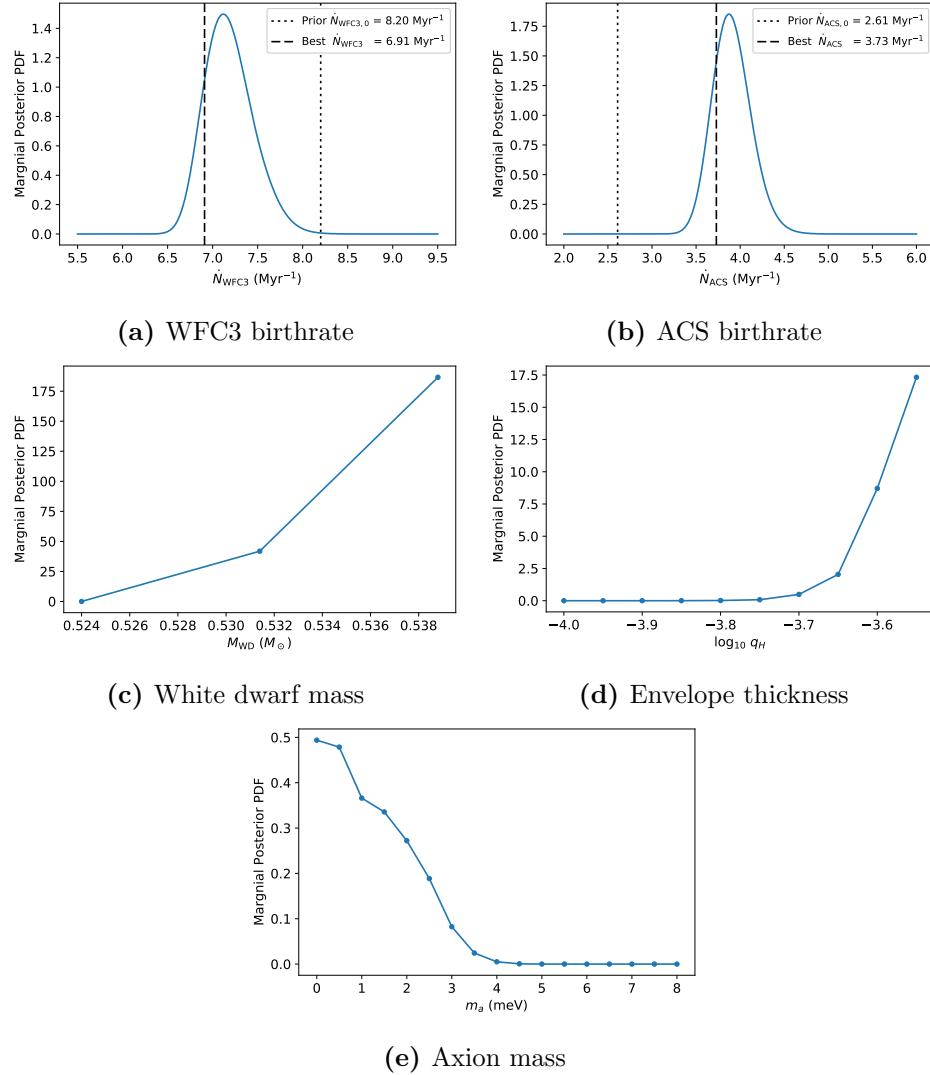


Figure 7.5: One-dimensional posterior probability density distributions for each parameter after marginalising over all other model parameters.

leaving the WFC3 field of view (i.e. the inner field) and entering the ACS field of view (i.e. the outer field).

Fig. 7.5c, Fig. 7.5d, and Fig. 7.5e show the one-dimensional posterior distributions for M_{WD} , $\log_{10} q_H$, and m_a , respectively. Some of the trends noted above for these parameters based on Fig. 7.3 and Fig. 7.4 can be seen clearly in Figs. 7.5c to 7.5e. Fig. 7.5c highlights the overall increased probability of larger M_{WD} values over the limited range of values considered in this analysis, which was noted above based on Fig. 7.3. Fig. 7.5d and Fig. 7.5e show that large $\log_{10} q_H$ and small m_a values are favoured, as is also clearly seen in the plot of the two-dimensional credible regions for these parameters, Fig. 7.4. Using the one-dimensional marginal posterior distributions to calculate the individual 95% confidence levels of each of these parameters, it is found that $-3.67 \leq \log_{10} q_H \leq -3.55$ and $m_a \leq 2.85$ meV at 95% confidence. This limit for the DFSZ axion mass corresponds to a limit for the coupling constant of $g_{aee} \leq 0.81 \times 10^{-13}$ for any axion or axion-like particle model with an axion-electron interaction.

7.5.2 Best-Fitting Model

The combination of parameter values for which the full joint posterior distribution is maximised on the discrete parameter grid for the combined analysis is summarised in Table 7.3. These parameter values correspond to the optimal model on the parameter grid for the combined analysis after accounting for all of the priors and before marginalising over any parameters. The 95% credible regions calculated from the one-dimensional marginal posterior distributions shown in Fig. 7.5 are given in Table 7.3 as errors on the parameter values of the best-fitting model. Note that these parameter values are not necessarily the values that optimise the marginal posteriors from which the credible regions were calculated, though it can be seen from Fig. 7.5 that these values are similar (and indeed exactly the same for M_{WD} , $\log_{10} q_H$, and m_a). Also note that although the values of M_{WD} , $\log_{10} q_H$, and m_a (but not \dot{N}_{WFC3} or \dot{N}_{ACS}) for the best-fitting model are restricted to the parameter grid of the cooling models, the limits of the credible regions are

7.5. Results

Parameter	Value
$\dot{N}_{\text{WFC3}} (\text{Myr}^{-1})$	$6.91_{-0.23}^{+0.82}$
$\dot{N}_{\text{ACS}} (\text{Myr}^{-1})$	$3.73_{-0.24}^{+0.62}$
$M_{\text{WD}} (M_{\odot})$	$0.5388_{-0.0106}^{+0.0000}$
$\log_{10} q_H$	$-3.55_{-0.12}^{+0.00}$
$m_a (\text{meV})$	$0.00_{-0.00}^{+2.85}$
$g_{aee} / 10^{-13}$	$0.00_{-0.00}^{+0.81}$

Table 7.3: Parameter values of optimal model for combined analysis. The values are reported for the combination of parameter values that maximises the joint posterior distribution on the full parameter grid. The errors indicate the 95% credible region calculated from the one-dimensional marginal posterior distribution for each parameter.

not restricted to the grid points because linear interpolation was used when calculating the credible regions. For easy reference, the g_{aee} values corresponding to the best-fitting value and credible region limits of m_a have also been given in Table 7.3.

To visually verify the goodness of fit of the optimal model, we plot the one-dimensional marginal cumulative number distributions predicted by the optimal model over the corresponding empirical distributions for each data component of each of the two data sets. These are shown in Figs. 7.6 and 7.7, with the model prediction shown as the red curve and the data shown as black points. In all of these cases, the incompleteness of the empirical observations is accounted for in the model instead of applying any completeness correction to the data.

Fig. 7.6 shows the cumulative number distribution with respect to the radial distance (R) from the cluster centre for the WFC3/UVIS data, with R given in units of WFC3/UVIS pixels. Note that the unit conversion between WFC3/UVIS pixels and arcseconds is 0.04 arcsec/pixel [263, 264]. The one-dimensional marginal number density distributions with respect to R predicted by the model for the WFC3 data is calculated by integrating the full

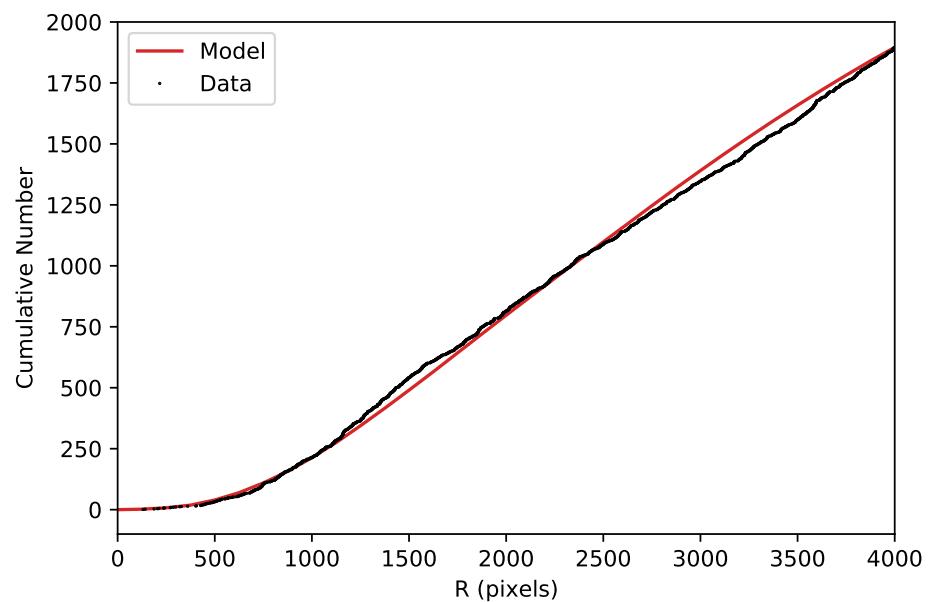


Figure 7.6: Cumulative number distribution of radial distance (R) from cluster centre for optimal model (red curve) and WFC3/UVIS data (black points).

7.5. Results

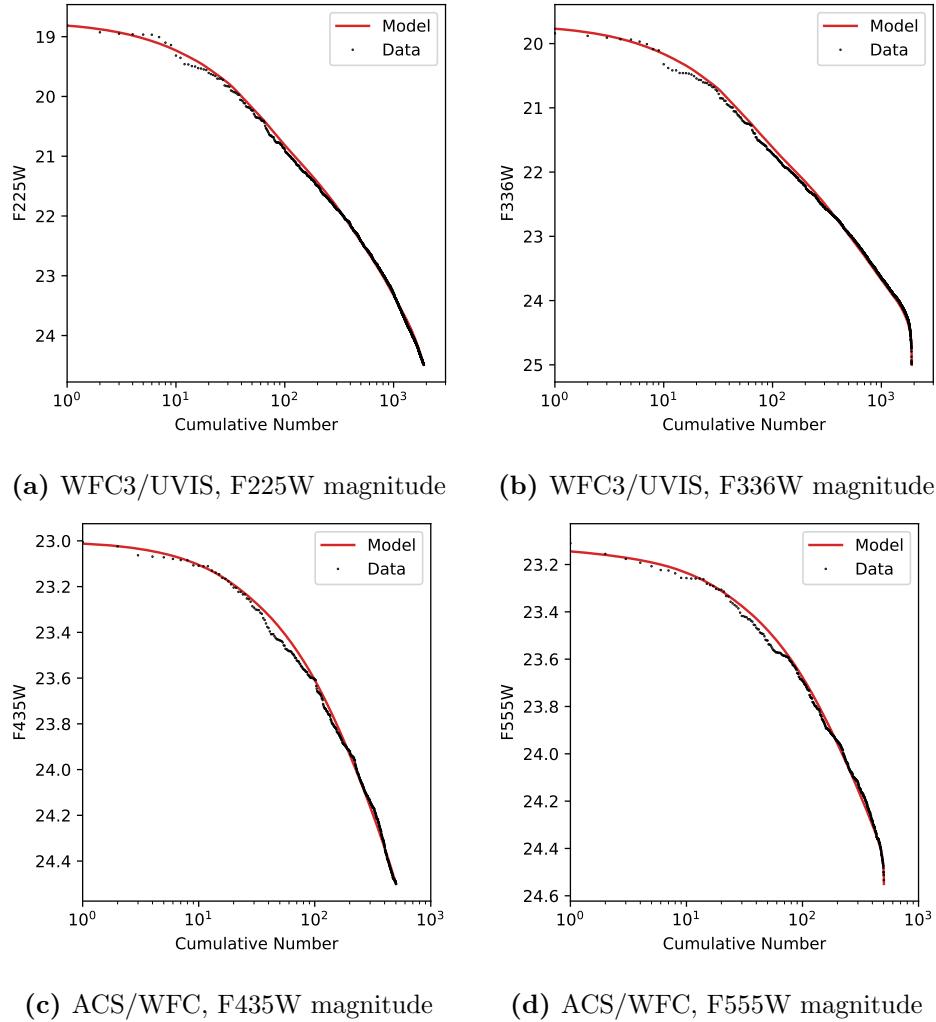


Figure 7.7: Inverse cumulative luminosity function of optimal model from combined fit (red curve) compared to the two sets of HST data used in the combined fit (black points). The different sub-figures show the empirical distributions for the WFC3/UVIS data (top row) with respect to F225W magnitude (left column; a) and F336W magnitude (right column; b) and the distributions for the ACS/WFC data (bottom row) with respect to F435W magnitude (left column; c) and F555W magnitude (right column; d).

7.5. Results

three-dimensional number density distribution $f_{\text{WFC3}}(\text{F225W}, \text{F336W}, R)$ with respect to the magnitudes F225W and F336W. The corresponding marginal cumulative number distribution with respect to R (plotted as the red curve in Fig. 7.6) is then calculated from this marginal number density distribution by integrating up to each value of R that is plotted. Note that since our models for the ACS/WFC data do not depend on R , there is no plot analogous to Fig. 7.6 to be made for the ACS/WFC data.

Fig. 7.7 shows the inverse cumulative luminosity functions for each magnitude of each data set. Note that the cumulative luminosity function with respect to a particular magnitude is simply the cumulative number distribution for that magnitude. For a particular magnitude, the model prediction is given by the cumulative number density distribution after marginalising over R (if applicable) and the other magnitude that f depends on. The top row of Fig. 7.7 shows the (inverse) distributions for the WFC3 data with respect to the two different WFC3 magnitudes, $m_1 = \text{F225W}$ in Fig. 7.7a and $m_2 = \text{F336W}$ in Fig. 7.7b. The theoretical distributions (predicted by the optimal model) that are shown in Figs. 7.7a and 7.7b are both calculated from the same distribution function $f_{\text{WFC3}}(\text{F225W}, \text{F336W}, R)$, but with a different magnitude marginalised over in each case (and R marginalised over in both cases). Likewise, the bottom row of Fig. 7.7 shows the (inverse) distributions for the ACS data with respect to the two different ACS magnitudes, $m_1 = \text{F435W}$ in Fig. 7.7c and $m_2 = \text{F555W}$ in Fig. 7.7d, and the theoretical distributions that are shown in Figs. 7.7c and 7.7d are both calculated from the same distribution function $f_{\text{ACS}}(\text{F435W}, \text{F555W})$, but with a different magnitude marginalised over in each case.

It can be seen from Figs. 7.6 and 7.7 that the distribution predicted by the optimal model for each data variable well reproduces the corresponding empirical distribution for that variable, indicating good agreement between the optimal model and the data. Fig. 7.6 shows that the form of the radial density distribution f_R used for the model was a reasonable choice. The small differences between the model distribution and the empirical distribution for the cumulative number with respect to R seen in Fig. 7.6 are well balanced across the range of R values, so allowing the King radius (or tidal

radius) to vary as another parameter of the model is not likely to provide much improvement to the fit. Fig. 7.7 shows how well the optimal cooling model fits the data. Note that the cumulative number is shown on a linear scale in Fig. 7.6, whereas a logarithmic scale is used for the cumulative number in each of the plots shown in Fig. 7.7. The minor differences between the model and the empirical distributions with respect to magnitude seen in Fig. 7.7 occur at smaller magnitudes where there are fewer white dwarfs. These early (small magnitude) regions of the distributions in Fig. 7.7 are emphasised in the plots due the logarithmic scale used for the cumulative number, but they are actually less important to the fit than the later (larger magnitude) regions where there are much more white dwarfs and where the model distributions very closely match the empirical distributions.

As a quick and simple quantitative test of the goodness of fit of the optimal model, we perform a set of one-sample Kolmogorov-Smirnov (KS) tests for each of the one-dimensional marginal distribution functions predicted by the optimal model for each of the two data sets used in the combined fit. The results of the KS tests for each of the cases considered are summarised in Table 7.4. The pairs of model and empirical distributions compared by these KS tests correspond to the distributions shown in Figs. 7.6 and 7.7 after normalisation. Note that the KS test actually compared the cumulative probability density distributions (measured by the cumulative fraction), rather than the (inverse) cumulative number distributions shown in Figs. 7.6 and 7.7. The p -value returned by the (one-sample) KS test is the probability that, for a sample drawn from the model distribution, the resultant sample distribution differs from the model distribution by at least as much as the observed empirical distribution differs from the model distribution. The large p -values found for the KS tests comparing the optimal model to the data (see Table 7.4) are indicative of a good fit. These p -values are all well above a reasonable threshold value of 10^{-4} . The KS test results are particularly good for the ACS data, and are also quite reasonable for the WFC3 data.

7.6. Discussion

Data Set	Data Variable	<i>p</i> -value
WFC3/UVIS	<i>R</i>	0.0156
	F225W	0.0362
	F336W	0.0132
ACS/WFC	F435W	0.3527
	F555W	0.2739

Table 7.4: Results of KS tests for combined analysis comparing the one-dimensional marginal cumulative probability distribution functions predicted by the optimal model to the corresponding empirical (cumulative fraction) distribution.

7.6 Discussion

The constraints found in this chapter could potentially be improved upon by applying non-uniform priors for the white dwarf birthrates in the unbinned likelihood analysis, but determining appropriate prior values for these birthrates requires further analysis that is beyond the scope of this work. The birthrate priors used by Goldsbury *et al.* [215], which were calculated from red giant branch stars that have just left the main sequence, seem to simultaneously overestimate the birthrate for the WFC3 field and underestimate it for the ACS field. Goldsbury *et al.* [215] used a WFC3 birthrate prior of $\dot{N}_{\text{WFC3}} = 8.2 \pm 0.3$ Myr and an ACS birthrate prior of $\dot{N}_{\text{ACS}} = 2.61 \pm 0.07$ Myr. In contrast, the optimal model found in our work has a WFC3 birthrate of $\dot{N}_{\text{WFC3}} = 6.91_{-0.23}^{+0.82}$ Myr, which is smaller than the corresponding Goldsbury *et al.* [215] prior value, and an ACS birthrate of $\dot{N}_{\text{ACS}} = 3.73_{-0.24}^{+0.62}$ Myr, which is larger than the corresponding Goldsbury *et al.* [215] prior value. A similar difference between the posterior and prior birthrate values can also be seen in the results of Goldsbury *et al.* [215], though it was not discussed in that work. From the one-dimensional posterior distributions plotted in Goldsbury *et al.* [215], it can be seen that the posterior distribution for the WFC3 birthrate is concentrated about smaller values than the corresponding prior value while the distribution for the ACS birthrate is concentrated about larger values than the corresponding prior

7.6. Discussion

value, with the average posterior value in each case differing by more than 3σ from the prior.

This may indicate that cluster relaxation has a measurable effect that needs to be accounted for when determining the white dwarf birthrate. Stars move away from the centre of a star cluster over time due to diffusion, a phenomena that was observed for 47 Tuc by Heyl *et al.* [261, 262]. Stars leaving the inner WFC3 field over time should cause \dot{N}_{WFC3} to decrease over time. The stars that leave the WFC3 field then enter the ACS field, which would cause \dot{N}_{ACS} to increase (as long as there are more stars entering than leaving the field). This effect could potentially be accounted for by modelling the change in \dot{N}_{WFC3} and \dot{N}_{ACS} over time as part of the analysis, in which case these parameters would really be an “effective” birthrate equal to the net rate of white dwarf formation minus the rate of white dwarfs leaving (or plus the rate of white dwarfs entering) the field of view. However, this is an additional complication that requires careful consideration, so it is left for future work.

Using horizontal branch stars instead of red giant branch stars could also potentially give a better estimate of the white dwarf birthrate, since horizontal branch stars are closer to reaching their white dwarf stage. However, since axions affect the length of time that stars spend on the horizontal branch, this effect would need to be accounted for in the stellar evolution models used to calculate the prior. While axions can also have an effect on stellar evolution along the red giant branch, this effect is less important at the start of the red giant branch (instead being more important for the tip of the red giant branch).

Stronger constraints could also potentially be achieved if the ACS data space could be extended to larger magnitudes, but this would require carefully accounting for the SMC. In this chapter, the issue of SMC contamination has been avoided by simply ending the ACS data space before the SMC begins to intersect the 47 Tuc white dwarf sequence. The ACS data space could be extended by modelling the SMC and incorporating the model for the SMC contaminants into the full model used for the analysis. However, as the ACS data alone favours similar cooling model parameters to the WFC3

data alone (see Appendix B), with both data sets favouring large $\log_{10} q_H$ and small m_a values, extending the ACS data space to larger magnitudes would likely only provide a minor improvement in the final credible regions.

Even without any of these potential improvements to the analysis, the bounds derived for the axion-electron coupling (and DFSZ axion mass) are an improvement upon previous bounds reported in the literature. For comparison, a recent summary of bounds (and hints) for various axion couplings (including g_{aee} and $g_{a\gamma\gamma}$) from stellar evolution is given by Di Luzio *et al.* [84] (see especially Table 1 of Di Luzio *et al.* [84]). The new bound on g_{aee} found in our work improves upon the leading bound from white dwarf cooling (derived using the Galactic white dwarf luminosity function) [48], as well as the slightly stronger bound from the tip of the red giant branch [92, 93]. Of particular interest is that our newly derived (DFSZ axion) bound of $m_a \leq 2.85$ meV (at 95% confidence) excludes the favoured range of $m_a \sim 4 - 10$ meV hinted at by the cooling anomaly reported for the white dwarf luminosity functions of the Galactic disc and halo [47, 49].

7.7 Conclusions

If axions exist and couple to electrons, then axions can be produced in the electron-degenerate core of a white dwarf via axion bremsstrahlung from electrons scattering off ions. The emission of these axions from the white dwarf would provide an extra energy loss mechanism that would affect the rate of white dwarf cooling. For a sufficiently large axion-electron coupling strength, this would have an observable effect on the shape of the white dwarf luminosity function, and hints of axions affecting the rate of white dwarf cooling have been suggested based on a cooling anomaly observed in Galactic white dwarf luminosity functions [47, 49].

In this chapter, we analysed the cooling of white dwarfs in the globular cluster 47 Tuc to search for indirect evidence of axions produced via axion bremsstrahlung from electrons in the white dwarf interiors. White dwarf cooling models were created by using the stellar evolution software MESA to perform simulations of white dwarf cooling that accounted for the energy loss

7.7. Conclusions

associated with the emission of axions produced in the white dwarf. A suite of such models was created for a grid of parameter values that varied the axion-electron coupling strength, the white dwarf mass, and the thickness of the H envelope of the white dwarf. These cooling models were compared to HST observations of young white dwarfs, for which axion emission is expected to be the dominant cooling mechanism, by performing an unbinned likelihood analysis similar to that of Goldsbury *et al.* [215] and Chapter 6.

The white dwarf mass and envelope thickness are nuisance parameters in the analysis which were constrained using prior information from Chapter 6. The results of Chapter 6, which analysed the cooling of older white dwarfs in 47 Tuc using a separate set of HST data that included white dwarfs old enough for their envelopes to have become convectively coupled to their cores, are particularly important for constraining the envelope thickness. There is some degeneracy between the axion-electron coupling constant and the envelope thickness, so including this prior information for the envelope thickness in particular is important for constraining the value of the coupling constant. From the analysis performed in this chapter, we find that models with thick envelopes and no axion emission are favoured.

We find that the axion-electron coupling constant is constrained to be $g_{aee} \leq 0.81 \times 10^{-13}$ at 95% confidence. This bound applies for any axion or axion-like particle model that includes a coupling to electrons. For a DFSZ model, this corresponds to a constraint on the axion mass (and angular parameter) of $m_a \sin^2 \beta \leq 2.85$ meV. The bound on g_{aee} found in this work is more stringent than previous astrophysical bounds from both the Galactic white dwarf luminosity function [48] and the tip of the red giant branch of globular clusters [92, 93]. We do not see evidence of any cooling anomaly in the cooling of 47 Tuc white dwarfs, and the updated bound on g_{aee} excludes the values favoured for the hint of axions suggested by the white dwarf luminosity functions of the Galactic disc and halo [49].

Chapter 8

Prospective axion search using ZTF J1901+1458

8.1 Introduction

The X-ray emission from magnetic hot white dwarfs may reveal evidence for axions or axion-like particles [53, 54], which have been a major focus of studies to go beyond the Standard Model and to explain dark matter [265–267]. The QCD axion, proposed to solve the strong CP problem [73, 74, 77, 78], is a well motivated addition to the Standard Model. Axion-like particles, which are pseudo-scalar particles with properties similar to the QCD axion but that do not necessarily relate to the strong CP problem, also arise naturally in many other extensions to the Standard Model, such as compactified string theories [268–275].

Interactions of axions with photons and nucleons are generic features of QCD axion models, including the benchmark KSVZ [79, 80] and DFSZ [81, 82] models, and are common features of axion-like particle models (see e.g. Di Luzio *et al.* [85] for a recent review). Many axion models also include a coupling of axions to electrons, including the DFSZ model.

Axions couple to photons through the Lagrangian density interaction term $\mathcal{L} \supset -g_{a\gamma\gamma}aF\tilde{F}$, where $g_{a\gamma\gamma}$ is the axion-photon coupling constant, a is the axion field, and F is the electromagnetic field strength tensor. The axion interaction with a fermion species occurs through the operator $\mathcal{L} \supset -ig_{aff}a\bar{f}\gamma_5f$, where g_{aff} is the axion-fermion coupling constant and f is the fermion field, such as for electrons, protons, or neutrons (i.e. $f = e, p, n$). An effective coupling constant g_{aNN}^{eff} can also be defined for the axion interaction

8.1. Introduction

with the nucleon doublet field $N = (p, n)^T$ in terms of the axion-proton and axion-neutron couplings (g_{app} and g_{ann} , respectively). For axion production through the ^{57}Fe nuclear transition, which is the focus of this chapter, it is convenient to adopt the definition $g_{aNN}^{\text{eff}} = 0.16g_{app} + 1.16g_{ann}$ for this coupling [56].

The axion couplings to nucleons, photons, and electrons each have a model-dependent relation to the axion mass, m_a . To summarize, typical values of the QCD axion couplings for the DFSZ (type I) model are³⁶ [85, 276, 277]

$$g_{aNN}^{\text{eff}} = 1.6 \times 10^{-8} (4.3 \sin^2 \beta - 2.2) \left(\frac{m_a}{1 \text{ eV}} \right), \quad (8.1)$$

$$g_{a\gamma\gamma} = 1.5 \times 10^{-10} \left(\frac{m_a}{1 \text{ eV}} \right) \text{ GeV}^{-1}, \quad (8.2)$$

$$g_{aee} = 3.0 \times 10^{-11} \sin^2 \beta \left(\frac{m_a}{1 \text{ eV}} \right), \quad (8.3)$$

where β is a parameter of the DFSZ model defined in the perturbative domain $\tan \beta \in [0.25, 170]$. The axion-nucleon and axion-photon couplings for the KSVZ model are similar (typically within a factor of a few) to those of the DFSZ model, while axion-like particles can have couplings much larger than either the KSVZ or DFSZ benchmark models [85]. For both benchmark QCD axion models, the axion-nucleon coupling is typically three orders of magnitude larger than the DFSZ axion-electron coupling. As X-ray observations of white dwarfs searching for evidence of axions have focused on axions produced through the axion-electron interaction, the greater strength of the axion-nucleon interaction highlights the potential power of similar X-ray searches for evidence of axions produced through nuclear interactions.

Astrophysical observations are often used to search for signatures indicative of the various possible axion interactions with Standard Model particles and to constrain the axion coupling constants. The axion-nucleon coupling

³⁶These expressions use the updated relation between the axion mass and symmetry breaking scale summarized in e.g. section 2.8 of Di Luzio *et al.* [85] instead of the relation used by Raffelt [51] to derive Eqs. (7.1) and (7.2) used for $g_{a\gamma\gamma}$ and g_{aee} in Chapter 7. Using this updated relation instead results in only very minor differences in the values of the coefficients in these equations.

8.1. Introduction

has been probed indirectly through the observed neutrino emission from SN 1987A [51, 278–283] and the cooling of neutron stars [284–289]. Helioscope experiments, which search for axions produced in the Sun, have also been used to constrain the axion-nucleon coupling. Axions can be produced in the Sun through the decay of excited nuclear states such as the first excited state of ^{57}Fe . The CERN Axion Solar Telescope (CAST [290, 291]) has searched for axions produced in this way, setting current constraints for the axion-nucleon coupling [55, 56]. A similar search has been proposed for the future planned International Axion Observatory (IAXO [56, 292, 293]), along with improved calculations for the axion flux from the ^{57}Fe transition [56] using updated nuclear matrix elements [294].

White dwarfs are another popular target of searches for axions produced from astrophysical sources and have typically been used to probe the axion-electron coupling, which enables the production of axions within a white dwarf through axion bremsstrahlung. This extra source of energy loss would modify the cooling of white dwarfs and thus the white dwarf luminosity function, which in comparison to the observed luminosity function has been used to constrain the axion-electron coupling (e.g. [47, 48, 215, 295, 296] and Chapter 7 of this dissertation). Furthermore, in the strong field surrounding magnetic white dwarfs, axions can convert to photons and vice versa. Axions produced by bremsstrahlung which then convert to photons in the surrounding magnetic field produce a blackbody-like spectrum in the X-ray, which was the focus of a 100 ks observation of the magnetic white dwarf RE J0317-853 [54] with *Chandra*.

In this chapter, we show that hot, highly magnetized white dwarfs are also ideal targets to probe the axion-nucleon coupling via the ^{57}Fe transition, and searches for the associated signal with an X-ray telescope can in fact probe this coupling with better sensitivity than both current and planned future helioscope searches for axions produced in the Sun through nuclear transitions. If the temperature of the white dwarf is sufficiently high, low-lying nuclear states may be excited within the star, and these states may decay through the emission of an axion. Among all of the low-lying nuclear states, the excited state of ^{57}Fe at 14.4 keV stands out through the

combination of low energy and relatively large abundance within the star [297]. In general, these axions would stream out of the star unimpeded and unnoticed. However, in the case of a strongly magnetized white dwarf like ZTF J1901+1458 [57], these axions have a chance to transform into X-ray photons at 14.4 keV as they pass through the magnetosphere of the white dwarf. White dwarfs are ideal targets for measurements of this signature because their thermal emission in the hard X-rays is negligible, and therefore they provide the chance for a very clean detection. We demonstrate the potential sensitivity of X-ray searches for the ^{57}Fe axion signal from magnetic white dwarfs by calculating the projected sensitivity to the axion-nucleon and axion-photon couplings that could be obtained with a 100 ks observation of ZTF J1901+1458 by *NuSTAR*.

8.2 Calculations

The spectrum of axions produced through the ^{57}Fe nuclear transition is a narrow peak at the nuclear excitation energy, $E^* = 14.4$ keV. The resultant photon number flux at Earth induced by axions being produced via the nuclear transition process in the core of a white dwarf and then converting to photons in the magnetosphere is [53]

$$\Phi_\gamma = \mathcal{N}_a M_{\text{WD}} \times p_{a \rightarrow \gamma} \times \frac{1}{4\pi d_{\text{WD}}^2}, \quad (8.4)$$

where \mathcal{N}_a is the emission rate per unit mass of axions produced in the white dwarf interior (which we consider to be isothermal), M_{WD} is the white dwarf mass, $p_{a \rightarrow \gamma}$ is the probability of an axion with energy E^* converting into a photon in the magnetosphere of the white dwarf, and d_{WD} is the distance to the white dwarf from the point of observation. The product $\mathcal{N}_a M_{\text{WD}}$ is the mass averaged number of axions produced by the white dwarf per unit time, and multiplying this quantity by E^* yields the luminosity of axion production under the assumption of an isothermal white dwarf core.

Modeling the core of the white dwarf as isothermal, the number density of axions produced from the ^{57}Fe nuclear transition in the white dwarf core

8.2. Calculations

is [56]

$$\mathcal{N}_a = \mathcal{N}_{57} \omega_1(T_c) \frac{1}{\tau_0} \frac{1}{(1+\alpha)} \frac{\Gamma_a}{\Gamma_\gamma}, \quad (8.5)$$

where \mathcal{N}_{57} is the ^{57}Fe number density per unit mass of stellar matter, ω_1 is the occupation number of the first excited state at the temperature T_c of the white dwarf core, τ_0 is the lifetime of the excited state, α is the internal conversion coefficient, and Γ_a/Γ_γ is the branching ratio of axion emission relative to photon emission.

The axion emission rate depends on the axion-nucleon coupling through the term Γ_a/Γ_γ . For axions produced by the ^{57}Fe transition, the axion flux from the Sun has recently been calculated by Di Luzio *et al.* [56] using the updated nuclear matrix elements of Avignone *et al.* [294]. This yielded an updated axion-to-photon branching ratio for ultrarelativistic axions of

$$\frac{\Gamma_a}{\Gamma_\gamma} = 2.32 \left(g_{aNN}^{\text{eff}} \right)^2, \quad (8.6)$$

where g_{aNN}^{eff} is the effective axion-nucleon coupling constant, defined as $g_{aNN}^{\text{eff}} = 0.16g_{app} + 1.16g_{ann}$ in terms of the axion couplings to protons and neutrons [56].

The core temperature of the white dwarf is another important parameter for the calculation of \mathcal{N}_a , as the occupation number ω_1 is temperature-dependent. We determined the core temperature of ZTF J1901+1458 using the published photometry and fitting technique of Caiazzo *et al.* [57], with T_c used as one of the free parameters instead of the effective temperature T_{eff} . We related T_c to the photon luminosity L_γ (and thus T_{eff} and radius) using [53]

$$kT_c \simeq (0.3 \text{ keV}) \left(\frac{L_\gamma}{10^{-4} L_\odot} \right)^{0.4}, \quad (8.7)$$

where k is the Boltzmann constant. The results of the joint fit for T_c along with the radius of the white dwarf, R_{WD} , and colour excess due to interstellar reddening, $E(B-V)$, are shown in Fig. 8.1. Based on these results, we use a core temperature of $kT_c = 2.85_{-0.63}^{+1.58}$ keV for our calculation of the axion emission rate.

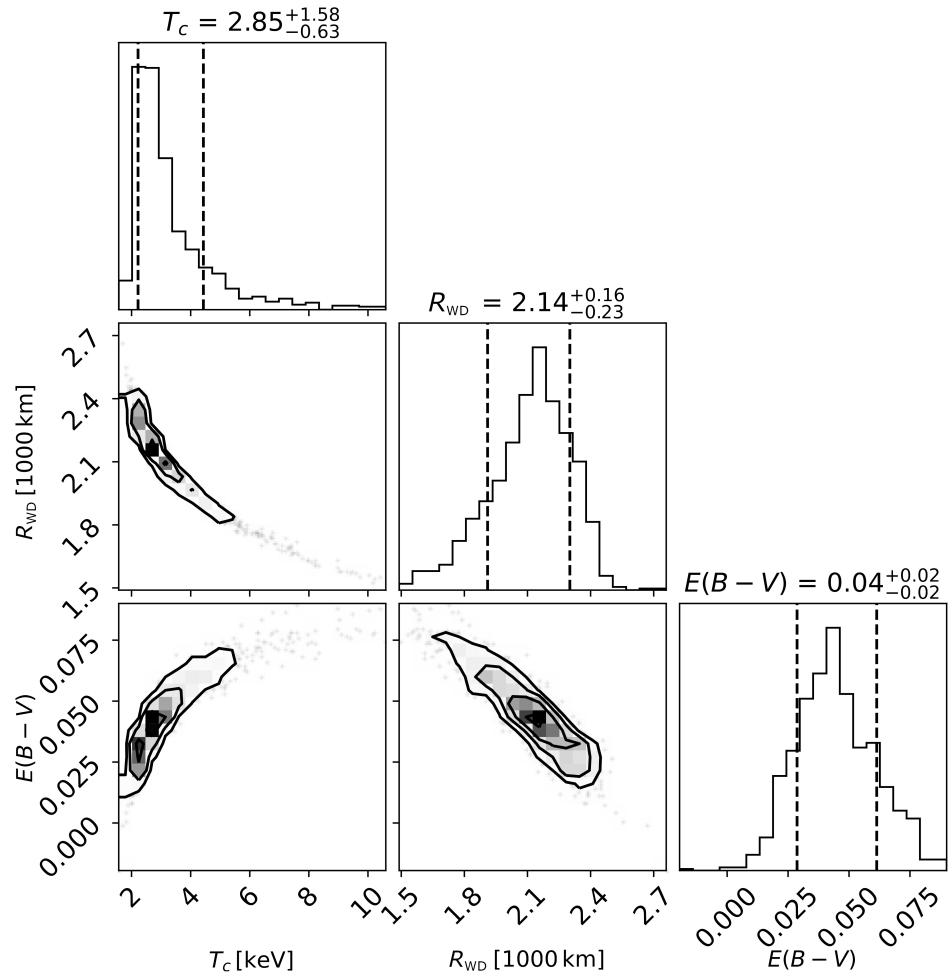


Figure 8.1: Constraints on the core temperature of ZTF J1901+1458 from the published photometry [57].

8.2. Calculations

The thermal occupation number of an excited state with excitation energy E^* as a function of temperature T is $\omega_1 = (2J_1+1)e^{-E^*/kT} / [(2J_0+1)+(2J_1+1)e^{-E^*/kT}]$, where J_0 and J_1 are the angular momenta of the ground and excited state, respectively [56]. These angular momenta are $J_0 = 1/2$ and $J_1 = 3/2$ for the ^{57}Fe ground and first excited state, respectively [298], giving a thermal occupation number of $\omega_1 = 2e^{-E^*/kT} / (1 + 2e^{-E^*/kT})$. This further simplifies to approximately $\omega_1 \sim 2e^{-E^*/kT_c}$ for the core temperature of ZTF J1901+1458 and an isothermal core.

A large core temperature can also broaden the width of the expected axion signal. In the case of the Sun, the energy of the axion is broadened by the thermal motion of the nuclei, yielding a width of a few eV [56]. As the temperature in the core of the white dwarf is larger than in the Sun, so is the thermal broadening (5 eV for ZTF J1901+1458); however, the varying gravitational potential through the core has a larger effect through the gravitational redshift, which is about 280 km/s at the surface and 980 km/s at the centre, yielding a width of 33 eV that dominates over the thermal effects. These broadening effects are negligible compared to the spectral resolution of *NuSTAR* (≈ 400 eV), so we work in terms of the total flux rather than the spectral flux and approximate that all of the axions are emitted with the same energy E^* .

The values for all of the other parameters used in the calculation of \mathcal{N}_a are taken from the literature. For the parameters characterizing the first excited state of ^{57}Fe , we use values from Röhlsberger [298] of $\tau_0 = 141$ ns and $\alpha = 8.56$. To get the number density of ^{57}Fe nuclei, we use the solar elemental abundances reported by Lodders [297]. We use a proto-solar hydrogen mass fraction of 0.71 and a number fraction of ^{57}Fe nuclei relative to protons of 7.34×10^{-7} (see Table 9 of Lodders [297]). The corresponding ^{57}Fe number density per unit mass is $\mathcal{N}_{57} = 6.24 \times 10^{50} M_{\odot}^{-1}$. As the observations of Caiazzo *et al.* [57] indicate that ZTF J1901+1458 is young (with its cooling age estimated to be only 10–100 Myr), is located in the solar neighbourhood (at about 40 pc), and has kinematics consistent with the Galactic disc, we conclude that, like other stars in the local disc, the ZTF J1901+1458 progenitor abundances were similar to the solar abundances. Furthermore, we

8.2. Calculations

performed stellar evolution simulations using MESA to verify that stellar evolution processes do not alter the iron abundance in the stellar core from the initial value. Finally, for the white dwarf parameters needed in the rest of the flux calculation, we use the values for ZTF J1901+1458 reported in Caiazzo *et al.* [57], for which the mass is $M_{\text{WD}} = 1.346 \pm 0.019 M_{\odot}$ and the distance is $d_{\text{WD}} = 41.4 \pm 0.1$ pc.

To determine the observable photon flux Φ_{γ} that would be induced by axions produced at the rate \mathcal{N}_a in the white dwarf, we must also calculate the probability of the axions converting into photons as they propagate outward through the magnetic field surrounding the white dwarf. Under the approximation that the axions travel along radial trajectories (relative to the centre of the white dwarf), the propagation of the axion-photon field is described by the equations [299]

$$\left[i\partial_r + E + \begin{pmatrix} \Delta_{\parallel} & \Delta_B \\ \Delta_B & \Delta_a \end{pmatrix} \right] \begin{pmatrix} A_{\parallel} \\ a \end{pmatrix} = 0, \quad (8.8)$$

$$\Delta_{\parallel}(r) = (7/2) (\alpha_{\text{EM}}/45\pi) E [B(r)/B_{\text{crit}}]^2 \sin^2 \Theta, \quad (8.9)$$

$$\Delta_B(r) = (1/2) g_{a\gamma\gamma} B(r) \sin \Theta, \quad (8.10)$$

$$\Delta_a = -m_a^2/(2E), \quad (8.11)$$

where r is the radial coordinate, $B(r)$ is the magnetic field strength, $a(r)$ is the axion field, and $A_{\parallel}(r)$ is the component of the electromagnetic vector potential corresponding to the parallel photon polarization (i.e. the polarization such that the electric field of the wave has a component parallel to the stellar magnetic field). The parameter $E \approx E^*$ is the axion energy, $\alpha_{\text{EM}} \approx 1/137$ is the electromagnetic fine structure constant, $B_{\text{crit}} = 4.414 \times 10^{13}$ G is the quantum electrodynamics (QED) critical field strength, and Θ is the angle between the magnetic field and the radial propagation direction. For details on the origin and solution of the axion-photon propagation equations for magnetic stars, see e.g. Raffelt and Stodolsky [299] or Lai and Heyl [300].

8.2. Calculations

The expression for Δ_{\parallel} describes strong-field QED effects in vacuum. In general, Δ_{\parallel} includes an additional term $-\omega_{\text{pl}}^2/(2E)$ arising from plasma effects, where the plasma frequency is $\omega_{\text{pl}} = 28.71\sqrt{Y_e \rho / (1 \text{ g cm}^{-3})}$ eV for an electron plasma [300], with Y_e the electron fraction and ρ the mass density. Modeling the white dwarf as surrounded by an isothermal plasma of fully ionized hydrogen ($Y_e = 1$) with the density profile from Gill and Heyl [301], of the form $\rho(r) = \rho_0 \exp[-(r - R_{\text{WD}})/H_{\rho}] + \rho_{\infty}$, the density drops from $\rho_0 = 10^{-10} \text{ g cm}^{-3}$ at the surface to $\rho_{\infty} = 10^{-20} \text{ g cm}^{-3}$ far from the white dwarf with a density scale height of $H_{\rho} = (2kT_{\text{eff}})/(m_p g_0)$, where m_p is the proton mass and g_0 is the surface gravity. For ZTF J1901+1458, with $g_0 = 4.0 \times 10^9 \text{ cm s}^{-2}$ and $T_{\text{eff}} = 4.6 \times 10^4 \text{ K}$ [57], the plasma frequency drops to $\omega_{\text{pl}} \approx 2.9 \times 10^{-9} \text{ eV} \ll E$, m_a with scale height $H_{\rho} \approx 0.019 \text{ km} \ll R_{\text{WD}}$. Thus, the plasma term is negligible compared to the vacuum QED and Δ_a terms for the scenario we consider.

To calculate the probability of an axion converting to a photon in the magnetosphere, we solve Eq. (8.8) numerically for an initial pure axion state. The probability $p_{a \rightarrow \gamma}$ is given by the squared magnitude of the asymptotic solution for A_{\parallel} . For this calculation, we model the external magnetic field of the white dwarf as a magnetic dipole, $B(r) = B_0 (R_{\text{WD}}/r)^3$, where B_0 is the magnetic field at the surface of the white dwarf. For this form of the magnetic field, $\sin \Theta$ has a fixed value along the radial trajectory, which we take to be unity. The values of the relevant magnetic field parameters for ZTF J1901+1458 are $B_0 \sim 800 \text{ MG}$ and $R_{\text{WD}} = 2,140_{-230}^{+160} \text{ km}$ [57].

As Θ always appears in Eq. (8.8) through the factor $B_0 \sin \Theta$, setting $\sin \Theta = 1$ effectively makes B_0 the transverse magnetic field at the surface, $B_{T,0} \equiv B_0 \sin \Theta$. For magnetic field strengths relevant to this work, $p_{a \rightarrow \gamma}$ is approximately $\propto B_{T,0}^{2/5}$ (Dessert *et al.* [53], Supplemental Material). This dependence is much weaker than the exponential dependence of the corresponding photon flux on T_c . As the core temperature is the dominant source of uncertainty in our sensitivity estimate, more detailed modeling of the magnetic field is beyond the scope of this work.

The key parameter values used to calculate the X-ray photon flux at Earth are summarized in Table 8.1. The observable photon flux induced

8.2. Calculations

Parameter	Value
E^*	14.4 keV
J_0	1/2
J_1	3/2
τ_0	141 ns
α	8.56
\mathcal{N}_{57}	$6.24 \times 10^{50} M_{\odot}^{-1}$
M_{WD}	$1.34 M_{\odot}$
d_{WD}	41.44 pc
B_0	800 MG
R_{WD}	2,100 km
kT_c	$2.85^{+1.58}_{-0.63}$ keV

Table 8.1: Parameter values used to calculate the X-ray flux induced by axions produced in the ^{57}Fe nuclear transition process for ZTF J1901+1458. The properties of the first excited state of ^{57}Fe come from Röhlsberger [298]. The abundances used to determine \mathcal{N}_{57} come from Lodders [297]. The ZTF J1901+1458 white dwarf parameters come from Caiazzo *et al.* [57], except for T_c which is calculated in this work.

by the ^{57}Fe transition depends on both the axion-nucleon coupling and the axion-photon coupling. The emission rate of axions produced by the ^{57}Fe nuclear transition in the white dwarf interior goes as $(g_{aNN}^{\text{eff}})^2$, while the probability of the axions converting into photons in the external magnetic field goes as $(g_{a\gamma\gamma})^2$. Thus, observations of the hard X-ray emission from magnetic white dwarfs are sensitive to the product of the couplings: $g_{aNN}^{\text{eff}} g_{a\gamma\gamma}$. Furthermore, the probability of an axion converting to a photon depends on the axion mass, so the sensitivity to $g_{aNN}^{\text{eff}} g_{a\gamma\gamma}$ of such X-ray observations will be a function of axion mass, which we evaluate numerically over a grid of axion mass values.

We calculate the sensitivity to $g_{aNN}^{\text{eff}} g_{a\gamma\gamma}$ that could be achieved from *NuSTAR* observations with a total exposure time of 100 ks. Using XSPEC simulations, we determined the minimum flux for a three-sigma detection of a narrow (width of 0.03 keV), Gaussian spectral line at 14.4 keV above the background (30-arcsecond extraction region and one arcminute off-axis) to

8.3. Results

be 1.6×10^{-6} photons s $^{-1}$ cm $^{-2}$. In the background, we also included the thermal emission from the white dwarf atmosphere modeled as a blackbody at a temperature $kT_{\text{eff}} = 3.88$ eV with a radius of 2,100 km. For a given axion mass, the value of $g_{aNN}^{\text{eff}} g_{a\gamma\gamma}$ for which the photon number flux given by Eq. (8.4) equals (or exceeds) the detector threshold sets the limit defining the sensitivity. In the case of a non-detection, this sensitivity would correspond to a constraint on $g_{aNN}^{\text{eff}} g_{a\gamma\gamma}$ (as a function of m_a) set by the observations.

8.3 Results

The potential *NuSTAR* signature of axion production through nuclear processes is a narrow emission line at 14.4 keV (the excitation energy of the ^{57}Fe nucleus). The projected sensitivity curve for a 100 ks *NuSTAR* observation of the white dwarf ZTF J1901+1458 is depicted in Fig. 8.2, where it is compared to both current constraints and the predicted sensitivity of the future observatory IAXO. For reference, the parameter relations of the benchmark KSVZ and DFSZ (with $\sin \beta \approx 1$) models are also shown. The blue shaded region shows the statistical uncertainty in the projected sensitivity curve due to the measurement error of T_c , where a dipolar magnetic field has been assumed. More complicated magnetic field geometries can additionally shift the projected sensitivity to a lesser extent in either direction through the axion-photon conversion probability (e.g. [53, 301, 302], noting that the sensitivity curve for $|g_{aNN}^{\text{eff}} g_{a\gamma\gamma}|$ is only proportional to the square root of the axion-photon conversion probability). However, the magnetic field structure of ZTF J1901+1458 is currently not well understood, so a more detailed analysis of the field geometry and its associated uncertainty is left for future work.

Compared to helioscope searches for an analogous ^{57}Fe axion signal, a 100 ks *NuSTAR* observation of ZTF J1901+1458 would be sensitive to $|g_{aNN}^{\text{eff}} g_{a\gamma\gamma}|$ values smaller than the existing constraints from CAST for all axion masses and would even be more sensitive than the planned search by the future proposed IAXO by an order of magnitude for small axion masses. The shaded orange region in Fig. 8.2 shows the current CAST constraints

8.3. Results

([55] with correction by [56]), and the solid black curves show the sensitivity limits for the four different configurations of both the proposed intermediate-stage BabyIAXO and the proposed fully operational IAXO (see Tab. I of Di Luzio *et al.* [56] for a summary of these configurations).

These helioscope searches for axions produced by the ^{57}Fe transition are sensitive to the same quantity, $|g_{aNN}^{\text{eff}} g_{a\gamma\gamma}|$, as our proposed search, making them the most directly comparable to our work. However, a current bound that is stronger than the CAST bound from ^{57}Fe axion searches is achieved by combining the SN 1987A constraint on g_{aNN}^{eff} with the independent constraint on $g_{a\gamma\gamma}$ from helioscope searches for axions produced through the (axion) Primakoff process. From an updated analysis of the SN 1987A neutrino signal [281], the bound on the axion-nucleon couplings is $g_{ann}^2 + 0.6g_{app}^2 + 0.5g_{ann}g_{app} \lesssim 8.3 \times 10^{-19}$. Following the procedure of Di Luzio *et al.* [56], this limit can be translated into a bound on g_{aNN}^{eff} by choosing the ratio g_{app}/g_{ann} that minimizes the left hand side while keeping g_{aNN}^{eff} constant ($g_{app}/g_{ann} \approx 0.2$), which gives $|g_{aNN}^{\text{eff}}| \lesssim 1.1 \times 10^{-9}$. This saturated SN 1987A limit for g_{aNN}^{eff} , in conjunction with the CAST Primakoff axion limit for $g_{a\gamma\gamma}$ [303], is shown as the red curve in Fig. 8.2, which falls between the lowest two IAXO curves for $m_a \lesssim 0.02$ eV. Our projected sensitivity would thus even improve upon current bounds from SN 1987A (and CAST) for small axion masses.

In addition to probing the combination of couplings $g_{aNN}^{\text{eff}} g_{a\gamma\gamma}$, X-ray observations of ZTF J1901+1458 could also probe the combination $g_{aee} g_{a\gamma\gamma}$ as white dwarfs can produce axions through electron bremsstrahlung, which yields a blackbody-like spectrum of axions. This was the focus of a 100 ks observation of a cooler magnetic white dwarf, RE J0317-853 [54], with the X-ray telescope *Chandra*. These axions are typically produced at lower energies (near the temperature of the core at a few keV), so the resulting X-rays lie squarely in the energy range probed by *Chandra*, but a search could be performed using *NuSTAR* observations as well.

ZTF J1901+1458 is a prime target to search for axion signatures in the X-rays because it is one of the hottest and most strongly magnetized white dwarfs known. Both of these properties increase the predicted strength of

8.3. Results

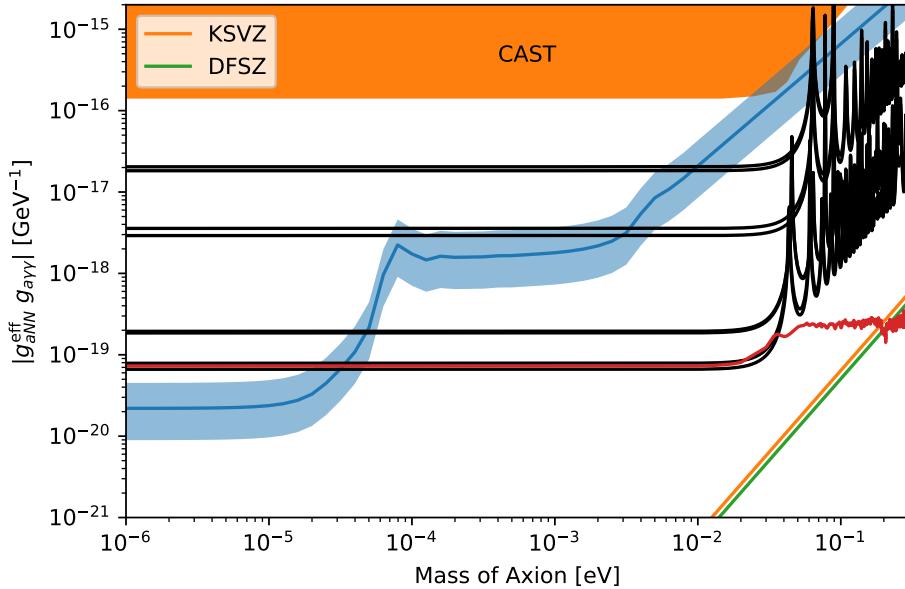


Figure 8.2: Projected sensitivity to the product of the axion-nucleon and axion-photon couplings that would be obtained with a 100 ks *NuSTAR* observation of ZTF J1901+1458 in blue. The band accounts for the uncertainty in the core temperature of the white dwarf. The shaded orange region is excluded by CAST ^{57}Fe axion searches [55, 56]. The black curves give sensitivity limits predicted for ^{57}Fe axion searches with the future IAXO [56] for both the intermediate-stage BabyIAXO (upper four) and the fully operational IAXO (lower four). The red curve shows the SN 1987A bound on g_{aNN}^{eff} [281] in combination with the current CAST Primakoff axion bound on $g_{a\gamma\gamma}$ [303].

8.4. Conclusions

the ^{57}Fe nuclear transition line and bremsstrahlung emission, and therefore the chance of detection, or the significance of the constraints in the case of a non-detection. Recently published photometry and spectroscopy [57] yield somewhat broad constraints on the core temperature of the white dwarf (see Fig. 8.1) that result in a broad range for the implied sensitivity curve of the axion search (the shadowed blue area in Fig. 8.2). Fortunately, follow-up ultraviolet spectroscopy observations are scheduled for the current HST cycle [304] that should provide stronger constraints on the effective temperature and therefore the core temperature of the white dwarf, which would narrow the band for the sensitivity that could be achieved with *NuSTAR* observations.

8.4 Conclusions

White dwarfs are a popular target of indirect searches for axions. While observations of white dwarfs have been used to constrain the coupling of axions to electrons, white dwarfs have not previously been identified as a target for probing the axion coupling to nucleons. In this chapter, we have shown for the first time that observations of white dwarfs can be used to probe the axion-nucleon coupling through searches for the X-ray signal from hot, highly magnetized white dwarfs that would be induced by the process of axions being produced in the core via the nuclear transition of the first excited state of ^{57}Fe and then converting into photons in the magnetosphere.

The recently discovered white dwarf ZTF J1901+1458 is a compelling target to search for an X-ray signal arising from axions produced via the ^{57}Fe transition. We have shown that a 100 ks observation of ZTF J1901+1458 by *NuSTAR* would probe the coupling of axions with nucleons and photons at a level below current helioscope constraints for all axion masses and would furthermore be more sensitive than what is predicted for planned future terrestrial experiments at small masses. This would provide a dramatic improvement in our knowledge of these particles that are critical to our understanding of the Standard Model and possibly dark matter as well.

Chapter 9

Conclusions

This dissertation has demonstrated various ways in which observations of white dwarfs can be used as probes of new phenomena in physics, particularly as tests of hypothesised particles such as axions in models beyond the Standard Model of particle physics. The statistical analysis of distributions constructed from large samples of white dwarfs is especially powerful for this purpose. If white dwarf cooling is well understood, then the white dwarf luminosity function in particular provides a good test of new physics that can affect the rate of white dwarf cooling, such as axion models that include an axion-electron coupling. A slower rate of white dwarf cooling over a particular luminosity range, as would be induced by a cooling delay, manifests as a larger number of white dwarfs over that range, while a faster rate of cooling, as would be induced by the emission of extra particles such as axions, manifests as a smaller number of white dwarfs. Observations of the emission spectra of individual white dwarfs can also be useful in looking for evidence of new physics, and in particular such observations can be used to probe axion couplings that the cooling rates of most white dwarfs are not sensitive to.

While white dwarf cooling is generally thought to be well-understood, some *Gaia* DR2 observations of ultramassive white dwarfs in the solar neighbourhood challenged the conventional understanding of white dwarf cooling, and in particular the understanding of core crystallisation. Based on *Gaia* DR2 observations, Cheng *et al.* [43] noted an excess of ultramassive white dwarfs on the Q branch, where these white dwarfs are in the process of undergoing core crystallisation, and furthermore found that a small fraction of ultramassive white dwarfs on the Q branch were moving faster than expected based on the young ages inferred from their photometry using cooling

models. Cheng *et al.* [43] argued that this could only be explained by an anomalous cooling delay, with $\sim 6\%$ of $\sim 1.08 - 1.23 M_{\odot}$ white dwarfs experiencing an extra cooling delay of ~ 8 Gyr, and suggested modification to the treatment of ^{22}Ne sedimentation as a possible mechanism for this delay. This cooling delay explanation argues that the anomalously fast-moving white dwarfs are actually much older than suggested by their photometric ages, as the velocities of stars in the local Galactic disc have been observed to become more dispersed over time, and that the transverse velocities measured for the fast-movers are actually what is expected based on their true ages according to the AVR_s of local disc stars. These anomalous features of ultramassive white dwarfs in the solar neighbourhood and the corresponding cooling delay explanation were re-investigated in Chapters 4 and 5 using the improved *Gaia* EDR3 observations. New features not previously considered were also analysed in these chapters.

In Chapter 4, the number density distribution of white dwarfs as a function of cooling age and mass was analysed for *Gaia* EDR3 white dwarfs within 200 pc of the Sun with masses in the range $0.95 - 1.25 M_{\odot}$. The cooling age and mass of each white dwarf were inferred from the photometry using white dwarf cooling models, and the full analysis was done for a variety of different publicly available cooling models. In each case, the sample was subdivided into the three equally-spaced mass bins $0.95 - 1.05$, $1.05 - 1.15$, and $1.15 - 1.25 M_{\odot}$, and the empirical distribution of cooling ages was constructed for each bin. For each of these mass bins, it was found that the number density increases over the first ~ 3 Gyr of cooling time. If a uniform birthrate is assumed, then this corresponds to an apparent excess number of white dwarfs both on the Q branch and below it, and the excess reaches a peak after core crystallisation has already ended. Furthermore, this peak coincides with a burst of star formation that occurred in the Galaxy $\sim 2 - 3$ Gyr ago, according to the star formation history determined from *Gaia* DR2 main sequence stars [106].

For white dwarfs of the large masses considered in Chapter 4, the distribution of cooling ages is expected to closely track the star formation rate if the white dwarfs formed through single stellar evolution, and for the lightest

two mass bins, $0.95 - 1.05$ and $1.05 - 1.15 M_{\odot}$, it was found that the empirical distribution of cooling ages is statistically consistent with this expectation for the thick-envelope Montreal models, i.e. the cooling models of Bédard *et al.* [109] with thick H envelopes for DA white dwarfs. For the more massive $1.15 - 1.25 M_{\odot}$ white dwarfs, it was found that the empirical distribution of cooling ages constructed using these same models is well-fitted by a linear combination of the distribution expected for single stellar evolution and the distribution expected for double white dwarf merger products, the latter of which was taken to be the convolution of the star formation rate from Mor *et al.* [106] and the merger delay time distribution from Cheng *et al.* [126], with the best fit given when $\sim 51\%$ of the $1.15 - 1.25 M_{\odot}$ white dwarfs that formed over the last 4 Gyr are the product of double white dwarf mergers. This merger fraction is further supported by the recent empirical evidence presented by Kilic *et al.* [125, 164] indicating a similarly large fraction of double white dwarf merger products amongst ultramassive white dwarfs.

It is interesting to note that the results of Chapter 4 favour cooling models with C/O cores even for white dwarfs more massive than the threshold mass ($\sim 1.05 - 1.15 M_{\odot}$) above which white dwarfs are typically thought to have O/Ne cores. Though simulations favour O/Ne cores at these high masses, they have also shown the potential for ultramassive white dwarfs to have C/O cores, and the composition of ultramassive white dwarfs is still an open question that warrants further examination. After the completion of the work described in Chapter 4, the La Plata group published new cooling models for ultramassive white dwarfs with C/O cores [159], and it would be interesting to redo the analysis of Chapter 4 (and possibly Chapter 5) using these new models. A similar analysis could also be done using custom MESA models for more control over the composition and other parameters of the cooling models, and it would be interesting to extend that analysis to larger masses than those spanned by the set of publicly available models considered in Chapter 4. More detailed investigation like this into the composition of ultramassive white dwarfs is not necessary for the purposes of this work, but it would be an interesting topic for future research to explore.

Overall, the results of Chapter 4 show that the empirical number density

distribution of cooling ages for massive and ultramassive white dwarfs is consistent with the expected distribution when, in addition to accounting for merger products in the case of the most massive white dwarfs, the time-dependent star formation history is properly accounted for. These results indicate that it is not necessary to invoke an extra ~ 8 Gyr cooling delay to explain the previously reported excess in the number of white dwarfs on the Q branch, and thus that conventional cooling models do not need to be modified to explain this feature.

In Chapter 5, the work of Chapter 4 was followed up with an analysis of the kinematics of the same sample of $0.95 - 1.25 M_{\odot}$ *Gaia* EDR3 white dwarfs within 200 pc of the Sun that was considered in Chapter 4. The sample was subdivided into the same mass bins of $0.95 - 1.05$, $1.05 - 1.15$, and $1.15 - 1.25 M_{\odot}$, and the masses and photometric cooling ages were determined using the best-fitting models from Chapter 4, i.e. the thick-envelope Montreal models [109]. The transverse motions of these white dwarfs were analysed, including in particular the three-dimensional components U_{\perp} , V_{\perp} , and W_{\perp} of the transverse velocity vectors, which had not previously been examined. Further subdividing the mass-binned subsamples into age bins, the empirical distributions of the three-dimensional transverse velocity components were constructed for each mass and age bin, and these distributions were compared to the expectation for stars in the local Galactic disc based on the observed AVR_s of local main-sequence stars.

As was found for *Gaia* DR2 data by Cheng *et al.* [43], a population of photometrically young, ultramassive white dwarfs that are moving anomalously fast was found in the analysis of *Gaia* EDR3 data in Chapter 5. The population of fast-movers found in Chapter 5 is concentrated to the mass range $1.15 - 1.25 M_{\odot}$ and photometric cooling age range $0.5 - 1.5$ Gyr, and the white dwarfs comprising this anomalous population are moving faster than expected, according to observed AVR_s, for stars from the local disc with the young ages inferred from photometry. Contrary to the explanation proposed by Cheng *et al.* [43], however, the results of Chapter 5 show that these fast-movers cannot be explained solely by an extra cooling delay in white dwarfs originating from the local disc. It was found in Chapter 5

that the dispersion of V_{\perp} for the fast-movers is too large to be achieved in the age of the Galactic disc and that the dispersion ratio of V_{\perp} to U_{\perp} is inconsistent with a disc in equilibrium. Both of these features indicate that the fast-movers do not originate from the local Galactic disc and that the explanation of the fast-movers lies in this non-local origin.

Of importance to note is that an extra cooling delay alone cannot explain either of these kinematic features, and even if other measurements of the local AVR or the age of the disc were to find that it is possible to achieve the large dispersion observed for V_{\perp} within the age of the disc, a cooling delay would still not account for the dispersion ratio of the fast-movers. A disc in equilibrium, like the Galactic disc, should have a dispersion ratio of $\sigma_V/\sigma_U \sim 2/3$ for the velocity components V and U of the full motion before projection onto the tangent plane (including both transverse and radial motion), with the dispersion of V expected to be notably smaller than the dispersion of U . A similar relation was shown in Chapter 5 to hold for the transverse velocity components V_{\perp} and U_{\perp} after projection onto the tangent plane; however, it was found empirically that the dispersion of V_{\perp} is actually larger than the dispersion of U_{\perp} for the fast-moving population. A cooling delay does not address this discrepancy, and while it is possible to incorporate an extra cooling delay as part of an explanation of the fast-movers that additionally has some other mechanism to address the kinematic features that the cooling delay does not explain, it is also possible to explain these kinematic features without invoking an extra cooling delay.

Some possible explanations of the fast-moving population, such as a halo origin or production through dynamical effects of stellar triple systems, were discussed in Chapter 5. However, the exact origin of this population is still unknown, and the task of further determining the correct explanation of the fast-movers has been left for future work. To this end, it would be useful to more precisely characterise the kinematic features of this population in the future, for instance by fitting models to the empirical distributions of the three-dimensional transverse velocity components. It would also be interesting to search for other features such as strong magnetic fields, which are associated with mergers, that might characterise the fast-movers and

provide more hints as to their origin.

While determining the precise origin of the fast-moving ultramassive white dwarfs and further characterising this population is an interesting topic for future research, it is beyond the scope of this work. The results of Chapter 5 are sufficient to show that the fast-movers do not originate from the local Galactic disc, which points to the explanation of this population being a question of dynamics rather than a problem with white dwarf cooling models. A cooling delay is not sufficient to explain the key kinematic features highlighted in Chapter 5, and it was furthermore found in Chapter 4 that the cooling age distributions are well-fitted by standard cooling models when the time-dependent star formation rate is properly accounted for, without needing to invoke an extra cooling delay. Neither the cooling age distributions analysed in Chapter 4 nor the distributions of the three-dimensional transverse velocity components analysed in Chapter 5 show specific evidence of an extra cooling delay. The results presented in these chapters can be explained without modifying standard white dwarf cooling models, and in particular do not indicate a need to modify the treatment of core crystallisation in white dwarf cooling models.

Properly determining the parameters of white dwarf cooling models, however, is still an important task if observations of the white dwarf luminosity function are to be used to constrain other exotic physics that would affect white dwarf cooling rates, such as axion models that couple to electrons. This task was the focus of the work presented in Chapter 6, where the main goal was to constrain the parameters that characterise white dwarf cooling in the globular cluster 47 Tuc for later use in the work of Chapter 7 to constrain axion models based on the cooling of white dwarfs in 47 Tuc. In Chapter 6, models of white dwarf cooling without axion emission were fit to observations of white dwarfs in 47 Tuc that are old enough that any potential cooling by axion emission would be negligible. The results of these fits were then used as priors in the analysis of Chapter 7, where models of white dwarf cooling that included energy loss due to axion emission were fit to separate observations of younger white dwarfs in 47 Tuc that are still in the cooling regime where the emission of axions would be a dominant

cooling process.

In Chapter 6, white dwarf cooling models were created by performing stellar evolution simulations using MESA for a parameter grid of different white dwarf mass values, envelope thickness values, and element diffusion scenarios. Three different treatments of element diffusion were considered: i) the standard MESA implementation that approximates the ions as an ideal gas, ii) a custom modified treatment of diffusion to account for non-ideal gas effects that suppress diffusion, and iii) the scenario of no diffusion. These cooling models were compared to deep HST observations of white dwarfs in 47 Tuc by performing an unbinned likelihood analysis similar to that of Goldsbury *et al.* [215]. As the white dwarf cooling sequence of 47 Tuc suffers from SMC contamination for this deep HST data, particularly at the late cooling times that are of the most interest for this analysis, a careful cleaning and associated calibration procedure was performed to remove these SMC contaminants and account for the reduction in completeness associated with this cleaning procedure in the analysis.

The data space used in the analysis performed in Chapter 6 extended to faint enough magnitudes (and thus old enough cooling times) to resolve the bump in the cumulative luminosity function associated with the convective coupling of the white dwarf envelope to the core. The thickness of this bump is sensitive to the H envelope thickness of the white dwarfs, so using a data space that includes white dwarfs old enough that their envelopes have convectively coupled to their cores is important for constraining the typical white dwarf envelope thickness for this population. At such late cooling times, the white dwarfs in 47 Tuc have also begun to undergo core crystallisation. As shown by the earlier work of Obertas *et al.* [30], which used the same deep HST observations of 47 Tuc as were analysed in Chapter 6, the convective coupling of the envelope to the core for white dwarfs in 47 Tuc coincides with the onset of core crystallisation in these white dwarfs. Both of these phenomena produce a bump in the white dwarf luminosity function, and Obertas *et al.* [30] showed that these features overlap in the case of 47 Tuc. At these late cooling times, the white dwarf cores should thus no longer be approximated as an ideal gas from a theoretical standpoint,

though the usual implementation of element diffusion in stellar evolution code like MESA approximates ions as an ideal gas even for simulations of white dwarf cooling in order to make the diffusion equations tractable. In addition to constraining the values of parameters like the envelope thickness that are important for modelling the cooling of white dwarfs in 47 Tuc, the analysis performed in Chapter 6 thus also provides a good test of the parameter regime over which the standard MESA implementation of element diffusion can be used.

The results of Chapter 6 show that the standard MESA implementation of element diffusion produces white dwarf cooling models that well reproduce the empirical cumulative white dwarf luminosity function even at cooling times old enough that the core has begun to crystallise. The best-fitting model from the work of Chapter 6 corresponds to the case of standard MESA diffusion, and models with very thick envelopes ($-3.75 \leq \log_{10} q_H \leq -3.55$) and white dwarf masses in the limited range $M_{\text{WD}} = 0.5314 - 0.5338 M_{\odot}$ are strongly favoured (at the 3σ level in the posterior probability distribution) for that diffusion scenario. The best-fitting model was specifically given by the case of standard diffusion with $M_{\text{WD}} = 0.5314 M_{\odot}$ and $\log_{10} q_H = -3.55$. There is some degeneracy between which diffusion scenario is used and the most likely H envelope thickness (as well as the white dwarf mass to a lesser degree), with the cases of modified diffusion and no diffusion favouring slightly thinner (though still quite thick) envelopes. The modified treatment of diffusion, in particular, produces some cooling models that are nearly as likely (within 1σ) as the best-fitting model that used the standard treatment of diffusion.

Further distinguishing between the standard and modified treatments of diffusion and breaking the degeneracy between the diffusion scenarios and the envelope thickness at a statistically significant level could potentially be accomplished in the future if new data is obtained with higher completeness at the faint magnitude end of the white dwarf cooling sequence, but that is beyond the scope of what can be accomplished with the deep HST data analysed in Chapter 6. The careful cleaning and calibration procedure performed in Chapter 6 has already pushed the limit of what can be done with the deep

HST data considered in that chapter. The data is magnitude-limited and its completeness decreases significantly with increasing magnitude at the faint end of the 47 Tuc white dwarf cooling sequence where the fit is most sensitive to the envelope thickness and core crystallisation. As the analysis of Chapter 6 cannot distinguish between the modified and standard diffusion scenarios at a statistically significant level, there is no reason to perform a more detailed analysis in this work of precisely how diffusion should be modified to best account for non-ideal gas effects, though this would be an interesting topic to explore in future work if better data becomes available in the future, such as from future observations performed by a newer telescope like JWST. This potential future work is also not necessary for the use of the Chapter 6 results in the analysis of Chapter 7.

In Chapter 7, simulations were performed using MESA to create white dwarf cooling models analogous to those of Chapter 6 but that also accounted for the additional energy loss associated with the emission of axions produced in the interior of the white dwarf via axion bremsstrahlung from electrons. The results of Chapter 6 were used in Chapter 7 to choose which treatment of element diffusion to use when creating white dwarf cooling models and as prior information to constrain the white dwarf mass and envelope thickness values. Only the standard MESA treatment of diffusion was used, and the cooling models were restricted to the limited white dwarf mass range $M_{\text{WD}} = 0.5240 - 0.5388 M_{\odot}$ based on the 5σ level of the Chapter 6 results for the standard diffusion scenario. Cooling models were created for a grid of parameter values, varying the white dwarf mass (over this limited range), the H envelope thickness, and the axion-electron coupling constant.

These cooling models were compared to HST observations of young white dwarfs in 47 Tuc using an unbinned likelihood analysis procedure similar to that of Chapter 6 (and Goldsbury *et al.* [215]). The data analysed in Chapter 7 are separate sets of HST observations that are independent from the deep HST data analysed in Chapter 6. The data considered in Chapter 7 also consist of observations by two HST instruments, WFC3 and ACS, that both observed 47 Tuc as part of the same observation scheme, and both the WFC3 and ACS data sets were analysed simultaneously in a combined

analysis to produce the final results of Chapter 7. These are furthermore the same data that were analysed by Goldsbury *et al.* [215] to study neutrino cooling in 47 Tuc.

The results of Chapter 7 favour white dwarf cooling models with thick envelopes and negligible axion emission. The best-fitting model from the analysis in Chapter 7 was found to correspond to parameter values (on the discrete parameter grid) of $M_{\text{WD}} = 0.5388 M_{\odot}$, $\log_{10} q_H = -3.55$, and $g_{aee} = 0$, and it was found that the cumulative white dwarf luminosity functions predicted by this model for both the WFC3 and ACS data sets well-reproduce the corresponding empirical white dwarf luminosity functions. The analysis of Chapter 7 furthermore provides an improved bound on the axion-electron coupling of $g_{aee} \leq 0.81 \times 10^{-13}$ at 95% confidence. For a DFSZ model, this corresponds to a bound on the axion mass (and angular parameter) of $m_a \sin^2 \beta \leq 2.85 \text{ meV}$ at 95% confidence. This new constraint on g_{aee} (and m_a for the DFSZ model) is more stringent than the previous bound from the Galactic white dwarf luminosity function [48], as well as the stronger (and previously leading) bounds from the tip of the red giant branch of globular clusters [92, 93]. Contemporaneous with this dissertation, the bound on g_{aee} from the tip of the red giant branch (at 95% confidence) has been updated by Carenza *et al.* [305] from $g_{aee} \leq 1.5 \times 10^{-13}$ [92] to $g_{aee} \leq 0.96 \times 10^{-13}$ [305] using the updated distances provided by *Gaia* DR3 for 21 of the 22 globular clusters used by Straniero *et al.* [92] to derive the previous bound. However, the new bound of $g_{aee} \leq 0.81 \times 10^{-13}$ found in Chapter 7 is more stringent than even the updated bound from the tip of the red giant branch [305]. This improved constraint on g_{aee} furthermore excludes the range of values favoured by the hints of axions from Galactic disc and halo white dwarf luminosity functions [49].

This bound on g_{aee} could potentially be further improved in the future by using non-uniform priors for the white dwarf birthrates. The analysis of Goldsbury *et al.* [215] used Gaussian birthrate priors that were determined from stars on the red giant branch, similar to what was done in Chapter 6, and the birthrate priors from Goldsbury *et al.* [215] were considered for the analysis in Chapter 7. However, it was shown in Chapter 7 that these

priors appear to simultaneously underestimate the birthrate in the inner WFC3 field and overestimate the birthrate in the outer ACS field by a similar, appreciable amount. This suggests that white dwarfs may be leaving the inner WFC3 field and entering the outer ACS field at a measurable rate, which could be due to a phenomenon like cluster relaxation [261, 262]. This would be an important phenomenon to account for in modelling and determining a prior for the birthrate if a non-uniform prior is to be used, and it also presents interesting possibilities for future research.

Astrophysical axion searches using observations of white dwarfs typically probe the axion-electron coupling, as was done in Chapter 7. While other astrophysical objects have been recognised as targets for axion searches that probe the coupling of axions to nucleons, such as helioscope searches that look for axions produced by the ^{57}Fe transition in the Sun, white dwarfs have not previously been recognised as a target for searches that probe this coupling. The work presented in Chapter 8 showed, for the first time, the potential for observations of very hot, magnetic white dwarfs to probe the coupling of axions to nucleons. The proposed mechanism was the production of axions through the decay of the first excited state of ^{57}Fe in the core of the white dwarf, with the emitted axions subsequently converting to photons in the strong magnetic field surrounding the white dwarf. The signal associated with this process would be a narrow peak in the X-ray spectrum of the white dwarf at ~ 14.4 keV, the excitation energy of the first excited state of ^{57}Fe . The white dwarf ZTF J1901+1458 was identified as a particularly good target to search for this signal, and the projected sensitivity that could be achieved by observing ZTF J1901+1458 with the X-ray telescope *NuSTAR* was calculated.

In Chapter 8, it was shown that a 100 ks observation of ZTF J1901+1458 with *NuSTAR* would be sensitive to the product of the effective axion-nucleon and axion-photon couplings, $|g_{aNN}^{\text{eff}} g_{a\gamma\gamma}|$, at a level below that of both current and future helioscope searches that look for an analogous signal for axions produced in the Sun through the ^{57}Fe transition. Such an observation would also be sensitive to $|g_{aNN}^{\text{eff}} g_{a\gamma\gamma}|$ at a level below the strongest current bound, given by combining the saturated bound on g_{aNN}^{eff}

from SN 1987A and the most recent bound on $g_{a\gamma\gamma}$ from CAST searches for Primakoff axions produced in the Sun.

While the sensitivity curve calculated in Chapter 8 is currently just a projection of what could be achieved, this work demonstrates the potential for probing the coupling of axions to nucleons using observations of very hot, magnetic white dwarfs. Future work will be able to determine new results for $|g_{aNN}^{\text{eff}} g_{a\gamma\gamma}|$ by applying the procedure detailed in Chapter 8 to actual data from *NuSTAR*, and observation time has been awarded by *NuSTAR* for this purpose [5]. Observation time has also been awarded by HST for ultraviolet spectroscopy observations of ZTF J1901+1458 [304], which will allow the temperature of ZTF J1901+1458 to be determined more precisely. The *NuSTAR* observations have in fact already been performed, but an accurate determination of the temperature of ZTF J1901+1458 from HST observations is still needed before proceeding with the final analysis of this *NuSTAR* data. The emission rate of axions produced by the ^{57}Fe transition is highly sensitive to the core temperature of the white dwarf, so the precise implications for $|g_{aNN}^{\text{eff}} g_{a\gamma\gamma}|$ rely on accurately determining this temperature.

The results also depend on the magnetic field strength of the magnetosphere through the axion-photon conversion probability, though this dependence on the magnetic field strength is less significant than the dependence on the core temperature. Given the current very large uncertainty in the value of the temperature of ZTF J1901+1458, in addition to the stronger dependence of the results on the core temperature compared to the magnetic field strength, the uncertainty associated with modelling the magnetic field of ZTF J1901+1458 was only a sub-dominant effect that could be neglected in the calculations performed in Chapter 8. For future calculations where the uncertainty in the temperature has been significantly reduced, however, accurately modelling the magnetic field and accounting for the uncertainty in the field strength will become a more important consideration.

It should also be noted that the mechanism proposed in Chapter 8 of axions being produced in the core of a hot white dwarf through the ^{57}Fe transition and then converting to photons in the magnetosphere of a strongly magnetic white dwarf would occur in any sufficiently hot and magnetic white

dwarf with enough ^{57}Fe in its core. While ZTF J1901+1458 is a particularly good target to search for the signal associated with this process, other white dwarfs may also make good targets to search for this signal. Furthermore, while the calculations of Chapter 8 assumed a 100 ks observation, longer observations could also be performed to probe smaller values of $|g_{aNN}^{\text{eff}} g_{a\gamma\gamma}|$.

Bibliography

- [1] L. Fleury, *Simulations of axion string-wall networks*, Master's thesis, McGill University, Montreal, Canada (2015).
- [2] L. Fleury, I. Caiazzo, and J. Heyl, The cooling of massive white dwarfs from Gaia EDR3, [MNRAS 511, 5984 \(2022\)](#), arXiv:2110.00598.
- [3] L. Fleury, I. Caiazzo, and J. Heyl, The origin of ultramassive white dwarfs: hints from Gaia EDR3, [MNRAS 520, 364 \(2023\)](#), arXiv:2205.01015.
- [4] L. Fleury, I. Caiazzo, and J. Heyl, Constraining axions with ZTF J1901+1458, [Phys. Rev. D 107, L101303 \(2023\)](#), arXiv:2208.00405.
- [5] I. Caiazzo, L. Fleury, and J. Heyl, Constraining Axions with ZTF J1901+1458, NuSTAR Proposal. Cycle 8, ID. #8204 (2021).
- [6] D. E. Winget and S. O. Kepler, Pulsating white dwarf stars and precision asteroseismology, [ARA&A 46, 157 \(2008\)](#), arXiv:0806.2573.
- [7] A. H. C  rsico, L. G. Althaus, M. M. Miller Bertolami, and S. O. Kepler, Pulsating white dwarfs: new insights, [A&ARv 27, 7 \(2019\)](#), arXiv:1907.00115.
- [8] S. E. Woosley and A. Heger, The Remarkable Deaths of 9-11 Solar Mass Stars, [ApJ 810, 34 \(2015\)](#), arXiv:1505.06712.
- [9] C. J. Hansen, S. D. Kawaler, and V. Trimble, *Stellar Interiors: Physical Principles, Structure, and Evolution*, 2nd ed. (Springer, New York, USA, 2004).

Bibliography

- [10] S. O. Kepler, S. J. Kleinman, A. Nitta, D. Koester, B. G. Castanheira, O. Giovannini, A. F. M. Costa, and L. Althaus, White dwarf mass distribution in the SDSS, *MNRAS* **375**, 1315 (2007), [arXiv:astro-ph/0612277](https://arxiv.org/abs/astro-ph/0612277).
- [11] J. Liebert, P. Bergeron, and J. B. Holberg, The Formation Rate and Mass and Luminosity Functions of DA White Dwarfs from the Palomar Green Survey, *ApJS* **156**, 47 (2005), [arXiv:astro-ph/0406657](https://arxiv.org/abs/astro-ph/0406657).
- [12] M. Kilic, P. Bergeron, A. Kosakowski, W. R. Brown, M. A. Agüeros, and S. Blouin, The 100 pc White Dwarf Sample in the SDSS Footprint, *ApJ* **898**, 84 (2020), [arXiv:2006.00323](https://arxiv.org/abs/2006.00323).
- [13] F. M. Jiménez-Esteban, S. Torres, A. Rebassa-Mansergas, P. Cruz, R. Murillo-Ojeda, E. Solano, C. Rodrigo, and M. E. Camisassa, Spectral classification of the 100 pc white dwarf population from Gaia-DR3 and the virtual observatory, *MNRAS* **518**, 5106 (2023), [arXiv:2211.08852](https://arxiv.org/abs/2211.08852).
- [14] A. Caron, P. Bergeron, S. Blouin, and S. K. Leggett, A spectrophotometric analysis of cool white dwarfs in the Gaia and pan-STARRS footprint, *MNRAS* **519**, 4529 (2023), [arXiv:2212.08014](https://arxiv.org/abs/2212.08014).
- [15] S. J. Kleinman, S. O. Kepler, D. Koester, I. Pelisoli, V. Peçanha, A. Nitta, J. E. S. Costa, J. Krzesinski, P. Dufour, F. R. Lachapelle, P. Bergeron, C.-W. Yip, H. C. Harris, D. J. Eisenstein, L. Althaus, and A. Córscico, SDSS DR7 White Dwarf Catalog, *ApJS* **204**, 5 (2013), [arXiv:1212.1222](https://arxiv.org/abs/1212.1222).
- [16] S. E. Woosley, A. Heger, and T. A. Weaver, The evolution and explosion of massive stars, *Reviews of Modern Physics* **74**, 1015 (2002).
- [17] G. Wallerstein, J. Iben, Icko, P. Parker, A. M. Boesgaard, G. M. Hale, A. E. Champagne, C. A. Barnes, F. Käppeler, V. V. Smith, R. D. Hoffman, F. X. Timmes, C. Sneden, R. N. Boyd, B. S. Meyer, and D. L. Lambert, Synthesis of the elements in stars: forty years of progress, *Reviews of Modern Physics* **69**, 995 (1997).

Bibliography

- [18] C. E. Fields, R. Farmer, I. Petermann, C. Iliadis, and F. X. Timmes, Properties of Carbon-Oxygen White Dwarfs From Monte Carlo Stellar Models, [ApJ 823, 46 \(2016\)](#), [arXiv:1603.06666](#).
- [19] M. Salaris, S. Cassisi, A. Pietrinferni, P. M. Kowalski, and J. Isern, A Large Stellar Evolution Database for Population Synthesis Studies. VI. White Dwarf Cooling Sequences, [ApJ 716, 1241 \(2010\)](#), [arXiv:1005.1791](#).
- [20] F. C. De Gerónimo, L. G. Althaus, A. H. Córscico, A. D. Romero, and S. O. Kepler, Asteroseismology of ZZ Ceti stars with fully evolutionary white dwarf models. I. The impact of the uncertainties from prior evolution on the period spectrum, [A&A 599, A21 \(2017\)](#), [arXiv:1611.10298](#).
- [21] L. Siess, Evolution of massive AGB stars. I. Carbon burning phase, [A&A 448, 717 \(2006\)](#).
- [22] L. Siess, Evolution of massive AGB stars. II. model properties at non-solar metallicity and the fate of Super-AGB stars, [A&A 476, 893 \(2007\)](#).
- [23] C. L. Doherty, L. Siess, J. C. Lattanzio, and P. Gil-Pons, Super asymptotic giant branch stars. I - Evolution code comparison, [MNRAS 401, 1453 \(2010\)](#).
- [24] C. L. Doherty, P. Gil-Pons, L. Siess, J. C. Lattanzio, and H. H. B. Lau, Super- and massive AGB stars - IV. Final fates - initial-to-final mass relation, [MNRAS 446, 2599 \(2015\)](#), [arXiv:1410.5431](#).
- [25] H. M. van Horn, Crystallization of White Dwarfs, [ApJ 151, 227 \(1968\)](#).
- [26] E. Garcia-Berro, M. Hernanz, R. Mochkovitch, and J. Isern, Theoretical white-dwarf luminosity functions for two phase diagrams of the carbon-oxygen dense plasma., [A&A 193, 141 \(1988\)](#).

Bibliography

- [27] L. Segretain, G. Chabrier, M. Hernanz, E. García-Berro, J. Isern, and R. Mochkovitch, Cooling Theory of Crystallized White Dwarfs, *ApJ* **434**, 641 (1994).
- [28] J. Isern, R. Mochkovitch, E. García-Berro, and M. Hernanz, The Physics of Crystallizing White Dwarfs, *ApJ* **485**, 308 (1997), [arXiv:astro-ph/9703028](https://arxiv.org/abs/astro-ph/9703028).
- [29] L. G. Althaus, E. García-Berro, J. Isern, A. H. Córscico, and M. M. Miller Bertolami, New phase diagrams for dense carbon-oxygen mixtures and white dwarf evolution, *A&A* **537**, A33 (2012), [arXiv:1110.5665](https://arxiv.org/abs/1110.5665).
- [30] A. Obertas, I. Caiazzo, J. Heyl, H. Richer, J. Kalirai, and P.-E. Tremblay, The onset of convective coupling and freezing in the white dwarfs of 47 Tucanae, *MNRAS* **474**, 677 (2018), [arXiv:1709.08097](https://arxiv.org/abs/1709.08097).
- [31] P.-E. Tremblay, G. Fontaine, N. P. G. Fusillo, B. H. Dunlap, B. T. Gänsicke, M. A. Hollands, J. J. Hermes, T. R. Marsh, E. Cukanovaite, and T. Cunningham, Core crystallization and pile-up in the cooling sequence of evolving white dwarfs, *Nature* **565**, 202 (2019), [arXiv:1908.00370](https://arxiv.org/abs/1908.00370).
- [32] M. S. Bessell, UBVRI passbands., *PASP* **102**, 1181 (1990).
- [33] M. S. Bessell, Standard Photometric Systems, *ARA&A* **43**, 293 (2005).
- [34] A. G. A. Brown et al. (Gaia Collaboration), Gaia Early Data Release 3. Summary of the contents and survey properties, *A&A* **649**, A1 (2021), [arXiv:2012.01533](https://arxiv.org/abs/2012.01533).
- [35] C. Rodrigo, E. Solano, and A. Bayo, *SVO Filter Profile Service Version 1.0*, IVOA Working Draft 15 October 2012 (2012).
- [36] C. Rodrigo and E. Solano, The SVO Filter Profile Service, in *XIV.0 Scientific Meeting (virtual) of the Spanish Astronomical Society*, SEA Scientific Meeting Series (2020) held

Bibliography

13 – 15 July 2020, online at <https://www.sea-astronomia.es/reunion-cientifica-2020>, id.182.

- [37] E. L. Fitzpatrick, Correcting for the Effects of Interstellar Extinction, *PASP* **111**, 63 (1999), arXiv:astro-ph/9809387.
- [38] J. A. Cardelli, G. C. Clayton, and J. S. Mathis, The Relationship between Infrared, Optical, and Ultraviolet Extinction, *ApJ* **345**, 245 (1989).
- [39] J. E. O'Donnell, R_ν -dependent Optical and Near-Ultraviolet Extinction, *ApJ* **422**, 158 (1994).
- [40] L. A. Valencic, G. C. Clayton, and K. D. Gordon, Ultraviolet Extinction Properties in the Milky Way, *ApJ* **616**, 912 (2004), arXiv:astro-ph/0408409.
- [41] G. V. Schultz and W. Wiemer, Interstellar reddening and IR-excesses of O and B stars., *A&A* **43**, 133 (1975).
- [42] T. Prusti et al. (Gaia Collaboration), The Gaia mission, *A&A* **595**, A1 (2016), arXiv:1609.04153.
- [43] S. Cheng, J. D. Cummings, and B. Ménard, A Cooling Anomaly of High-mass White Dwarfs, *ApJ* **886**, 100 (2019), arXiv:1905.12710.
- [44] J. A. Munn, H. C. Harris, T. von Hippel, M. Kilic, J. W. Liebert, K. A. Williams, S. DeGennaro, E. Jeffery, K. Dame, A. Gianninas, and W. R. Brown, A Deep Proper Motion Catalog Within the Sloan Digital Sky Survey Footprint. II. The White Dwarf Luminosity Function, *AJ* **153**, 10 (2017), arXiv:1611.06275.
- [45] M. Kilic, J. A. Munn, H. C. Harris, T. von Hippel, J. W. Liebert, K. A. Williams, E. Jeffery, and S. DeGennaro, The Ages of the Thin Disk, Thick Disk, and the Halo from Nearby White Dwarfs, *ApJ* **837**, 162 (2017), arXiv:1702.06984.

Bibliography

- [46] N. Rowell and N. C. Hambly, White dwarfs in the SuperCOSMOS Sky Survey: the thin disc, thick disc and spheroid luminosity functions, *MNRAS* **417**, 93 (2011), arXiv:1102.3193.
- [47] J. Isern, E. García-Berro, S. Torres, and S. Catalán, Axions and the Cooling of White Dwarf Stars, *ApJ* **682**, L109 (2008), arXiv:0806.2807.
- [48] M. M. Miller Bertolami, B. E. Melendez, L. G. Althaus, and J. Isern, Revisiting the axion bounds from the Galactic white dwarf luminosity function, *J. Cosmol. Astropart. Phys.* **2014** (10), 069, arXiv:1406.7712.
- [49] J. Isern, E. García-Berro, S. Torres, R. Cojocaru, and S. Catalán, Axions and the luminosity function of white dwarfs: the thin and thick discs, and the halo, *MNRAS* **478**, 2569 (2018), arXiv:1805.00135.
- [50] J. Binney and S. Tremaine, *Galactic Dynamics: Second Edition* (Princeton University Press, Princeton, NJ USA, 2008).
- [51] G. G. Raffelt, Astrophysical methods to constrain axions and other novel particle phenomena, *Phys. Rep.* **198**, 1 (1990).
- [52] G. G. Raffelt, *Stars as laboratories for fundamental physics: The astrophysics of neutrinos, axions, and other weakly interacting particles* (University of Chicago Press, Chicago, USA, 1996).
- [53] C. Dessert, A. J. Long, and B. R. Safdi, X-Ray Signatures of Axion Conversion in Magnetic White Dwarf Stars, *Phys. Rev. Lett.* **123**, 061104 (2019), arXiv:1903.05088.
- [54] C. Dessert, A. J. Long, and B. R. Safdi, No Evidence for Axions from Chandra Observation of the Magnetic White Dwarf RE J0317-853, *Phys. Rev. Lett.* **128**, 071102 (2022), arXiv:2104.12772.
- [55] S. Andriamonje et al. (CAST Collaboration), Search for 14.4 keV solar axions emitted in the M1-transition of ^{57}Fe nuclei with CAST, *J. Cosmol. Astropart. Phys.* **2009** (12), 002, arXiv:0906.4488.

Bibliography

- [56] L. Di Luzio *et al.*, Probing the axion-nucleon coupling with the next generation of axion helioscopes, *European Physical Journal C* **82**, 120 (2022), arXiv:2111.06407.
- [57] I. Caiazzo *et al.*, A highly magnetized and rapidly rotating white dwarf as small as the Moon, *Nature* **595**, 39 (2021), arXiv:2107.08458.
- [58] L. Mestel, On the theory of white dwarf stars. I. The energy sources of white dwarfs, *MNRAS* **112**, 583 (1952).
- [59] S. L. Shapiro and S. A. Teukolsky, *Black holes, white dwarfs and neutron stars. The physics of compact objects* (Wiley, 1983).
- [60] H. M. van Horn, Cooling of White Dwarfs, in *White Dwarfs*, IAU Symposium, Vol. 42, edited by W. J. Luyten (1971) p. 97.
- [61] M. Schwarzschild, *Structure and evolution of the stars*. (Princeton University Press, 1958).
- [62] N. Itoh, T. Adachi, M. Nakagawa, Y. Kohyama, and H. Munakata, Neutrino Energy Loss in Stellar Interiors. III. Pair, Photo-, Plasma, and Bremsstrahlung Processes, *ApJ* **339**, 354 (1989).
- [63] N. Itoh, H. Mutoh, A. Hikita, and Y. Kohyama, Neutrino Energy Loss in Stellar Interiors. IV. Plasma Neutrino Process for Strongly Degenerate Electrons, *ApJ* **395**, 622 (1992).
- [64] N. Itoh, H. Hayashi, A. Nishikawa, and Y. Kohyama, Neutrino Energy Loss in Stellar Interiors. VII. Pair, Photo-, Plasma, Bremsstrahlung, and Recombination Neutrino Processes, *ApJS* **102**, 411 (1996).
- [65] M. Nakagawa, Y. Kohyama, and N. Itoh, Axion Bremsstrahlung in Dense Stars, *ApJ* **322**, 291 (1987).
- [66] M. Nakagawa, T. Adachi, Y. Kohyama, and N. Itoh, Axion Bremsstrahlung in Dense Stars. II. Phonon Contributions, *ApJ* **326**, 241 (1988).

Bibliography

- [67] K.-H. Böhm, Outer Convection Zones and Internal Temperatures of Cool White Dwarfs, *Ap&SS* **2**, 375 (1968).
- [68] G. Fontaine, P. Brassard, and P. Bergeron, The Potential of White Dwarf Cosmochronology, *PASP* **113**, 409 (2001).
- [69] E. E. Mamajek, G. Torres, A. Prsa, P. Harmanec, M. Asplund, P. D. Bennett, N. Capitaine, J. Christensen-Dalsgaard, E. Depagne, W. M. Folkner, M. Haberreiter, S. Hekker, J. L. Hilton, V. Kostov, D. W. Kurtz, J. Laskar, B. D. Mason, E. F. Milone, M. M. Montgomery, M. T. Richards, J. Schou, and S. G. Stewart, IAU 2015 Resolution B2 on Recommended Zero Points for the Absolute and Apparent Bolometric Magnitude Scales, arXiv e-prints (2015), [arXiv:1510.06262](https://arxiv.org/abs/1510.06262).
- [70] F. K. Guo, R. Horsley, U. G. Meißner, Y. Nakamura, H. Perlitz, P. E. L. Rakow, G. Schierholz, A. Schiller, and J. M. Zanotti, Electric Dipole Moment of the Neutron from 2 +1 Flavor Lattice QCD, *Phys. Rev. Lett.* **115**, 062001 (2015), [arXiv:1502.02295](https://arxiv.org/abs/1502.02295).
- [71] C. A. Baker, D. D. Doyle, P. Geltenbort, K. Green, M. G. D. van der Grinten, P. G. Harris, P. Iaydjiev, S. N. Ivanov, D. J. R. May, J. M. Pendlebury, J. D. Richardson, D. Shiers, and K. F. Smith, Improved Experimental Limit on the Electric Dipole Moment of the Neutron, *Phys. Rev. Lett.* **97**, 131801 (2006), [arXiv:hep-ex/0602020](https://arxiv.org/abs/hep-ex/0602020).
- [72] C. A. Baker, D. D. Doyle, P. Geltenbort, K. Green, M. G. D. van der Grinten, P. G. Harris, P. Iaydjiev, S. N. Ivanov, D. J. R. May, J. M. Pendlebury, J. D. Richardson, D. Shiers, and K. F. Smith, Reply to Comment on ‘Improved Experimental Limit on the Electric Dipole Moment of the Neutron’, *Phys. Rev. Lett.* **98**, 149102 (2007), [arXiv:0704.1354](https://arxiv.org/abs/0704.1354).
- [73] R. D. Peccei and H. R. Quinn, CP conservation in the presence of pseudoparticles, *Phys. Rev. Lett.* **38**, 1440 (1977).
- [74] R. D. Peccei and H. R. Quinn, Constraints imposed by CP conservation in the presence of pseudoparticles, *Phys. Rev. D* **16**, 1791 (1977).

Bibliography

- [75] R. D. Peccei, QCD, strong CP and axions., *Journal of Korean Physical Society* **29**, S199 (1996), [arXiv:hep-ph/9606475](https://arxiv.org/abs/hep-ph/9606475).
- [76] G. 't Hooft, Computation of the quantum effects due to a four-dimensional pseudoparticle, *Phys. Rev. D* **14**, 3432 (1976).
- [77] S. Weinberg, A new light boson?, *Phys. Rev. Lett.* **40**, 223 (1978).
- [78] F. Wilczek, Problem of strong P and T invariance in the presence of instantons, *Phys. Rev. Lett.* **40**, 279 (1978).
- [79] J. E. Kim, Weak-interaction singlet and strong CP invariance, *Phys. Rev. Lett.* **43**, 103 (1979).
- [80] M. A. Shifman, A. I. Vainshtein, and V. I. Zakharov, Can confinement ensure natural CP invariance of strong interactions?, *Nucl. Phys. B* **166**, 493 (1980).
- [81] M. Dine, W. Fischler, and M. Srednicki, A simple solution to the strong CP problem with a harmless axion, *Phys. Lett. B* **104**, 199 (1981).
- [82] A. R. Zhitnitsky, On Possible Suppression of the Axion Hadron Interactions. (In Russian), *Sov. J. Nucl. Phys.* **31**, 260 (1980), [*Yad. Fiz.* 31, 497 (1980)].
- [83] C. Burgess and G. Moore, *The Standard Model: A Primer* (Cambridge University Press, UK, 2007).
- [84] L. Di Luzio, M. Fedele, M. Giannotti, F. Mescia, and E. Nardi, Stellar evolution confronts axion models, *J. Cosmol. Astropart. Phys.* **2022** (02), 035, [arXiv:2109.10368](https://arxiv.org/abs/2109.10368).
- [85] L. Di Luzio, M. Giannotti, E. Nardi, and L. Visinelli, The landscape of QCD axion models, *Phys. Rep.* **870**, 1 (2020), [arXiv:2003.01100](https://arxiv.org/abs/2003.01100).
- [86] H. Primakoff, Photo-Production of Neutral Mesons in Nuclear Electric Fields and the Mean Life of the Neutral Meson, *Physical Review* **81**, 899 (1951).

Bibliography

- [87] A. Friedland, M. Giannotti, and M. Wise, Constraining the Axion-Photon Coupling with Massive Stars, *Phys. Rev. Lett.* **110**, 061101 (2013), [arXiv:1210.1271](https://arxiv.org/abs/1210.1271).
- [88] G. G. Raffelt and D. S. P. Dearborn, Bounds on hadronic axions from stellar evolution, *Phys. Rev. D* **36**, 2211 (1987).
- [89] A. Ayala, I. Domínguez, M. Giannotti, A. Mirizzi, and O. Straniero, Revisiting the Bound on Axion-Photon Coupling from Globular Clusters, *Phys. Rev. Lett.* **113**, 191302 (2014), [arXiv:1406.6053](https://arxiv.org/abs/1406.6053).
- [90] O. Straniero, A. Ayala, M. Giannotti, A. Mirizzi, and I. Dominguez, Axion-Photon Coupling: Astrophysical Constraints, in *Proceedings, 11th Patras Workshop on Axions, WIMPs and WISPs (Axion-WIMP 2015)*, DESY-PROC, edited by I. G. Irastorza, J. Redondo, J. M. Carmona, S. Cebrian, T. Dafni, F. J. Iguaz, and G. Luzon (Verlag Deutsches Elektronen-Synchrotron, Hamburg, Germany, 2015) pp. 77–81.
- [91] M. J. Dolan, F. J. Hiskens, and R. R. Volkas, Advancing globular cluster constraints on the axion-photon coupling, *J. Cosmol. Astropart. Phys.* **2022** (10), 096, [arXiv:2207.03102](https://arxiv.org/abs/2207.03102).
- [92] O. Straniero, C. Pallanca, E. Dalessandro, I. Domínguez, F. R. Ferraro, M. Giannotti, A. Mirizzi, and L. Piersanti, The RGB tip of galactic globular clusters and the revision of the axion-electron coupling bound, *A&A* **644**, A166 (2020), [arXiv:2010.03833](https://arxiv.org/abs/2010.03833).
- [93] F. Capozzi and G. Raffelt, Axion and neutrino bounds improved with new calibrations of the tip of the red-giant branch using geometric distance determinations, *Phys. Rev. D* **102**, 083007 (2020), [arXiv:2007.03694](https://arxiv.org/abs/2007.03694).
- [94] D. E. Winget, C. J. Hansen, and H. M. van Horn, Do pulsating PG1159-035 stars put constraints on stellar evolution?, *Nature* **303**, 781 (1983).

Bibliography

- [95] M. Giannotti, I. Irastorza, J. Redondo, and A. Ringwald, Cool WISPs for stellar cooling excesses, *J. Cosmol. Astropart. Phys.* **2016** (5), 057, [arXiv:1512.08108](https://arxiv.org/abs/1512.08108).
- [96] D. G. York et al. (SDSS Collaboration), The Sloan Digital Sky Survey: Technical Summary, *AJ* **120**, 1579 (2000), [arXiv:astro-ph/0006396](https://arxiv.org/abs/astro-ph/0006396).
- [97] D. J. Eisenstein, J. Liebert, H. C. Harris, S. J. Kleinman, A. Nitta, N. Silvestri, S. A. Anderson, J. C. Barentine, H. J. Brewington, J. Brinkmann, M. Harvanek, J. Krzesiński, J. Nielsen, Eric H., D. Long, D. P. Schneider, and S. A. Snedden, A Catalog of Spectroscopically Confirmed White Dwarfs from the Sloan Digital Sky Survey Data Release 4, *ApJS* **167**, 40 (2006), [arXiv:astro-ph/0606700](https://arxiv.org/abs/astro-ph/0606700).
- [98] N. P. Gentile Fusillo, P.-E. Tremblay, B. T. Gänsicke, C. J. Manser, T. Cunningham, E. Cukanovaite, M. Hollands, T. Marsh, R. Raddi, S. Jordan, S. Toonen, S. Geier, M. Barstow, and J. D. Cummings, A Gaia Data Release 2 catalogue of white dwarfs and a comparison with SDSS, *MNRAS* **482**, 4570 (2019), [arXiv:1807.03315](https://arxiv.org/abs/1807.03315).
- [99] N. P. Gentile Fusillo, P. E. Tremblay, E. Cukanovaite, A. Vorontseva, R. Lallement, M. Hollands, B. T. Gänsicke, K. B. Burdge, J. McCleery, and S. Jordan, A catalogue of white dwarfs in Gaia EDR3, *MNRAS* **508**, 3877 (2021), [arXiv:2106.07669](https://arxiv.org/abs/2106.07669).
- [100] C. Babusiaux et al. (Gaia Collaboration), Gaia Data Release 2. Observational Hertzsprung-Russell diagrams, *A&A* **616**, A10 (2018), [arXiv:1804.09378](https://arxiv.org/abs/1804.09378).
- [101] E. B. Bauer, J. Schwab, L. Bildsten, and S. Cheng, Multi-gigayear White Dwarf Cooling Delays from Clustering-enhanced Gravitational Sedimentation, *ApJ* **902**, 93 (2020), [arXiv:2009.04025](https://arxiv.org/abs/2009.04025).
- [102] M. E. Caplan, C. J. Horowitz, and A. Cumming, Neon Cluster Formation and Phase Separation during White Dwarf Cooling, *ApJ* **902**, L44 (2020), [arXiv:2010.00036](https://arxiv.org/abs/2010.00036).

Bibliography

- [103] S. Blouin, J. Daligault, D. Saumon, A. Bédard, and P. Brassard, Toward precision cosmochemistry. A new C/O phase diagram for white dwarfs, *A&A* **640**, L11 (2020), [arXiv:2007.13669](https://arxiv.org/abs/2007.13669).
- [104] S. Blouin, J. Daligault, and D. Saumon, ^{22}Ne Phase Separation as a Solution to the Ultramassive White Dwarf Cooling Anomaly, *ApJ* **911**, L5 (2021), [arXiv:2103.12892](https://arxiv.org/abs/2103.12892).
- [105] M. E. Camisassa, L. G. Althaus, S. Torres, A. H. Córscico, A. Rebassa-Mansergas, P.-E. Tremblay, S. Cheng, and R. Raddi, Forever young white dwarfs: When stellar ageing stops, *A&A* **649**, L7 (2021), [arXiv:2008.03028](https://arxiv.org/abs/2008.03028).
- [106] R. Mor, A. C. Robin, F. Figueras, S. Roca-Fàbregas, and X. Luri, Gaia DR2 reveals a star formation burst in the disc 2-3 Gyr ago, *A&A* **624**, L1 (2019), [arXiv:1901.07564](https://arxiv.org/abs/1901.07564).
- [107] J. Isern, The Star Formation History in the Solar Neighborhood as Told by Massive White Dwarfs, *ApJ* **878**, L11 (2019), [arXiv:1905.10779](https://arxiv.org/abs/1905.10779).
- [108] S. Cheng, [Wd models](#) (2020).
- [109] A. Bédard, P. Bergeron, P. Brassard, and G. Fontaine, On the Spectral Evolution of Hot White Dwarf Stars. I. A Detailed Model Atmosphere Analysis of Hot White Dwarfs from SDSS DR12, *ApJ* **901**, 93 (2020), [arXiv:2008.07469](https://arxiv.org/abs/2008.07469).
- [110] M. E. Camisassa, L. G. Althaus, A. H. Córscico, F. C. De Gerónimo, M. M. Miller Bertolami, M. L. Novarino, R. D. Rohrmann, F. C. Wachlin, and E. García-Berro, The evolution of ultra-massive white dwarfs, *A&A* **625**, A87 (2019), [arXiv:1807.03894](https://arxiv.org/abs/1807.03894).
- [111] M. E. Camisassa, L. G. Althaus, R. D. Rohrmann, E. García-Berro, S. Torres, A. H. Córscico, and F. C. Wachlin, Updated Evolutionary Sequences for Hydrogen-deficient White Dwarfs, *ApJ* **839**, 11 (2017), [arXiv:1703.05340](https://arxiv.org/abs/1703.05340).

Bibliography

- [112] I. Renedo, L. G. Althaus, M. M. Miller Bertolami, A. D. Romero, A. H. Córscico, R. D. Rohrmann, and E. García-Berro, New Cooling Sequences for Old White Dwarfs, [ApJ 717, 183 \(2010\)](#), [arXiv:1005.2170](#).
- [113] G. R. Lauffer, A. D. Romero, and S. O. Kepler, New full evolutionary sequences of H- and He-atmosphere massive white dwarf stars using MESA, [MNRAS 480, 1547 \(2018\)](#), [arXiv:1807.04774](#).
- [114] L. G. Althaus, P. Gil-Pons, A. H. Córscico, M. Miller Bertolami, F. De Gerónimo, M. E. Camisassa, S. Torres, J. Gutierrez, and A. Rebassa-Mansergas, The formation of ultra-massive carbon-oxygen core white dwarfs and their evolutionary and pulsational properties, [A&A 646, A30 \(2021\)](#), [arXiv:2011.10439](#).
- [115] J. Schwab, Evolutionary Models for the Remnant of the Merger of Two Carbon-Oxygen Core White Dwarfs, [ApJ 906, 53 \(2021\)](#), [arXiv:2011.03546](#).
- [116] S. C. Yoon, P. Podsiadlowski, and S. Rosswog, Remnant evolution after a carbon-oxygen white dwarf merger, [MNRAS 380, 933 \(2007\)](#), [arXiv:0704.0297](#).
- [117] P. Lorén-Aguilar, J. Isern, and E. García-Berro, High-resolution smoothed particle hydrodynamics simulations of the merger of binary white dwarfs, [A&A 500, 1193 \(2009\)](#).
- [118] G. Wagstaff, M. M. Miller Bertolami, and A. Weiss, Impact of convective boundary mixing on the TP-AGB, [MNRAS 493, 4748 \(2020\)](#), [arXiv:2002.01860](#).
- [119] H.-W. Rix, D. W. Hogg, D. Boubert, A. G. A. Brown, A. Casey, R. Drimmel, A. Everall, M. Fouesneau, and A. M. Price-Whelan, Selection Functions in Astronomical Data Modeling, with the Space Density of White Dwarfs as a Worked Example, [AJ 162, 142 \(2021\)](#), [arXiv:2106.07653](#).

Bibliography

- [120] M. Schmidt, Space Distribution and Luminosity Functions of Quasi-Stellar Radio Sources, *ApJ* **151**, 393 (1968).
- [121] J. Choi, A. Dotter, C. Conroy, M. Cantiello, B. Paxton, and B. D. Johnson, Mesa Isochrones and Stellar Tracks (MIST). I. Solar-scaled Models, *ApJ* **823**, 102 (2016), arXiv:1604.08592.
- [122] J. Heyl, I. Caiazzo, and H. B. Richer, Reconstructing the Pleiades with Gaia EDR3, *ApJ* **926**, 132 (2022), arXiv:2110.03837.
- [123] K. D. Temmink, S. Toonen, E. Zapartas, S. Justham, and B. T. Gänsicke, Looks can be deceiving. Underestimating the age of single white dwarfs due to binary mergers, *A&A* **636**, A31 (2020), arXiv:1910.05335.
- [124] A. I. Bogomazov and A. V. Tutukov, Merging of components in close binaries: Type Ia supernovae, massive white dwarfs, and Ap stars, *Astronomy Reports* **53**, 214 (2009), arXiv:0901.4899.
- [125] M. Kilic, P. Bergeron, S. Blouin, and A. Bédard, The most massive white dwarfs in the solar neighbourhood, *MNRAS* **503**, 5397 (2021), arXiv:2103.06906.
- [126] S. Cheng, J. D. Cummings, B. Ménard, and S. Toonen, Double White Dwarf Merger Products among High-mass White Dwarfs, *ApJ* **891**, 160 (2020), arXiv:1910.09558.
- [127] G. Strömberg, The Motions of the Stars Within 20 Parsecs of the Sun, *ApJ* **104**, 12 (1946).
- [128] N. G. Roman, Some characteristics of the spectra of F-, G-, and K-type stars, *AJ* **55**, 182 (1950).
- [129] N. G. Roman, A Correlation Between the Spectroscopic and Dynamical Characteristics of the Late F- and Early G- Type Stars, *ApJ* **112**, 554 (1950).

Bibliography

- [130] R. Wielen, The Diffusion of Stellar Orbits Derived from the Observed Age-Dependence of the Velocity Dispersion, [A&A 60, 263 \(1977\)](#).
- [131] B. Nordström, M. Mayor, J. Andersen, J. Holmberg, F. Pont, B. R. Jørgensen, E. H. Olsen, S. Udry, and N. Mowlavi, The Geneva-Copenhagen survey of the Solar neighbourhood. Ages, metallicities, and kinematic properties of $\sim 14\,000$ F and G dwarfs, [A&A 418, 989 \(2004\)](#), [arXiv:astro-ph/0405198](#).
- [132] J. Holmberg, B. Nordström, and J. Andersen, The Geneva-Copenhagen survey of the Solar neighbourhood II. New uvby calibrations and rediscussion of stellar ages, the G dwarf problem, age-metallicity diagram, and heating mechanisms of the disk, [A&A 475, 519 \(2007\)](#), [arXiv:0707.1891](#).
- [133] J. Holmberg, B. Nordström, and J. Andersen, The Geneva-Copenhagen survey of the solar neighbourhood. III. Improved distances, ages, and kinematics, [A&A 501, 941 \(2009\)](#), [arXiv:0811.3982](#).
- [134] G. M. Seabroke and G. Gilmore, Revisiting the relations: Galactic thin disc age-velocity dispersion relation, [MNRAS 380, 1348 \(2007\)](#), [arXiv:0707.1027](#).
- [135] C. Soubiran, O. Bienaymé, T. V. Mishenina, and V. V. Kovtyukh, Vertical distribution of Galactic disk stars. IV. AMR and AVR from clump giants, [A&A 480, 91 \(2008\)](#), [arXiv:0712.1370](#).
- [136] M. Aumer and J. J. Binney, Kinematics and history of the solar neighbourhood revisited, [MNRAS 397, 1286 \(2009\)](#), [arXiv:0905.2512](#).
- [137] L. Casagrande, R. Schönrich, M. Asplund, S. Cassisi, I. Ramírez, J. Meléndez, T. Bensby, and S. Feltzing, New constraints on the chemical evolution of the solar neighbourhood and Galactic disc(s). Improved astrophysical parameters for the Geneva-Copenhagen Survey, [A&A 530, A138 \(2011\)](#), [arXiv:1103.4651](#).

Bibliography

- [138] S. Sharma, J. Bland-Hawthorn, J. Binney, K. C. Freeman, M. Steinmetz, C. Boeche, O. Bienaymé, B. K. Gibson, G. F. Gilmore, E. K. Grebel, A. Helmi, G. Kordopatis, U. Munari, J. F. Navarro, Q. A. Parker, W. A. Reid, G. M. Seabroke, A. Siebert, F. Watson, M. E. K. Williams, R. F. G. Wyse, and T. Zwitter, Kinematic Modeling of the Milky Way Using the RAVE and GCS Stellar Surveys, *ApJ* **793**, 51 (2014), [arXiv:1405.7435](https://arxiv.org/abs/1405.7435).
- [139] J. Yu and C. Liu, The age-velocity dispersion relation of the Galactic discs from LAMOST-Gaia data, *MNRAS* **475**, 1093 (2018), [arXiv:1712.03965](https://arxiv.org/abs/1712.03965).
- [140] R. Raddi, S. Torres, A. Rebassa-Mansergas, J. Maldonado, M. E. Camisassa, D. Koester, N. P. Gentile Fusillo, P.-E. Tremblay, M. Dimpel, U. Heber, T. Cunningham, and J.-J. Ren, Kinematic properties of white dwarfs. Galactic orbital parameters and age-velocity dispersion relation, *A&A* **658**, A22 (2022), [arXiv:2111.01145](https://arxiv.org/abs/2111.01145).
- [141] M. A. C. Perryman, L. Lindegren, J. Kovalevsky, E. Hoeg, U. Bastian, P. L. Bernacca, M. Crézé, F. Donati, M. Grenon, M. Grewing, F. van Leeuwen, H. van der Marel, F. Mignard, C. A. Murray, R. S. Le Poole, H. Schrijver, C. Turon, F. Arenou, M. Froeschlé, and C. S. Petersen, The HIPPARCOS Catalogue, *A&A* **323**, L49 (1997).
- [142] W. Dehnen, The Distribution of Nearby Stars in Velocity Space Inferred from HIPPARCOS Data, *AJ* **115**, 2384 (1998), [arXiv:astro-ph/9803110](https://arxiv.org/abs/astro-ph/9803110).
- [143] W. Dehnen and J. J. Binney, Local stellar kinematics from HIPPARCOS data, *MNRAS* **298**, 387 (1998), [arXiv:astro-ph/9710077](https://arxiv.org/abs/astro-ph/9710077).
- [144] E. Chereul, M. Crézé, and O. Bienaymé, The distribution of nearby stars in phase space mapped by Hipparcos. Clustering and streaming among A-F type stars, *A&AS* **135**, 5 (1999), [arXiv:astro-ph/9809264](https://arxiv.org/abs/astro-ph/9809264).

Bibliography

- [145] J. Skuljan, J. B. Hearnshaw, and P. L. Cottrell, Velocity distribution of stars in the solar neighbourhood, *MNRAS* **308**, 731 (1999), [arXiv:astro-ph/9905002](https://arxiv.org/abs/astro-ph/9905002).
- [146] B. Famaey, A. Jorissen, X. Luri, M. Mayor, S. Udry, H. Dejonghe, and C. Turon, Local kinematics of K and M giants from CORAVEL/Hipparcos/Tycho-2 data. Revisiting the concept of super-clusters, *A&A* **430**, 165 (2005), [arXiv:astro-ph/0409579](https://arxiv.org/abs/astro-ph/0409579).
- [147] B. Famaey, A. Siebert, and A. Jorissen, On the age heterogeneity of the Pleiades, Hyades, and Sirius moving groups, *A&A* **483**, 453 (2008), [arXiv:0712.1470](https://arxiv.org/abs/0712.1470).
- [148] T. Antoja, F. Figueras, D. Fernández, and J. Torra, Origin and evolution of moving groups. I. Characterization in the observational kinematic-age-metallicity space, *A&A* **490**, 135 (2008), [arXiv:0809.0511](https://arxiv.org/abs/0809.0511).
- [149] J. Bovy, Galactic rotation in Gaia DR1, *MNRAS* **468**, L63 (2017), [arXiv:1610.07610](https://arxiv.org/abs/1610.07610).
- [150] I. Kushniruk, T. Schirmer, and T. Bensby, Kinematic structures of the solar neighbourhood revealed by Gaia DR1/TGAS and RAVE, *A&A* **608**, A73 (2017), [arXiv:1709.06905](https://arxiv.org/abs/1709.06905).
- [151] D. Katz et al. (Gaia Collaboration), Gaia Data Release 2. Mapping the Milky Way disc kinematics, *A&A* **616**, A11 (2018), [arXiv:1804.09380](https://arxiv.org/abs/1804.09380).
- [152] N. Rowell and M. Kilic, The kinematics of Galactic disc white dwarfs in Gaia DR2, *MNRAS* **484**, 3544 (2019), [arXiv:1901.04948](https://arxiv.org/abs/1901.04948).
- [153] D. Mikkola, P. J. McMillan, D. Hobbs, and J. Wimarsson, The velocity distribution of white dwarfs in Gaia EDR3, *MNRAS* **512**, 6201 (2022), [arXiv:2202.07672](https://arxiv.org/abs/2202.07672).
- [154] T. Antoja, A. Helmi, O. Bienayme, J. Bland-Hawthorn, B. Famaey, K. Freeman, B. K. Gibson, G. Gilmore, E. K. Grebel, I. Minchev,

Bibliography

U. Munari, J. Navarro, Q. Parker, W. Reid, G. M. Seabroke, A. Siebert, A. Siviero, M. Steinmetz, M. Williams, R. Wyse, and T. Zwitter, Kinematic groups beyond the solar neighbourhood with RAVE, [MNRAS 426, L1 \(2012\)](#), [arXiv:1205.0546](#).

[155] M. E. Caplan, I. F. Freeman, C. J. Horowitz, A. Cumming, and E. P. Bellinger, Cooling Delays from Iron Sedimentation and Iron Inner Cores in White Dwarfs, [ApJ 919, L12 \(2021\)](#), [arXiv:2108.11389](#).

[156] C. J. Horowitz, Nuclear and dark matter heating in massive white dwarf stars, [Phys. Rev. D 102, 083031 \(2020\)](#), [arXiv:2008.03291](#).

[157] S. Blouin and J. Daligault, Phase Separation in Ultramassive White Dwarfs, [ApJ 919, 87 \(2021\)](#), [arXiv:2107.07094](#).

[158] A. Vallenari et al. (Gaia Collaboration), Gaia Data Release 3. Summary of the content and survey properties, [A&A 674, A1 \(2023\)](#), [arXiv:2208.00211](#).

[159] M. E. Camisassa, L. G. Althaus, D. Koester, S. Torres, P. G. Pons, and A. H. Córscico, The evolution of ultra-massive carbon-oxygen white dwarfs, [MNRAS 511, 5198 \(2022\)](#), [arXiv:2202.03495](#).

[160] R. P. Olling and W. Dehnen, The Oort Constants Measured from Proper Motions, [ApJ 599, 275 \(2003\)](#), [arXiv:astro-ph/0301486](#).

[161] F. Wang, H. W. Zhang, Y. Huang, B. Q. Chen, H. F. Wang, and C. Wang, Local stellar kinematics and Oort constants from the LAMOST A-type stars, [MNRAS 504, 199 \(2021\)](#), [arXiv:2103.10232](#).

[162] T. P. Robitaille et al. (Astropy Collaboration), Astropy: A community Python package for astronomy, [A&A 558, A33 \(2013\)](#), [arXiv:1307.6212](#).

[163] A. M. Price-Whelan et al. (Astropy Collaboration), The Astropy Project: Building an Open-science Project and Status of the v2.0 Core Package, [AJ 156, 123 \(2018\)](#), [arXiv:1801.02634](#).

Bibliography

- [164] M. Kilic, A. G. Moss, A. Kosakowski, P. Bergeron, A. A. Conly, W. R. Brown, S. Toonen, K. A. Williams, and P. Dufour, The merger fraction of ultramassive white dwarfs, *MNRAS* **518**, 2341 (2023), [arXiv:2211.05938](https://arxiv.org/abs/2211.05938).
- [165] C. Gallart, E. J. Bernard, C. B. Brook, T. Ruiz-Lara, S. Cassisi, V. Hill, and M. Monelli, Uncovering the birth of the Milky Way through accurate stellar ages with Gaia, *Nature Astronomy* **3**, 932 (2019), [arXiv:1901.02900](https://arxiv.org/abs/1901.02900).
- [166] P. P. Eggleton and A. A. Tokovinin, A catalogue of multiplicity among bright stellar systems, *MNRAS* **389**, 869 (2008), [arXiv:0806.2878](https://arxiv.org/abs/0806.2878).
- [167] D. Raghavan, H. A. McAlister, T. J. Henry, D. W. Latham, G. W. Marcy, B. D. Mason, D. R. Gies, R. J. White, and T. A. ten Brummelaar, A Survey of Stellar Families: Multiplicity of Solar-type Stars, *ApJS* **190**, 1 (2010), [arXiv:1007.0414](https://arxiv.org/abs/1007.0414).
- [168] A. Tokovinin, From Binaries to Multiples. II. Hierarchical Multiplicity of F and G Dwarfs, *AJ* **147**, 87 (2014), [arXiv:1401.6827](https://arxiv.org/abs/1401.6827).
- [169] M. Moe and R. Di Stefano, Mind Your Ps and Qs: The Interrelation between Period (P) and Mass-ratio (Q) Distributions of Binary Stars, *ApJS* **230**, 15 (2017), [arXiv:1606.05347](https://arxiv.org/abs/1606.05347).
- [170] H. Sana, J. B. Le Bouquin, S. Lacour, J. P. Berger, G. Duvert, L. Gauchet, B. Norris, J. Olofsson, D. Pickel, G. Zins, O. Absil, A. de Koter, K. Kratter, O. Schnurr, and H. Zinnecker, Southern Massive Stars at High Angular Resolution: Observational Campaign and Companion Detection, *ApJS* **215**, 15 (2014), [arXiv:1409.6304](https://arxiv.org/abs/1409.6304).
- [171] S. Toonen, S. Portegies Zwart, A. S. Hamers, and D. Bandopadhyay, The evolution of stellar triples. The most common evolutionary pathways, *A&A* **640**, A16 (2020), [arXiv:2004.07848](https://arxiv.org/abs/2004.07848).
- [172] J. Stegmann, F. Antonini, and M. Moe, Evolution of massive stellar

Bibliography

triples and implications for compact object binary formation, *MNRAS* **516**, 1406 (2022), arXiv:2112.10786.

[173] S. Toonen, A. Hamers, and S. Portegies Zwart, The evolution of hierarchical triple star-systems, *Computational Astrophysics and Cosmology* **3**, 6 (2016), arXiv:1612.06172.

[174] P. Hut and J. N. Bahcall, Binary-single star scattering. I - Numerical experiments for equal masses, *ApJ* **268**, 319 (1983).

[175] S. Toonen, T. C. N. Boekholt, and S. Portegies Zwart, Stellar triples on the edge. Comprehensive overview of the evolution of destabilised triples leading to stellar and binary exotica, *A&A* **661**, A61 (2022), arXiv:2108.04272.

[176] R. A. Mardling and S. J. Aarseth, Tidal interactions in star cluster simulations, *MNRAS* **321**, 398 (2001).

[177] H. B. Perets and K. M. Kratter, The Triple Evolution Dynamical Instability: Stellar Collisions in the Field and the Formation of Exotic Binaries, *ApJ* **760**, 99 (2012), arXiv:1203.2914.

[178] H. Glanz and H. B. Perets, Simulations of common envelope evolution in triple systems: circumstellar case, *MNRAS* **500**, 1921 (2021), arXiv:2004.00020.

[179] E. Grishin and H. B. Perets, Chaotic dynamics of wide triples induced by galactic tides: a novel channel for producing compact binaries, mergers, and collisions, *MNRAS* **512**, 4993 (2022), arXiv:2112.11475.

[180] A. S. Hamers, H. Glanz, and P. Neunteufel, A Statistical View of the Stable and Unstable Roche Lobe Overflow of a Tertiary Star onto the Inner Binary in Triple Systems, *ApJS* **259**, 25 (2022), arXiv:2110.00024.

[181] A. S. Hamers, H. B. Perets, T. A. Thompson, and P. Neunteufel, Return of the TEDI: Revisiting the Triple Evolution Dynamical Instability Channel in Triple Stars, *ApJ* **925**, 178 (2022), arXiv:2107.13620.

Bibliography

- [182] H. von Zeipel, Sur l'application des séries de M. Lindstedt à l'étude du mouvement des comètes périodiques, *Astronomische Nachrichten* **183**, 345 (1910).
- [183] Y. Kozai, Secular perturbations of asteroids with high inclination and eccentricity, *AJ* **67**, 591 (1962).
- [184] M. L. Lidov, The evolution of orbits of artificial satellites of planets under the action of gravitational perturbations of external bodies, *Planet. Space Sci.* **9**, 719 (1962).
- [185] S. Naoz, The Eccentric Kozai-Lidov Effect and Its Applications, *ARA&A* **54**, 441 (2016), arXiv:1601.07175.
- [186] T. Ito and K. Ohtsuka, The Lidov-Kozai Oscillation and Hugo von Zeipel, *Monographs on Environment, Earth and Planets* **7**, 1 (2019), arXiv:1911.03984.
- [187] O. Blaes, M. H. Lee, and A. Socrates, The Kozai Mechanism and the Evolution of Binary Supermassive Black Holes, *ApJ* **578**, 775 (2002), arXiv:astro-ph/0203370.
- [188] T. A. Thompson, Accelerating Compact Object Mergers in Triple Systems with the Kozai Resonance: A Mechanism for “Prompt” Type Ia Supernovae, Gamma-Ray Bursts, and Other Exotica, *ApJ* **741**, 82 (2011), arXiv:1011.4322.
- [189] A. S. Hamers, O. R. Pols, J. S. W. Claeys, and G. Nelemans, Population synthesis of triple systems in the context of mergers of carbon-oxygen white dwarfs, *MNRAS* **430**, 2262 (2013), arXiv:1301.1469.
- [190] J. M. Antognini, B. J. Shappee, T. A. Thompson, and P. Amaro-Seoane, Rapid eccentricity oscillations and the mergers of compact objects in hierarchical triples, *MNRAS* **439**, 1079 (2014), arXiv:1308.5682.

Bibliography

- [191] F. Antonini, S. Toonen, and A. S. Hamers, Binary Black Hole Mergers from Field Triples: Properties, Rates, and the Impact of Stellar Evolution, *ApJ* **841**, 77 (2017), [arXiv:1703.06614](#).
- [192] B. Liu and D. Lai, Spin-Orbit Misalignment of Merging Black Hole Binaries with Tertiary Companions, *ApJ* **846**, L11 (2017), [arXiv:1706.02309](#).
- [193] B. Liu and D. Lai, Black Hole and Neutron Star Binary Mergers in Triple Systems: Merger Fraction and Spin-Orbit Misalignment, *ApJ* **863**, 68 (2018), [arXiv:1805.03202](#).
- [194] K. Silsbee and S. Tremaine, Lidov-Kozai Cycles with Gravitational Radiation: Merging Black Holes in Isolated Triple Systems, *ApJ* **836**, 39 (2017), [arXiv:1608.07642](#).
- [195] A. S. Hamers, B. Bar-Or, C. Petrovich, and F. Antonini, The Impact of Vector Resonant Relaxation on the Evolution of Binaries near a Massive Black Hole: Implications for Gravitational-wave Sources, *ApJ* **865**, 2 (2018), [arXiv:1805.10313](#).
- [196] B.-M. Hoang, S. Naoz, B. Kocsis, F. A. Rasio, and F. Dosopoulou, Black Hole Mergers in Galactic Nuclei Induced by the Eccentric Kozai-Lidov Effect, *ApJ* **856**, 140 (2018), [arXiv:1706.09896](#).
- [197] L. Randall and Z.-Z. Xianyu, Induced Ellipticity for Inspiring Binary Systems, *ApJ* **853**, 93 (2018), [arXiv:1708.08569](#).
- [198] L. Randall and Z.-Z. Xianyu, An Analytical Portrait of Binary Mergers in Hierarchical Triple Systems, *ApJ* **864**, 134 (2018), [arXiv:1802.05718](#).
- [199] S. Toonen, H. B. Perets, and A. S. Hamers, Rate of WD-WD head-on collisions in isolated triples is too low to explain standard type Ia supernovae, *A&A* **610**, A22 (2018), [arXiv:1709.00422](#).
- [200] G. Fragione and A. Loeb, Black hole-neutron star mergers from triples, *MNRAS* **486**, 4443 (2019), [arXiv:1903.10511](#).

Bibliography

- [201] J. Iben, I. and A. V. Tutukov, On the production by triple stars of binary blue stragglers and cataclysmic variables, in *11th European Workshop on White Dwarfs*, Astronomical Society of the Pacific Conference Series, Vol. 169, edited by S. E. Solheim and E. G. Meistas (1999) p. 432.
- [202] J. A. Ahumada and E. Lapasset, New catalogue of blue stragglers in open clusters, *Boletin de la Asociacion Argentina de Astronomia La Plata Argentina* **48**, 177 (2005).
- [203] J. A. Ahumada and E. Lapasset, New catalogue of blue stragglers in open clusters, [A&A 463, 789 \(2007\)](#).
- [204] F. de Marchi, F. de Angeli, G. Piotto, G. Carraro, and M. B. Davies, Search and analysis of blue straggler stars in open clusters, [A&A 459, 489 \(2006\)](#), [arXiv:astro-ph/0608464](#).
- [205] F. R. Ferraro, B. Lanzoni, E. Dalessandro, G. Beccari, M. Pasquato, P. Miocchi, R. T. Rood, S. Sigurdsson, A. Sills, E. Vesperini, M. Mapelli, R. Contreras, N. Sanna, and A. Mucciarelli, Dynamical age differences among coeval star clusters as revealed by blue stragglers, [Nature 492, 393 \(2012\)](#), [arXiv:1212.5071](#).
- [206] J. Parada, H. Richer, J. Heyl, J. Kalirai, and R. Goldsbury, Formation and Evolution of Blue Stragglers in 47 Tucanae, [ApJ 830, 139 \(2016\)](#), [arXiv:1609.02115](#).
- [207] C. Soubiran, T. Cantat-Gaudin, M. Romero-Gómez, L. Casamiquela, C. Jordi, A. Vallenari, T. Antoja, L. Balaguer-Núñez, D. Bossini, A. Bragaglia, R. Carrera, A. Castro-Ginard, F. Figueras, U. Heiter, D. Katz, A. Krone-Martins, J. F. Le Campion, A. Moitinho, and R. Sordo, Open cluster kinematics with Gaia DR2, [A&A 619, A155 \(2018\)](#), [arXiv:1808.01613](#).
- [208] C. Soubiran, T. Cantat-Gaudin, M. Romero-Gómez, L. Casamiquela, C. Jordi, A. Vallenari, T. Antoja, L. Balaguer-Núñez, D. Bossini,

Bibliography

A. Bragaglia, R. Carrera, A. Castro-Ginard, F. Figueras, U. Heiter, D. Katz, A. Krone-Martins, J. F. Le Campion, A. Moitinho, and R. Sordo, Open cluster kinematics with Gaia DR2 (Corrigendum), [A&A 623, C2 \(2019\)](#).

[209] N. Lodieu, A. Pérez-Garrido, R. L. Smart, and R. Silvotti, A 5D view of the α Per, Pleiades, and Praesepe clusters, [A&A 628, A66 \(2019\)](#), [arXiv:1906.03924](#).

[210] W. E. Harris, A Catalog of Parameters for Globular Clusters in the Milky Way, [AJ 112, 1487 \(1996\)](#).

[211] D. A. Vandenberg and R. A. Bell, Theoretical isochrones for globular clusters with predicted BVRI and Stromgren photometry., [ApJS 58, 561 \(1985\)](#).

[212] J. S. Kalirai and H. B. Richer, Star clusters as laboratories for stellar and dynamical evolution, [Philosophical Transactions of the Royal Society of London Series A 368, 755 \(2010\)](#), [arXiv:0911.0789](#).

[213] B. M. S. Hansen, J. S. Kalirai, J. Anderson, A. Dotter, H. B. Richer, R. M. Rich, M. M. Shara, G. G. Fahrlman, J. R. Hurley, I. R. King, D. Reitzel, and P. B. Stetson, An age difference of two billion years between a metal-rich and a metal-poor globular cluster, [Nature 500, 51 \(2013\)](#), [arXiv:1308.0032](#).

[214] S. Chen, H. Richer, I. Caiazzo, and J. Heyl, Distances to the Globular Clusters 47 Tucanae and NGC 362 Using Gaia DR2 Parallaxes, [ApJ 867, 132 \(2018\)](#), [arXiv:1807.07089](#).

[215] R. Goldsbury, J. Heyl, H. B. Richer, J. S. Kalirai, and P. E. Tremblay, Constraining White Dwarf Structure and Neutrino Physics in 47 Tucanae, [ApJ 821, 27 \(2016\)](#), [arXiv:1602.06286](#).

[216] M. Tassoul, G. Fontaine, and D. E. Winget, Evolutionary Models for Pulsation Studies of White Dwarfs, [ApJS 72, 335 \(1990\)](#).

Bibliography

- [217] B. Paxton, L. Bildsten, A. Dotter, F. Herwig, P. Lesaffre, and F. Timmes, Modules for Experiments in Stellar Astrophysics (MESA), *ApJS* **192**, 3 (2011), [arXiv:1009.1622](https://arxiv.org/abs/1009.1622).
- [218] B. Paxton, M. Cantiello, P. Arras, L. Bildsten, E. F. Brown, A. Dotter, C. Mankovich, M. H. Montgomery, D. Stello, F. X. Timmes, and R. Townsend, Modules for Experiments in Stellar Astrophysics (MESA): Planets, Oscillations, Rotation, and Massive Stars, *ApJS* **208**, 4 (2013), [arXiv:1301.0319](https://arxiv.org/abs/1301.0319).
- [219] B. Paxton, P. Marchant, J. Schwab, E. B. Bauer, L. Bildsten, M. Cantiello, L. Dessart, R. Farmer, H. Hu, N. Langer, R. H. D. Townsend, D. M. Townsley, and F. X. Timmes, Modules for Experiments in Stellar Astrophysics (MESA): Binaries, Pulsations, and Explosions, *ApJS* **220**, 15 (2015), [arXiv:1506.03146](https://arxiv.org/abs/1506.03146).
- [220] B. Paxton, J. Schwab, E. B. Bauer, L. Bildsten, S. Blinnikov, P. Duffell, R. Farmer, J. A. Goldberg, P. Marchant, E. Sorokina, A. Thoul, R. H. D. Townsend, and F. X. Timmes, Modules for Experiments in Stellar Astrophysics (MESA): Convective Boundaries, Element Diffusion, and Massive Star Explosions, *ApJS* **234**, 34 (2018), [arXiv:1710.08424](https://arxiv.org/abs/1710.08424).
- [221] B. Paxton, R. Smolec, J. Schwab, A. Gautschy, L. Bildsten, M. Cantiello, A. Dotter, R. Farmer, J. A. Goldberg, A. S. Jermyn, S. M. Kanbur, P. Marchant, A. Thoul, R. H. D. Townsend, W. M. Wolf, M. Zhang, and F. X. Timmes, Modules for Experiments in Stellar Astrophysics (MESA): Pulsating Variable Stars, Rotation, Convective Boundaries, and Energy Conservation, *ApJS* **243**, 10 (2019), [arXiv:1903.01426](https://arxiv.org/abs/1903.01426).
- [222] A. S. Jermyn, E. B. Bauer, J. Schwab, R. Farmer, W. H. Ball, E. P. Bellinger, A. Dotter, M. Joyce, P. Marchant, J. S. G. Mombarg, W. M. Wolf, T. L. Sunny Wong, G. C. Cinquegrana, E. Farrell, R. Smolec, A. Thoul, M. Cantiello, F. Herwig, O. Toloza, L. Bildsten, R. H. D.

Bibliography

Townsend, and F. X. Timmes, Modules for Experiments in Stellar Astrophysics (MESA): Time-dependent Convection, Energy Conservation, Automatic Differentiation, and Infrastructure, *ApJS* **265**, 15 (2023), [arXiv:2208.03651](https://arxiv.org/abs/2208.03651).

[223] J. S. Kalirai, H. B. Richer, J. Anderson, A. Dotter, G. G. Fahlman, B. M. S. Hansen, J. Hurley, I. R. King, D. Reitzel, R. M. Rich, M. M. Shara, P. B. Stetson, and K. A. Woodley, A Deep, Wide-field, and Panchromatic View of 47 Tuc and the SMC with HST: Observations and Data Analysis Methods, *AJ* **143**, 11 (2012), [arXiv:1112.1426](https://arxiv.org/abs/1112.1426).

[224] R. Goldsbury, H. B. Richer, J. Anderson, A. Dotter, A. Sarajedini, and K. Woodley, The ACS Survey of Galactic Globular Clusters. X. New Determinations of Centers for 65 Clusters, *AJ* **140**, 1830 (2010), [arXiv:1008.2755](https://arxiv.org/abs/1008.2755).

[225] J. S. Kalirai, J. Anderson, A. Dotter, H. B. Richer, G. G. Fahlman, B. M. S. Hansen, J. Hurley, I. N. Reid, R. M. Rich, and M. M. Shara, Ultra-Deep Hubble Space Telescope Imaging of the Small Magellanic Cloud: The Initial Mass Function of Stars with $M \gtrsim 1 M_\odot$, *ApJ* **763**, 110 (2013), [arXiv:1212.1159](https://arxiv.org/abs/1212.1159).

[226] T. H. Jarrett, M. E. Cluver, M. J. I. Brown, D. A. Dale, C. W. Tsai, and F. Masci, The WISE Extended Source Catalog (WXSC). I. The 100 Largest Galaxies, *ApJS* **245**, 25 (2019), [arXiv:1910.11793](https://arxiv.org/abs/1910.11793).

[227] H. B. Richer, J. Heyl, J. Anderson, J. S. Kalirai, M. M. Shara, A. Dotter, G. G. Fahlman, and R. M. Rich, A Dynamical Signature of Multiple Stellar Populations in 47 Tucanae, *ApJ* **771**, L15 (2013), [arXiv:1306.1226](https://arxiv.org/abs/1306.1226).

[228] P. B. Stetson, *User's Manual for DAOPHOT II*, Dominion Astrophysical Observatory (1998), documentation available at <http://www.star.bris.ac.uk/~mbt/daophot/>.

[229] D. Reimers, Circumstellar absorption lines and mass loss from red

Bibliography

giants., Memoires of the Societe Royale des Sciences de Liege **8**, 369 (1975).

[230] T. Bloecker, Stellar evolution of low and intermediate-mass stars. I. Mass loss on the AGB and its consequences for stellar evolution., [A&A](#) **297**, 727 (1995).

[231] J. Heyl, J. Kalirai, H. B. Richer, P. Marigo, E. Antolini, R. Goldsbury, and J. Parada, When do stars in 47 Tucanae lose their mass?, [ApJ](#) **810**, 127 (2015), [arXiv:1502.07306](#).

[232] Y. Chen, L. Girardi, X. Fu, A. Bressan, B. Aringer, P. Dal Tio, G. P astorelli, P. Marigo, G. Costa, and X. Zhang, YBC: a stellar bolometric corrections database with variable extinction coefficients. Application to PARSEC isochrones, [A&A](#) **632**, A105 (2019), [arXiv:1910.09037](#).

[233] A. Bressan, P. Marigo, L. Girardi, B. Salasnich, C. Dal Cero, S. Rubele, and A. Nanni, PARSEC: stellar tracks and isochrones with the PAdova and TRieste Stellar Evolution Code, [MNRAS](#) **427**, 127 (2012), [arXiv:1208.4498](#).

[234] F. Castelli and R. L. Kurucz, New Grids of ATLAS9 Model Atmospheres, in *Modelling of Stellar Atmospheres*, IAU Symposium, Vol. 210, edited by N. Piskunov, W. W. Weiss, and D. F. Gray (2003) p. A20, [arXiv:astro-ph/0405087](#).

[235] F. Allard, D. Homeier, B. Freytag, and C. M. Sharp, Atmospheres From Very Low-Mass Stars to Extrasolar Planets, in *EAS Publications Series*, EAS Publications Series, Vol. 57, edited by C. Reylé, C. Charbonnel, and M. Schultheis (2012) pp. 3–43, [arXiv:1206.1021](#).

[236] B. Aringer, L. Girardi, W. Nowotny, P. Marigo, and M. T. Lederer, Synthetic photometry for carbon rich giants. I. Hydrostatic dust-free models, [A&A](#) **503**, 913 (2009), [arXiv:0905.4415](#).

[237] B. Aringer, L. Girardi, W. Nowotny, P. Marigo, and A. Bressan, Synthetic photometry for M and K giants and stellar evolution: hydro-

Bibliography

static dust-free model atmospheres and chemical abundances, [MNRAS 457, 3611 \(2016\)](#), [arXiv:1601.07025](#).

[238] Y. Chen, A. Bressan, L. Girardi, P. Marigo, X. Kong, and A. Lanza, PARSEC evolutionary tracks of massive stars up to 350 M_{\odot} at metallicities $0.0001 \leq Z \leq 0.04$, [MNRAS 452, 1068 \(2015\)](#), [arXiv:1506.01681](#).

[239] M. Riello, F. de Angeli, D. W. Evans, P. Montegriffo, J. M. Carrasco, G. Busso, L. Palaversa, P. Burgess, C. Diener, M. Davidson, N. Rowell, C. Fabricius, C. Jordi, M. Bellazzini, E. Pancino, D. L. Harrison, C. Cacciari, F. van Leeuwen, N. C. Hambly, S. T. Hodgkin, P. J. Osborne, G. Altavilla, M. A. Barstow, A. G. A. Brown, M. Castellani, S. Cowell, F. de Luise, G. Gilmore, G. Giuffrida, S. Hidalgo, G. Holland, S. Marinoni, C. Pagani, A. M. Piersimoni, L. Pulone, S. Ragaini, M. Rainer, P. J. Richards, N. Sanna, N. A. Walton, M. Weiler, and A. Yoldas, VizieR Online Data Catalog: Gaia Early Data Release 3 photometric passbands (Riello+, 2021), [VizieR Online Data Catalog](#), [J/A+A/649/A3 \(2021\)](#).

[240] M. Riello, F. De Angeli, D. W. Evans, P. Montegriffo, J. M. Carrasco, G. Busso, L. Palaversa, P. W. Burgess, C. Diener, M. Davidson, N. Rowell, C. Fabricius, C. Jordi, M. Bellazzini, E. Pancino, D. L. Harrison, C. Cacciari, F. van Leeuwen, N. C. Hambly, S. T. Hodgkin, P. J. Osborne, G. Altavilla, M. A. Barstow, A. G. A. Brown, M. Castellani, S. Cowell, F. De Luise, G. Gilmore, G. Giuffrida, S. Hidalgo, G. Holland, S. Marinoni, C. Pagani, A. M. Piersimoni, L. Pulone, S. Ragaini, M. Rainer, P. J. Richards, N. Sanna, N. A. Walton, M. Weiler, and A. Yoldas, Gaia Early Data Release 3. Photometric content and validation, [A&A 649, A3 \(2021\)](#), [arXiv:2012.01916](#).

[241] D. Harbeck, G. H. Smith, and E. K. Grebel, CN Abundance Variations on the Main Sequence of 47 Tucanae, [AJ 125, 197 \(2003\)](#), [arXiv:astro-ph/0210364](#).

Bibliography

- [242] I. B. Thompson, J. Kaluzny, S. M. Rucinski, W. Krzeminski, W. Pych, A. Dotter, and G. S. Burley, The Cluster AgeS Experiment (CASE). IV. Analysis of the Eclipsing Binary V69 in the Globular Cluster 47 Tuc, [AJ **139**, 329 \(2010\)](#), [arXiv:0910.4262](#).
- [243] F. R. Ferraro, E. Lapenna, A. Mucciarelli, B. Lanzoni, E. Dalessandro, C. Pallanca, and D. Massari, Weighing Stars: The Identification of an Evolved Blue Straggler Star in the Globular Cluster 47 Tucanae, [ApJ **816**, 70 \(2016\)](#), [arXiv:1512.00649](#).
- [244] J. Parada, H. Richer, J. Heyl, J. Kalirai, and R. Goldsbury, Dynamical Estimate of Post-main-sequence Stellar Masses in 47 Tucanae, [ApJ **826**, 88 \(2016\)](#), [arXiv:1605.05740](#).
- [245] J. Christensen-Dalsgaard, On the red-giant luminosity bump, [MNRAS **453**, 666 \(2015\)](#), [arXiv:1507.05531](#).
- [246] S. Khan, O. J. Hall, A. Miglio, G. R. Davies, B. Mosser, L. Girardi, and J. Montalbán, The Red-giant Branch Bump Revisited: Constraints on Envelope Overshooting in a Wide Range of Masses and Metallicities, [ApJ **859**, 156 \(2018\)](#), [arXiv:1804.06669](#).
- [247] S. R. Bjork and B. Chaboyer, Theoretical Uncertainties in Red Giant Branch Evolution: The Red Giant Branch Bump, [ApJ **641**, 1102 \(2006\)](#), [arXiv:astro-ph/0512508](#).
- [248] A. Di Cecco, G. Bono, P. B. Stetson, A. Pietrinferni, R. Becucci, S. Cassisi, S. Degl'Innocenti, G. Iannicola, P. G. Prada Moroni, R. Buonanno, A. Calamida, F. Caputo, M. Castellani, C. E. Corsi, I. Ferraro, M. Dall'Ora, M. Monelli, M. Nonino, A. M. Piersimoni, L. Pulone, M. Romaniello, M. Salaris, A. R. Walker, and M. Zoccali, On the ΔV^{bump}_{HB} Parameter in Globular Clusters, [ApJ **712**, 527 \(2010\)](#), [arXiv:1002.2074](#).
- [249] S. Cassisi, A. Marín-Franch, M. Salaris, A. Aparicio, M. Monelli, and A. Pietrinferni, The magnitude difference between the main sequence

Bibliography

turn off and the red giant branch bump in Galactic globular clusters, *A&A* **527**, A59 (2011), arXiv:1012.0419.

[250] F. Troisi, G. Bono, P. B. Stetson, A. Pietrinferni, A. Weiss, M. Fabrizio, I. Ferraro, A. Di Cecco, G. Iannicola, R. Buonanno, A. Calamida, F. Caputo, C. E. Corsi, M. Dall’Ora, A. Kunder, M. Monelli, M. Nonino, A. M. Piersimoni, L. Pulone, M. Romaniello, A. R. Walker, and M. Zoccali, On a New Parameter to Estimate the Helium Content in Old Stellar Systems, *PASP* **123**, 879 (2011), arXiv:1106.2734.

[251] M. Joyce and B. Chaboyer, Investigating the Consistency of Stellar Evolution Models with Globular Cluster Observations via the Red Giant Branch Bump, *ApJ* **814**, 142 (2015), arXiv:1510.07648.

[252] X. Fu, A. Bressan, P. Marigo, L. Girardi, J. Montalbán, Y. Chen, and A. Nanni, New PARSEC data base of α -enhanced stellar evolutionary tracks and isochrones - I. Calibration with 47 Tuc (NGC 104) and the improvement on RGB bump, *MNRAS* **476**, 496 (2018), arXiv:1801.07137.

[253] J. B. Holberg and P. Bergeron, Calibration of Synthetic Photometry Using DA White Dwarfs, *AJ* **132**, 1221 (2006).

[254] P. Bergeron, F. Wesemael, and A. Beauchamp, Photometric Calibration of Hydrogen- and Helium-Rich White Dwarf Models, *PASP* **107**, 1047 (1995).

[255] S. Blouin, P. Dufour, and N. F. Allard, A New Generation of Cool White Dwarf Atmosphere Models. I. Theoretical Framework and Applications to DZ Stars, *ApJ* **863**, 184 (2018), arXiv:1807.06616.

[256] P. E. Tremblay, P. Bergeron, and A. Gianninas, An Improved Spectroscopic Analysis of DA White Dwarfs from the Sloan Digital Sky Survey Data Release 4, *ApJ* **730**, 128 (2011), arXiv:1102.0056.

[257] P. M. Kowalski and D. Saumon, Found: The Missing Blue Opacity in

Bibliography

Atmosphere Models of Cool Hydrogen White Dwarfs, *ApJ* **651**, L137 (2006), arXiv:astro-ph/0609720.

[258] R. W. Michie, The dynamics of spherical stellar systems, IV, *MNRAS* **126**, 499 (1963).

[259] I. R. King, The structure of star clusters. III. Some simple dynamical models, *AJ* **71**, 64 (1966).

[260] R. Goldsbury, J. Heyl, and H. Richer, Quantifying Mass Segregation and New Core Radii for 54 Milky Way Globular Clusters, *ApJ* **778**, 57 (2013), arXiv:1308.3706.

[261] J. Heyl, H. B. Richer, E. Antolini, R. Goldsbury, J. Kalirai, J. Parada, and P.-E. Tremblay, A Measurement of Diffusion in 47 Tucanae, *ApJ* **804**, 53 (2015), arXiv:1502.01890.

[262] J. Heyl, I. Caiazzo, H. Richer, J. Anderson, J. Kalirai, and J. Parada, Deep HST Imaging in 47 Tucanae: A Global Dynamical Model, *ApJ* **850**, 186 (2017), arXiv:1710.10666.

[263] K. C. Sahu et al., *WFC3 Data Handbook*, STScI, Baltimore (2021), documentation available at <https://hst-docs.stsci.edu/wfc3dhb>.

[264] M. Marinelli and L. Dressel, *WFC3 Instrument Handbook*, STScI, Baltimore (2024), documentation available at <https://hst-docs.stsci.edu/wfc3ihb>.

[265] J. Preskill, M. B. Wise, and F. Wilczek, Cosmology of the invisible axion, *Phys. Lett. B* **120**, 127 (1983).

[266] L. F. Abbott and P. Sikivie, A cosmological bound on the invisible axion, *Phys. Lett. B* **120**, 133 (1983).

[267] M. Dine and W. Fischler, The not-so-harmless axion, *Phys. Lett. B* **120**, 137 (1983).

[268] E. Witten, Some properties of O(32) superstrings, *Phys. Lett. B* **149**, 351 (1984).

Bibliography

- [269] J. P. Conlon, The QCD axion and moduli stabilisation, *J. High Energy Phys.* **2006** (05), 078, arXiv:hep-th/0602233.
- [270] A. Arvanitaki, S. Dimopoulos, S. Dubovsky, N. Kaloper, and J. March-Russell, String axiverse, *Phys. Rev. D* **81**, 123530 (2010), arXiv:0905.4720.
- [271] B. Samir Acharya, K. Bobkov, and P. Kumar, An M theory solution to the strong CP-problem, and constraints on the axiverse, *J. High Energy Phys.* **2010** (11), 105, arXiv:1004.5138.
- [272] T. Higaki and T. Kobayashi, Note on moduli stabilization, supersymmetry breaking, and axiverse, *Phys. Rev. D* **84**, 045021 (2011), arXiv:1106.1293.
- [273] M. Cicoli, M. D. Goodsell, and A. Ringwald, The type IIB string axiverse and its low-energy phenomenology, *J. High Energy Phys.* **2012** (10), 146, arXiv:1206.0819.
- [274] M. Demirtas, C. Long, L. McAllister, and M. Stillman, The Kreuzer-Skarke axiverse, *J. High Energy Phys.* **2020** (04), 138, arXiv:1808.01282.
- [275] V. M. Mehta, M. Demirtas, C. Long, D. J. E. Marsh, L. McAllister, and M. J. Stott, Superradiance in string theory, *J. Cosmol. Astropart. Phys.* **2021** (07), 033, arXiv:2103.06812.
- [276] G. G. di Cortona, E. Hardy, J. P. Vega, and G. Villadoro, The QCD axion, precisely, *J. High Energy Phys.* **2016** (01), 034, arXiv:1511.02867.
- [277] G. G. Raffelt, Axions—motivation, limits and searches, *Journal of Physics A: Mathematical and Theoretical* **40**, 6607 (2007), arXiv:hep-ph/0611118.
- [278] M. S. Turner, Axions from SN1987A, *Phys. Rev. Lett.* **60**, 1797 (1988).
- [279] A. Burrows, M. S. Turner, and R. P. Brinkmann, Axions and SN 1987A, *Phys. Rev. D* **39**, 1020 (1989).

Bibliography

- [280] G. Raffelt and D. Seckel, Bounds on exotic-particle interactions from SN1987A, *Phys. Rev. Lett.* **60**, 1793 (1988).
- [281] P. Carenza, T. Fischer, M. Giannotti, G. Guo, G. Martínez-Pinedo, and A. Mirizzi, Improved axion emissivity from a supernova via nucleon-nucleon bremsstrahlung, *J. Cosmol. Astropart. Phys.* **2019** (10), 016, [Erratum: *J. Cosmol. Astropart. Phys.* 2020 (05), E01], arXiv:1906.11844.
- [282] P. Carenza, B. Fore, M. Giannotti, A. Mirizzi, and S. Reddy, Enhanced Supernova Axion Emission and Its Implications, *Phys. Rev. Lett.* **126**, 071102 (2021), arXiv:2010.02943.
- [283] T. Fischer, P. Carenza, B. Fore, M. Giannotti, A. Mirizzi, and S. Reddy, Observable signatures of enhanced axion emission from protoneutron stars, *Phys. Rev. D* **104**, 103012 (2021), arXiv:2108.13726.
- [284] J. Keller and A. Sedrakian, Axions from cooling compact stars: Pair-breaking processes, *Nucl. Phys. A* **897**, 62 (2013), arXiv:1205.6940.
- [285] A. Sedrakian, Axion cooling of neutron stars, *Phys. Rev. D* **93**, 065044 (2016), arXiv:1512.07828.
- [286] K. Hamaguchi, N. Nagata, K. Yanagi, and J. Zheng, Limit on the axion decay constant from the cooling neutron star in Cassiopeia A, *Phys. Rev. D* **98**, 103015 (2018), arXiv:1806.07151.
- [287] M. V. Beznogov, E. Rrapaj, D. Page, and S. Reddy, Constraints on axion-like particles and nucleon pairing in dense matter from the hot neutron star in HESS J1731-347, *Phys. Rev. C* **98**, 035802 (2018), arXiv:1806.07991.
- [288] A. Sedrakian, Axion cooling of neutron stars. II. Beyond hadronic axions, *Phys. Rev. D* **99**, 043011 (2019), arXiv:1810.00190.
- [289] L. B. Leinson, Axion mass limit from observations of the neutron star in Cassiopeia A, *J. Cosmol. Astropart. Phys.* **2014** (08), 031, arXiv:1405.6873.

Bibliography

- [290] M. Arik et al. (CAST Collaboration), Search for Solar Axions by the CERN Axion Solar Telescope with ${}^3\text{He}$ Buffer Gas: Closing the Hot Dark Matter Gap, *Phys. Rev. Lett.* **112**, 091302 (2014), [arXiv:1307.1985](https://arxiv.org/abs/1307.1985).
- [291] M. Arik et al. (CAST Collaboration), New solar axion search using the CERN Axion Solar Telescope with ${}^4\text{He}$ filling, *Phys. Rev. D* **92**, 021101 (2015), [arXiv:1503.00610](https://arxiv.org/abs/1503.00610).
- [292] A. Abeln et al. (IAXO Collaboration), Conceptual design of BabyI-AXO, the intermediate stage towards the International Axion Observatory, *J. High Energy Phys.* **2021** (05), 137, [arXiv:2010.12076](https://arxiv.org/abs/2010.12076).
- [293] E. Armengaud et al. (IAXO Collaboration), Physics potential of the International Axion Observatory (IAXO), *J. Cosmol. Astropart. Phys.* **2019** (06), 047, [arXiv:1904.09155](https://arxiv.org/abs/1904.09155).
- [294] I. Avignone, F. T., R. J. Creswick, J. D. Vergados, P. Pirinen, P. C. Srivastava, and J. Suhonen, Estimating the flux of the 14.4 keV solar axions, *J. Cosmol. Astropart. Phys.* **2018** (01), 021, [arXiv:1711.06979](https://arxiv.org/abs/1711.06979).
- [295] B. M. S. Hansen, H. Richer, J. Kalirai, R. Goldsbury, S. Frewen, and J. Heyl, Constraining Neutrino Cooling Using the Hot White Dwarf Luminosity Function in the Globular Cluster 47 Tucanae, *ApJ* **809**, 141 (2015), [arXiv:1507.05665](https://arxiv.org/abs/1507.05665).
- [296] J. Isern, S. Catalan, E. Garcia-Berro, and S. Torres, Axions and the white dwarf luminosity function, *J. Phys.: Conf. Ser.* **172**, 012005 (2009), [arXiv:0812.3043](https://arxiv.org/abs/0812.3043).
- [297] K. Lodders, Solar elemental abundances, in *Oxford Research Encyclopedia of Planetary Science* (Oxford University Press, 2020) [arXiv:1912.00844](https://arxiv.org/abs/1912.00844).
- [298] R. Röhlsberger, *Nuclear Condensed Matter Physics with Synchrotron Radiation* (Springer-Verlag Berlin Heidelberg, Germany, 2004).

- [299] G. Raffelt and L. Stodolsky, Mixing of the photon with low-mass particles, *Phys. Rev. D* **37**, 1237 (1988).
- [300] D. Lai and J. Heyl, Probing axions with radiation from magnetic stars, *Phys. Rev. D* **74**, 123003 (2006), [arXiv:astro-ph/0609775](https://arxiv.org/abs/astro-ph/0609775).
- [301] R. Gill and J. S. Heyl, Constraining the photon-axion coupling constant with magnetic white dwarfs, *Phys. Rev. D* **84**, 085001 (2011), [arXiv:1105.2083](https://arxiv.org/abs/1105.2083).
- [302] A. J. Millar, S. Baum, M. Lawson, and M. C. D. Marsh, Axion-photon conversion in strongly magnetised plasmas, *J. Cosmol. Astropart. Phys.* **2021** (11), 013, [arXiv:2107.07399](https://arxiv.org/abs/2107.07399).
- [303] V. Anastassopoulos et al. (CAST Collaboration), New CAST Limit on the Axion-Photon Interaction, *Nature Phys.* **13**, 584 (2017), [arXiv:1705.02290](https://arxiv.org/abs/1705.02290).
- [304] I. Caiazzo, K. Burdge, L. Ferrario, J. Fuller, B. T. Gaensicke, J. Heyl, A. Kawka, S. R. Kulkarni, T. A. Prince, H. B. Richer, S. Vennes, and D. T. Wickramasinghe, The radius and magnetic field structure of the smallest white dwarf, HST Proposal. Cycle 29, ID. #16753 (2021).
- [305] P. Carenza, M. Giannotti, J. Isern, A. Mirizzi, and O. Straniero, Axion Astrophysics, arXiv e-prints (2024), [arXiv:2411.02492](https://arxiv.org/abs/2411.02492).
- [306] J. M. Burgers, *Flow Equations for Composite Gases* (Academic Press, Inc., New York, 1969).
- [307] A. A. Thoul, J. N. Bahcall, and A. Loeb, Element Diffusion in the Solar Interior, *ApJ* **421**, 828 (1994), [arXiv:astro-ph/9304005](https://arxiv.org/abs/astro-ph/9304005).
- [308] J. Iben, I. and J. MacDonald, The effects of diffusion due to gravity and due to composition gradients on the rate of hydrogen burning in a cooling degenerate dwarf. I - The case of a thick helium buffer layer, *ApJ* **296**, 540 (1985).

Appendix A

Supplemental material for the cooling of old white dwarfs in 47 Tucanae

A.1 SMC Contamination

Let the subscript “D” indicate the full **SHARP**-cleaned dataset before any proper motion cuts or CMD cuts, and let the subscripts “S” and “T” denote the subsets of the true SMC stars and the true 47 Tuc stars, respectively, within the full dataset. If a proper motion cut has been applied to one of these populations (“D”, “S”, or “T”), that will be denoted with the subscript “ps” for the SMC proper motion cut and “pt” for the 47 Tuc proper motion cut. If a CMD cut has been applied, that will be denoted with the subscript “cs” for the SMC CMD cut, “cw” for the 47 Tuc white dwarf CMD cut, and “cm” for the 47 Tuc main-sequence CMD cut.

The 47 Tuc white dwarf sample used in the main analysis of Chapter 6 is selected by applying both the 47 Tuc proper motion cut and the 47 Tuc CMD cut, and the number of stars in this sample is $N_{D,pt,cw}$. Assuming SMC stars are the only contaminants in this sample after **SHARP** cleaning, then

$$N_{D,pt,cw} = N_{T,pt,cw} + N_{S,pt,cw}, \quad (\text{A.1})$$

where $N_{T,pt,cw}$ is the true number of 47 Tuc white dwarfs in the sample and $N_{S,pt,cw}$ is the true number of SMC contaminants in the sample. We want to estimate $N_{S,pt,cw}$ using numbers that can actually be calculated by applying cuts to the full dataset.

A.1. SMC Contamination

If the SMC proper motion cut is chosen such that all of the objects selected by this cut, when applied to the full dataset, are actually SMC stars, then

$$N_{S,ps,cs} = N_{D,ps,cs}, \quad (A.2)$$

$$N_{S,ps,cw} = N_{D,ps,cw}, \quad (A.3)$$

where $N_{S,ps,cs}$ is the number of objects selected when both the SMC proper motion cut and SMC CMD cut are applied to the true SMC stars, $N_{D,ps,cs}$ is the number of objects selected when both the SMC proper motion cut and SMC CMD cut are applied to the full dataset, $N_{S,ps,cw}$ is the number of objects selected when both the SMC proper motion cut and 47 Tuc white dwarf CMD cut are applied to the true SMC stars, and $N_{D,ps,cw}$ is the number of objects selected when both the SMC proper motion cut and 47 Tuc white dwarf CMD cut are applied to the full dataset.

Likewise, if the SMC CMD cut is chosen such that all of the objects selected by this cut, when applied to the full dataset, are actually SMC stars, then

$$N_{S,pt,cs} = N_{D,pt,cs}, \quad (A.4)$$

where $N_{S,pt,cs}$ is the number of objects selected when both the 47 Tuc proper motion cut and SMC CMD cut are applied to the full dataset and $N_{D,pt,cs}$ is the number of objects selected when both the 47 Tuc proper motion cut and SMC CMD cut are applied to the true SMC stars.

In practice, Eqs. (A.2) to (A.4) are only approximately true, as some 47 Tuc stars could potentially survive the relevant cuts, particularly the SMC CMD cut in combination with the 47 Tuc proper motion. This makes $N_{D,pt,cs}$ in particular an upper limit on $N_{S,pt,cs}$, which will translate to our estimate of $N_{S,pt,cw}$ really being an upper limit on $N_{S,pt,cw}$. However, the SMC cuts are chosen such that Eqs. (A.2) to (A.4) are good approximations.

Assuming the ratio of SMC stars that survive the 47 Tuc white dwarf CMD cut to SMC stars that survive the SMC CMD cut is the same for both

the 47 Tuc and SMC proper motion cuts, as it should be, then

$$N_{S,pt,cw} = N_{S,pt,cs} \frac{N_{S,ps,cw}}{N_{S,ps,cs}} \quad (A.5)$$

$$\approx N_{D,pt,cs} \frac{N_{D,ps,cw}}{N_{D,ps,cs}}. \quad (A.6)$$

The expression on the right-hand side of Eq. (A.6) is equivalent to the definition of N_{contam} in Eq. (6.23), used in Section 6.4.3 to estimate the number of SMC contaminants in the 47 Tuc white dwarf data space.

Note that the numbers corresponding to samples selected using a cut designed to select 47 Tuc stars, i.e. $N_{D,pt,cs}$ (47 Tuc proper motion cut) and $N_{D,ps,cw}$ (47 Tuc white dwarf CMD cut), both occur in the numerator. These are the numbers for which 47 Tuc stars are most likely to be miscounted as SMC stars. Whereas $N_{D,ps,cs}$, which appears in the denominator of Eq. (A.6), is calculated from the sample that is least likely to contain any misclassified 47 Tuc stars as it is selected by applying both the SMC proper motion cut and the SMC CMD cut. Thus, Eq. (A.6) is really an upper limit on $N_{S,pt,cw}$, though this upper limit should also be close to the value of $N_{S,pt,cw}$. As the value of N_{contam} is found in Section 6.4.3 to be small at all magnitudes of interest, knowing the upper limit on $N_{S,pt,cw}$ is sufficient for our purpose of determining that the number of SMC contaminants in the 47 Tuc white dwarf data space after proper motion cleaning is small enough to be neglected.

A.2 Element Diffusion

The implementation of element diffusion in MESA assumes an ideal gas law through the form of the diffusion equations that MESA solves. While this is a reasonable approximation for many regimes of stellar evolution, we are interested in white dwarf cooling regimes where the core is in a liquid or even solid state. As part of our parameter grid of cooling models, we thus consider a modified form of element diffusion that allows for non-ideal gas behaviour. To understand how this modified diffusion is implemented, it is

useful to first review the treatment of diffusion in standard MESA.

A.2.1 Diffusion in MESA

MESA’s implementation of element diffusion is described in detail in the instrument papers Paxton *et al.* [219] (hereafter [MESA III](#)) and Paxton *et al.* [220] (hereafter [MESA IV](#)), more specifically Section 9 of [MESA III](#) and both Section 3 and Appendix C of [MESA IV](#). Some changes were made to MESA’s treatment of element diffusion between [MESA III](#) and [MESA IV](#) to properly account for degenerate electrons, for which an ideal gas approximation is not appropriate. However, this update did not address the potential inaccuracies of the ideal gas approximation for other particle species. The updated code documented in [MESA IV](#) directly solves the equations of Burgers [306] describing diffusion, rather than the rescaled versions of those equations introduced by Thoul *et al.* [307] that were solved in earlier versions of MESA and documented in [MESA III](#). As part of this update, one of Burgers’ equations for electrons was also dropped by treating the gravitational acceleration as a fixed input to the diffusion equations instead of a variable that needed to be solved for in the diffusion routine. Note that despite these updates, many of the details of element diffusion described by [MESA III](#) still apply for [MESA IV](#). Of particular note is that the schematic derivation of Burgers’ equations given in Section 9.1.1 of [MESA III](#) is still relevant for [MESA IV](#).

The equations of Burgers [306] that describe element diffusion and heat flow in a plasma, under the assumptions³⁷ of spherical symmetry (appropriate for MESA) and that all of the particles in the plasma obey the ideal gas

³⁷It has also been assumed that there is no magnetic field.

law, are

$$\begin{aligned} \frac{dP_s}{dr} + \rho_s g - \rho_{es} E \\ = \sum_{t \neq s} K_{st} (w_t - w_s) + \sum_{t \neq s} K_{st} z_{st} \left(\frac{m_t r_s - m_s r_t}{m_s + m_t} \right), \end{aligned} \quad (\text{A.7})$$

$$\begin{aligned} \frac{5}{2} n_s k_B \frac{dT}{dr} \\ = -\frac{2}{5} K_{ss} z_{ss}'' r_s - \frac{5}{2} \sum_{t \neq s} K_{st} z_{st} \frac{m_t}{(m_s + m_t)} (w_t - w_s) \\ - \sum_{t \neq s} K_{st} \left[\frac{(3m_s^2 + m_t^2 z_{st}')}{(m_s + m_t)^2} + \frac{4}{5} \frac{m_s m_t}{(m_s + m_t)^2} z_{st}'' \right] r_s \\ + \sum_{t \neq s} K_{st} \frac{m_s m_t}{(m_s + m_t)^2} \left(3 + z_{st}' - \frac{4}{5} z_{st}'' \right) r_t, \end{aligned} \quad (\text{A.8})$$

where the subscripts s and t denote plasma species (which can be either electrons or ions). The parameters in these equations are defined as follows. Independent of species, r is the radial distance from the centre of the white dwarf, $g(r)$ is the gravitational acceleration, $E(r)$ is the quasi-static electric field, k_B is Boltzmann's constant, and $T(r)$ is the temperature, where it is assumed that all species are in thermal equilibrium and thus have the same temperature. For a specific species labelled by the index s , the parameter $P_s(r)$ is the partial pressure, $\rho_s(r)$ is the mass density, $\rho_{es}(r)$ is the charge density (given by $\rho_{es} = e Z_s n_s$, where Z_s is the charge number and e is the charge of an electron), $n_s(r)$ is the number density, and m_s is the mass of one particle. The parameters K_{st} , z_{st} , z_{st}' , and z_{st}'' are resistance coefficients, which arise due to particle collisions, w_s are the diffusion velocities, and r_s are the heat flow vectors. These resistance coefficients, diffusion velocities, and heat flow vectors are defined below, and the interpretation of these quantities is discussed.

A.2. Element Diffusion

The resistance coefficients are defined through the equations

$$K_{st} = K_{ts} = \frac{2}{3} n_s n_t \mu_{st} \alpha_{st} \Sigma_{st}^{(11)}, \quad (\text{A.9})$$

$$\Sigma_{st}^{(12)} / \Sigma_{st}^{(11)} = \frac{5}{2} (1 - z_{st}), \quad (\text{A.10})$$

$$\Sigma_{st}^{(13)} / \Sigma_{st}^{(11)} = \frac{25}{4} - \frac{25}{2} z_{st} + \frac{5}{2} z'_{st}, \quad (\text{A.11})$$

$$\Sigma_{st}^{(22)} / \Sigma_{st}^{(11)} = z''_{st}, \quad (\text{A.12})$$

where $\mu_{st} = m_s m_t / (m_s + m_t)$ is the reduced mass of species s and t , $\alpha_{st}^2 = 2k_B T / \mu_{st}$, and $\Sigma_{st}^{(\ell j)}$ are cross-sections that result from taking moments of the collision terms in the Boltzmann equation in Burgers' derivation of Eqs. (A.7) and (A.8). Note that Burgers [306] derived Eqs. (A.7) and (A.8) by taking moments of the Boltzmann equation using a 13-moment approximation as a closure scheme and assuming an approximately Maxwellian distribution function.

The cross-sections that appear in the definitions of the resistance terms, i.e. in Eqs. (A.9) to (A.12), are given by the expression

$$\Sigma_{st}^{(\ell j)} = \frac{4\pi}{\pi^{3/2}} \int_0^\infty d\nu \exp\left(\frac{-\nu^2}{\alpha_{st}^2}\right) \frac{\nu^{2j+3}}{\alpha_{st}^{2j+4}} S_{st}^{(\ell)} \quad (\text{A.13})$$

in terms of the collision integrals

$$S_{st}^{(\ell)} = 2\pi \int_0^\infty db \left(1 - \cos^\ell \chi_{st}\right) b, \quad (\text{A.14})$$

where ν is the relative velocity of the colliding particles, b is the impact parameter, and $\chi_{st}(\nu, b)$ is the angle of deviation, which depends on both ν and b .

The diffusion velocities and heat flow vectors are defined relative to the mean mass flow velocity of the gas as a whole. Let \mathbf{u}_s be the mean velocity of species s , and let \mathbf{u} be the mean velocity of the gas, averaged over all plasma species. Both \mathbf{u}_s and \mathbf{u} are functions of position and time. For a species s described by the distribution function f_s , the i th Cartesian component of

A.2. Element Diffusion

the mean species velocity is formally defined as

$$u_{si} = \frac{1}{n_s} \int d\xi \xi_i f_s, \quad (\text{A.15})$$

where $f_s(\mathbf{x}, \boldsymbol{\xi}, t)$ is a function of the three-dimensional position vector \mathbf{x} (with components x_i), the three-dimensional velocity vector $\boldsymbol{\xi}$ (with components ξ_i), and time t . The mean velocity of the gas is

$$\mathbf{u} = \frac{1}{\rho} \sum_s \rho_s \mathbf{u}_s, \quad (\text{A.16})$$

where $\rho = \sum_s \rho_s$ is the mass density of the whole gas, whereas ρ_s is the mass density of a particular species.

The three-dimensional diffusion velocities are defined as

$$\mathbf{w}_s = \mathbf{u}_s - \mathbf{u}, \quad (\text{A.17})$$

which is simply the mean species velocity relative to the mean velocity of the plasma. The residual heat flow vectors are defined by Burgers [306] in terms of their Cartesian components as

$$r_{si} = \left[\frac{m_s}{2n_s k_B T} \int d^3 \boldsymbol{\xi} (\xi_i - u_i) |\boldsymbol{\xi} - \mathbf{u}|^2 f_s \right] - \frac{5}{2} w_{si}. \quad (\text{A.18})$$

The heat flow vectors \mathbf{r}_s represent the kinetic energy carried by diffusing particles along a temperature gradient.

The form of Burgers' equations given by Eqs. (A.7) and (A.8) assumes spherical symmetry (as does MESA). Under this assumption, $\mathbf{w}_s = w_s \hat{\mathbf{r}}$ and $\mathbf{r}_s = r_s \hat{\mathbf{r}}$, and thus Eqs. (A.7) and (A.8) depend on the magnitudes of the diffusion velocities and heat flow vectors, $w_s = |\mathbf{w}_s|$ and $r_s = |\mathbf{r}_s|$. For a plasma with N_S species, Burgers' equations represent a set of $2N_S$ equations in $2N_S + 2$ unknowns (w_s for N_S species, r_s for N_S species, g , and E). This system of equations is closed by including the two additional constraints of

A.2. Element Diffusion

no net flow of mass or electric current,

$$\sum_s \rho_s w_s = 0, \quad (\text{A.19})$$

$$\sum_s \rho_{es} w_s = 0. \quad (\text{A.20})$$

Taking the set of $2N_S$ equations given by Eqs. (A.7) and (A.8) together with Eqs. (A.19) and (A.20) gives a set of $2N_S + 2$ equations in $2N_S + 2$ unknowns.

To accommodate the case of degenerate electrons, the copy of Eq. (A.7) for electrons is dropped from this system of equations. Dropping Eq. (A.7) for electrons circumvents the difficulty presented by the dP_e/dr term when trying to apply Eq. (A.7) for degenerate electrons, in which case dP_e/dr does not take the simple analytic form of an ideal gas. However, dropping this equation reduces the total number of equations in the system to $2N_S + 1$. To close the system of equations, the number of unknowns is also reduced by one by treating g as a fixed input into the MESA diffusion routine, given by $g = Gm/r^2$, instead of treating g as a variable that needs to be solved for.

As of [MESA IV](#), the MESA implementation of diffusion then casts Eq. (A.7) for the remaining $N_S - 1$ ion species into the form

$$\begin{aligned} n_s k_B T \frac{d \ln T}{dr} + n_s k_B T \frac{d \ln n_s}{dr} + n_s A_s m_p g - n_s Z_s e E \\ = \sum_{t \neq s} K_{st} (w_t - w_s) + \sum_{t \neq s} K_{st} z_{st} \left(\frac{A_t r_s - A_s r_t}{A_s + A_t} \right) \end{aligned} \quad (\text{A.21})$$

using the ideal gas law, $P_s = n_s k_B T$, and the relations $\rho = n_s A_s m_p$ and $\rho_{es} = n_s Z_s e$, where A_s is the atomic mass number and m_p is the proton mass. Also substituting these density relations in the conservation equa-

tions, Eqs. (A.19) and (A.20), gives

$$\sum_s n_s A_s w_s = 0, \quad (\text{A.22})$$

$$\sum_s n_s Z_s w_s = 0. \quad (\text{A.23})$$

The implementation of the ideal gas law in Eq. (A.21) follows the procedure of Burgers [306] in deriving the left-hand side of Eq. (A.8), where the temperature of each species was defined as $T_s \equiv P_s/(n_s k_B)$ and thermal equilibrium was assumed between all species so that $T \equiv T_s$. Note that the quantities P_s and n_s are defined in terms of moments of a Maxwellian distribution function. If the distribution function of a species differs sufficiently from the assumed Maxwellian form, then the ideal gas law no longer holds, and if the species remains in thermal equilibrium with the surroundings while failing to satisfy this law, then the Burgers treatment does not assign the correct temperature in Eqs. (A.8) and (A.21). MESA avoids this issue with Eq. (A.21) in the case of degenerate electrons, for which the Fermi-Dirac distribution no longer reduces to a Maxwellian form, by simply not including Eq. (A.21) for electrons in the system of equations. But this ideal gas assumption persists for the ion species.

This ideal gas assumption is also still present for both ions and electrons in the heat flow equation, Eq. (A.8). As discussed in [MESA IV](#), however, temperature gradients in the core of a white dwarf are typically small, and thus the heat flow vectors become negligible ($r_s \ll w_s$ for all w_s) in white dwarf cores, which is where electrons are most strongly degenerate (and where the liquid or solid state of the ions becomes a concern). In this case, the N_S equations given by Eq. (A.8) can be neglected and the $N_S - 1$ equations given by Eq. (A.21) can be simplified to

$$k_B T \left(\frac{d \ln T}{dr} + \frac{d \ln n_s}{dr} \right) + A_s m_p g = Z_s e E + \frac{1}{n_s} \sum_{t \neq s} K_{st} (w_t - w_s), \quad (\text{A.24})$$

where the heat flow vectors have been set to $r_s = 0$. As the dependence on r_s has been removed from Eq. (A.24), the set of these equations in combination

with the two conservation relations Eqs. (A.19) and (A.20) is a set of $N_S + 1$ equations in $N_S + 1$ unknowns.

The MESA diffusion routine packs Eqs. (A.8) and (A.21) to (A.23) into a single matrix equation modelled after the approach of Thoul *et al.* [307] but without the re-scaling of that approach. This matrix formulation is documented in Appendix C of [MESA IV](#).

In solving for the diffusion velocities, MESA separates the diffusion velocities into two terms following the approach of Iben and MacDonald [308]

$$w_i = w_i^g - \sum_j \sigma_{ij} \frac{d \ln C_j}{dr}, \quad (\text{A.25})$$

where $C_j \equiv n_j/n_e$ is the concentration of species j . The term w_i^g captures the effects of gravitational settling, while the other term captures the effects of ordinary diffusion.

A.2.2 Modified Diffusion

We modify diffusion in MESA by multiplying the velocity diffusion term for concentration diffusion (i.e. ordinary diffusion) in Eq. (A.25) by a position-dependent correction factor f_{SIG} that can account for non-ideal gas effects of the ions, so that Eq. (A.25) becomes

$$w_i = w_i^g - f_{\text{SIG}} \sum_j \sigma_{ij} \frac{d \ln C_j}{dr}. \quad (\text{A.26})$$

We use a correction factor of the general form

$$f_{\text{SIG}}(r) = \frac{f_{\text{SIG}}^0}{1 + \frac{\Gamma(r)}{A}}, \quad (\text{A.27})$$

which is parameterised by the constants A and f_{SIG}^0 and is a function of the plasma coupling parameter Γ . We take the fiducial values of these parameters to be $A = 0.0625$ and $f_{\text{SIG}}^0 = 1$. We performed a small suite of molecular dynamics simulations of hydrogen-helium plasmas with different values of Γ to obtain this function form and estimate the parameter A .

MESA provides a control parameter `diffusion_SIG_factor` that multiplies the term in Eq. (A.25) for concentration diffusion, $\sum_j \sigma_{ij} \frac{d \ln C_j}{dr}$, by a user-specified constant factor, which by default is set to unity. To allow for possible non-ideal gas effects of the ions, we want to use a re-scaling factor similar to `diffusion_SIG_factor`, but it should have the flexibility to vary with position. In the MESA routine that implements diffusion, MESA privately uses another parameter `SIG_factor` that is set to the value of `diffusion_SIG_factor` for each cell in the position grid, and it is this parameter `SIG_factor` that is actually used to multiply the concentration diffusion term. This enables us to implement our modification to diffusion in MESA by modifying what value is assigned to `SIG_factor`, with `SIG_factor` = f_{SIG} and `diffusion_SIG_factor` = f_{SIG}^0 related through Eq. (A.27) instead of being equal at all positions.

In MESA revision r15140, the relevant modules through which diffusion is implemented are the private modules `element_diffusion`, `diffusion`, and `diffusion_support`. The module `element_diffusion` uses subroutines defined in the module `diffusion`, which in turn use subroutines defined in the module `diffusion_support`. The parameter `SIG_factor` that we want to modify is defined in the `get_matrix_coeffs` subroutine contained in the private `diffusion_support` module. To avoid modifying private files in MESA (which would change the whole installation), we make local copies of the relevant code and make changes to the local version, which is then implemented via the `other_diffusion` subroutine of the module `run_stars_extras`.

Appendix B

Supplemental material for axion constraints from white dwarfs in 47 Tucanae

B.1 Overview

In Appendix B.2 and Appendix B.3, results are presented for analyses analogous to the analysis in Chapter 7 but performed separately for each of the data sets used in the combined analysis of Chapter 7. The results of fitting only the WFC3/UVIS data are presented in Appendix B.2, while the results of fitting only the ACS/WFC data presented in Appendix B.3. For each of these data sets, the analysis of that data in isolation is performed for three different cases of priors, which are summarised in Table B.1. In the first case, all priors are taken to be uniform. Case 2 (in comparison to case 1) tests the effect of using the Gaussian prior from Goldsbury *et al.* [215] for the birthrate, with the other parameters taken to have the same uniform priors as in case 1. These birthrate priors are $\dot{N}_{\text{WFC3}} = 8.2 \pm 0.3$ Myr for the WFC3 data and $\dot{N}_{\text{ACS}} = 2.61 \pm 0.07$ Myr for the ACS data. Case 3 (in comparison to case 1) tests the effect of using the results of Chapter 6 for the priors on M_{WD} and $\log_{10} q_H$, with the birthrate taken to have the same uniform prior as in case 1. In all three cases, a uniform prior is used for m_a . Note that case 3 is the combination of priors that was used in Chapter 7 for the combined (WFC3 and ACS) analysis. In Appendix B.4, the final results of the WFC3 only and ACS only analyses for case 3 of the priors are compared to the results of the combined analysis from Chapter 7.

B.1. Overview

	\dot{N}	M_{WD}	$\log_{10} q_H$
Case 1	Uniform	Uniform	Uniform
Case 2	Gaussian	Uniform	Uniform
Case 3	Uniform	Chapter 6	Chapter 6

Table B.1: Summary of priors used in each case for the analysis of only the WFC3 data and only the ACS data. The Gaussian prior for the birthrate (\dot{N}) is the same prior used for the corresponding field in Goldsbury *et al.* [215]. The Chapter 6 priors for the white dwarf mass (M_{WD}) and envelope thickness parameter ($\log_{10} q_H$) refer to the joint posterior distribution from Chapter 6 for the standard diffusion scenario after marginalising over the birthrate.

For the WFC3 only analyses in Appendix B.2, plots of the posterior distributions and credible regions for each case of priors are shown in Appendix B.2.1. Three figures are shown for each case in Appendix B.2.1. The first figure shows slices of the three-dimensional joint posterior distribution after marginalising over the birthrate, the second figure shows the two-dimensional joint credible regions for $\log_{10} q_H$ and m_a , and the third figures shows the one-dimensional marginal posterior distributions of all the parameters. These figures are analogous to Fig. 7.3, Fig. 7.4, and Fig. 7.5, respectively, from Chapter 7. The corresponding figures for the WFC3 only analyses are: i) Fig. B.1, Fig. B.2, and Fig. B.3 for WFC3 case 1, ii) Fig. B.4, Fig. B.5, and Fig. B.6 for WFC3 case 2, and iii) Fig. B.7, Fig. B.8, and Fig. B.9 for WFC3 case 3. Note that for Figs. B.1, B.4 and B.7 showing the joint posterior distribution (analogous to Fig. 7.3), all three M_{WD} slices are shown instead of just the two slices shown in Fig. 7.3. In Appendix B.2.2, plots of the (inverse) cumulative luminosity functions and cumulative radial distributions for the best-fitting model of each case are shown. The cumulative luminosity functions for the best-fitting models of all three cases are compared in Fig. B.10, which shows a set of plots analogous to (the top row of) Fig. 7.7. Each row of Fig. B.10 corresponds to a separate case. Likewise, Fig. B.11 shows a set of plots analogous to Fig. 7.6, with each row showing the cumulative radial distribution for the best-fitting model of a

B.1. Overview

particular case. Tables summarising the best-fitting model results for each WFC3 case, Table B.2 and Table B.3, are also given in Appendix B.2.2. Table B.2 gives the best-fitting model parameters for each case. Table B.3 gives the p -values from KS tests comparing the distributions predicted by the best-fitting model of each case to the corresponding empirical distribution. Finally, the results of the WFC3 only analyses are also discussed in Appendix B.2.2.

The structure of Appendix B.3 for the ACS only analyses is the same as Appendix B.2. Plots of the posterior distributions and credible regions are given in Appendix B.3.1 for each case. Like for the WFC3 only analysis, the three figures shown for each case of the ACS only analysis in Appendix B.3.1 are analogous to Fig. 7.3, Fig. 7.4, and Fig. 7.5 from Chapter 7. The equivalent figures for the ACS only analyses are: i) Fig. B.12, Fig. B.13, and Fig. B.14 for ACS case 1, ii) Fig. B.15, Fig. B.16, and Fig. B.17 for ACS case 2, and iii) Fig. B.18, Fig. B.19, and Fig. B.20 for ACS case 3. Results for the best-fitting models of all of these cases are given in Appendix B.3.2, including plots of the cumulative luminosity functions and tables summarising the results. The cumulative luminosity functions for the best-fitting models of the three ACS cases are compared in Fig. B.21. This plot is analogous to Fig. B.10 of the WFC3 only analysis and contains a set of plots analogous to (the bottom row of) Fig. 7.7, with each row of Fig. B.21 corresponding to a different case. Note that for the ACS only analysis, there is no radial distribution figure (like Fig. B.11 or Fig. 7.6) because the ACS analysis does not include R dependence. Table B.4 gives the best-fitting model parameters for the ACS cases, and Table B.5 gives the p -values of KS tests comparing the best-fitting model distributions to the corresponding empirical distributions. Table B.4 and Table B.5 are both given in Appendix B.3.2, along with a discussion of the results.

B.2 WFC3/UVIS Data Only

B.2.1 Posterior Distributions

B.2.1.1 Case 1

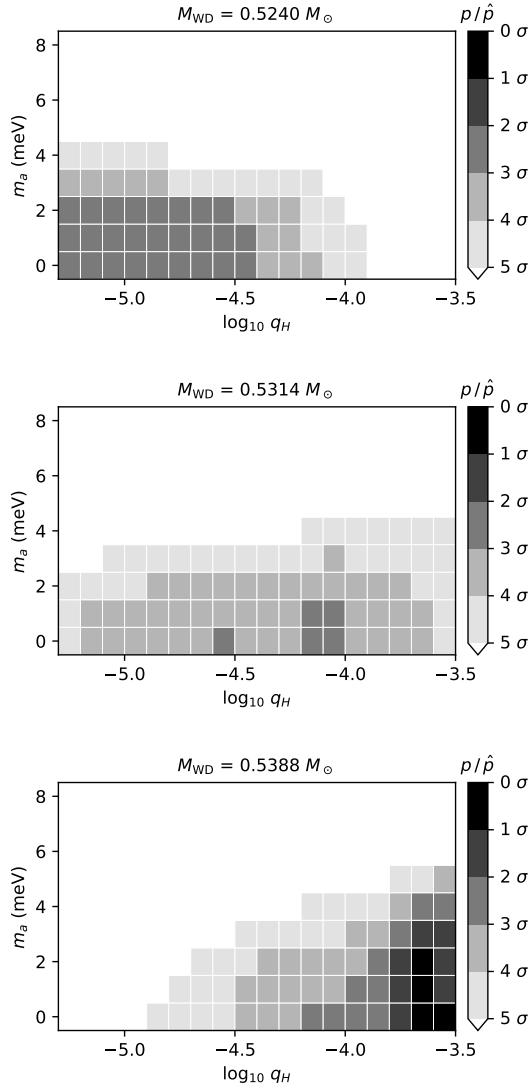


Figure B.1: WFC3, case 1: joint posterior probability density distribution after marginalising over the birthrate.

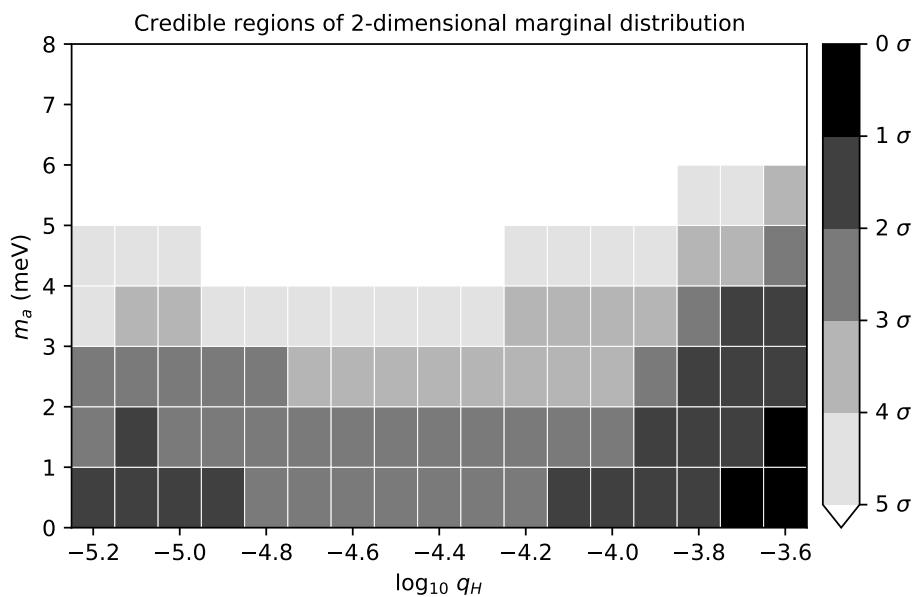


Figure B.2: WFC3, case 1: two-dimensional joint credible regions of axion mass (m_a) and envelope thickness (q_H) after marginalising over the other parameters.

B.2. WFC3/UVIS Data Only

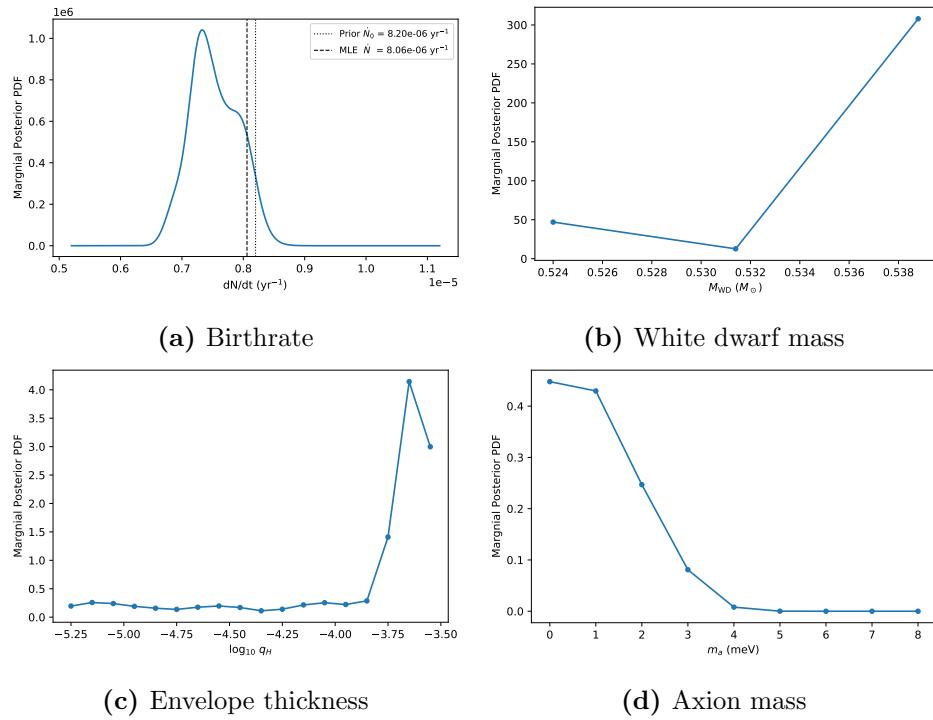


Figure B.3: WFC3, case 1: one-dimensional posterior density distributions for each parameter after marginalising over all other model parameters.

B.2.1.2 Case 2

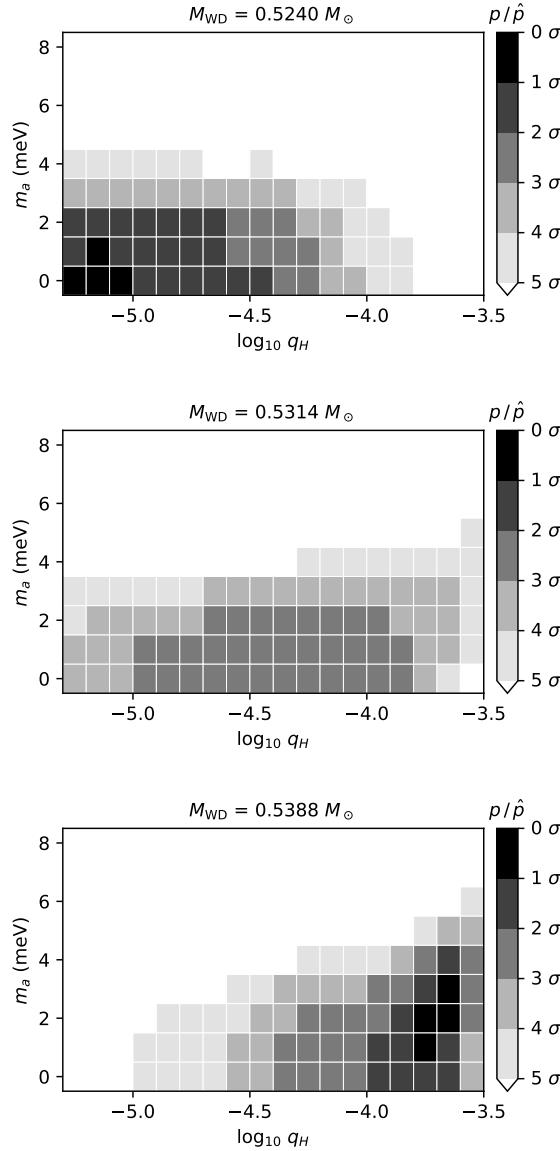


Figure B.4: WFC3, case 2: joint posterior probability density distribution after marginalising over the birthrate.

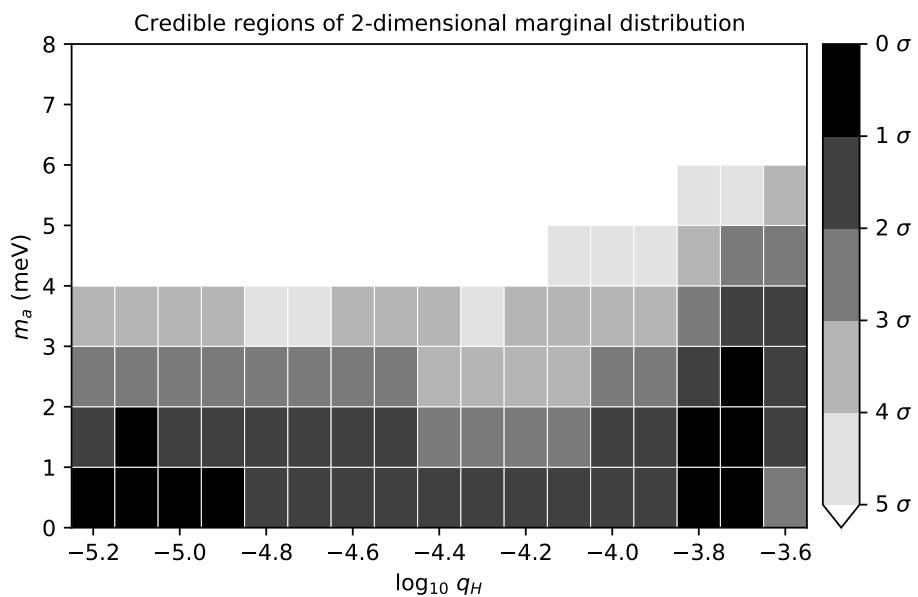


Figure B.5: WFC3, case 2: two-dimensional joint credible regions of axion mass (m_a) and envelope thickness (q_H) after marginalising over the other parameters.

B.2. WFC3/UVIS Data Only

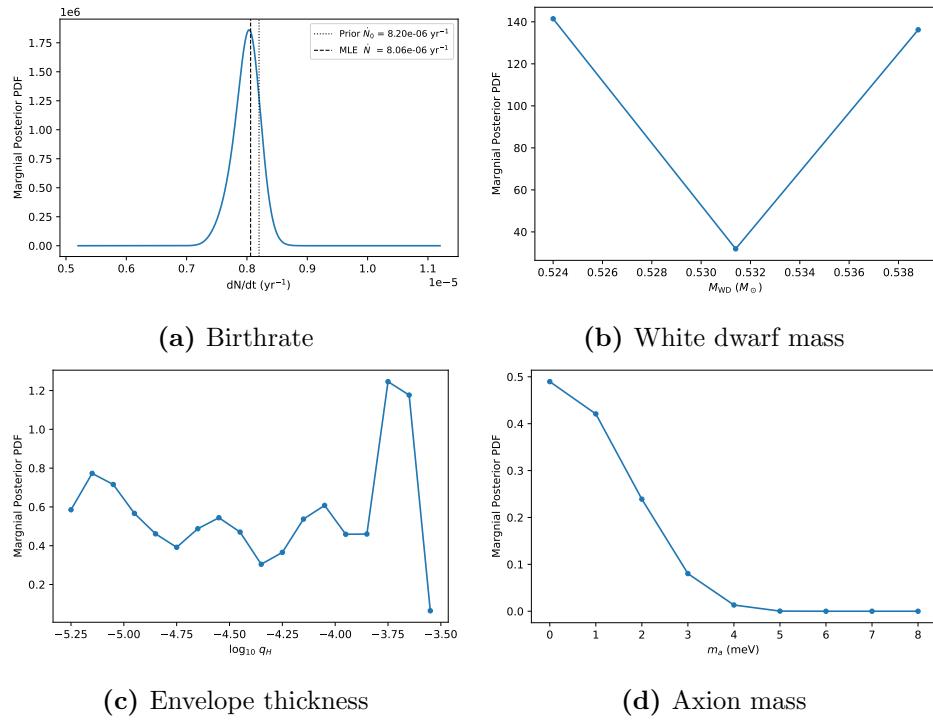


Figure B.6: WFC3, case 2: one-dimensional posterior density distributions for each parameter after marginalising over all other model parameters.

B.2.1.3 Case 3

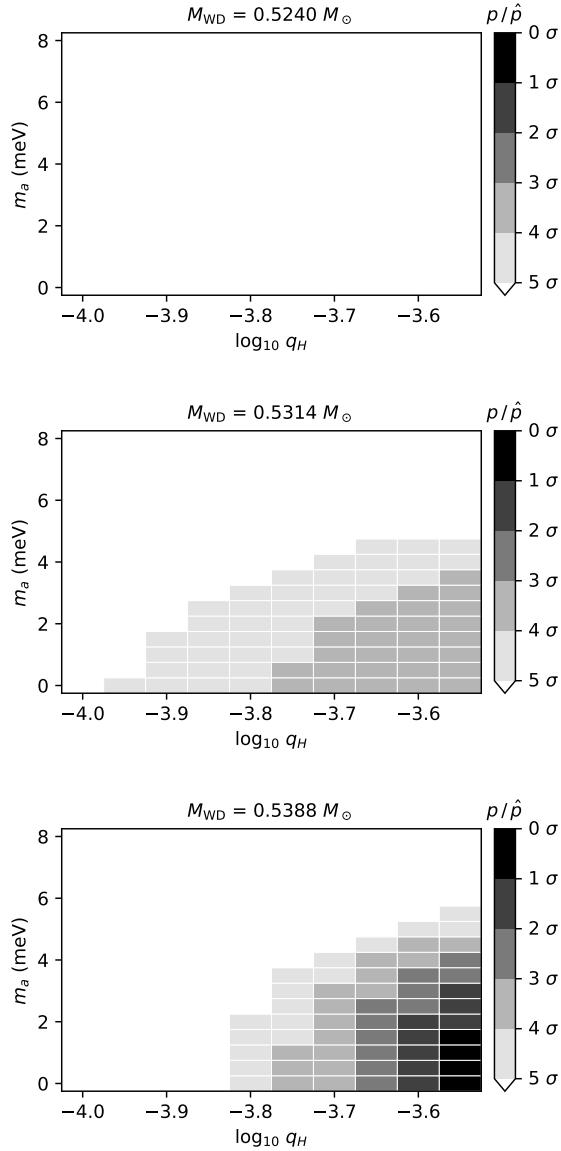


Figure B.7: WFC3, case 3: joint posterior probability density distribution after marginalising over the birthrate.

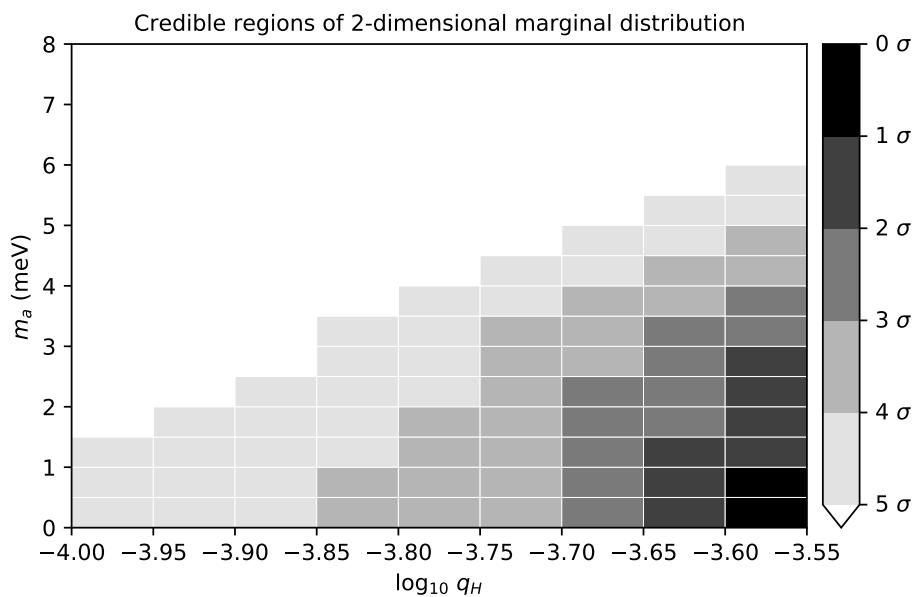


Figure B.8: WFC3, case 3: two-dimensional joint credible regions of axion mass (m_a) and envelope thickness (q_H) after marginalising over the other parameters.

B.2. WFC3/UVIS Data Only

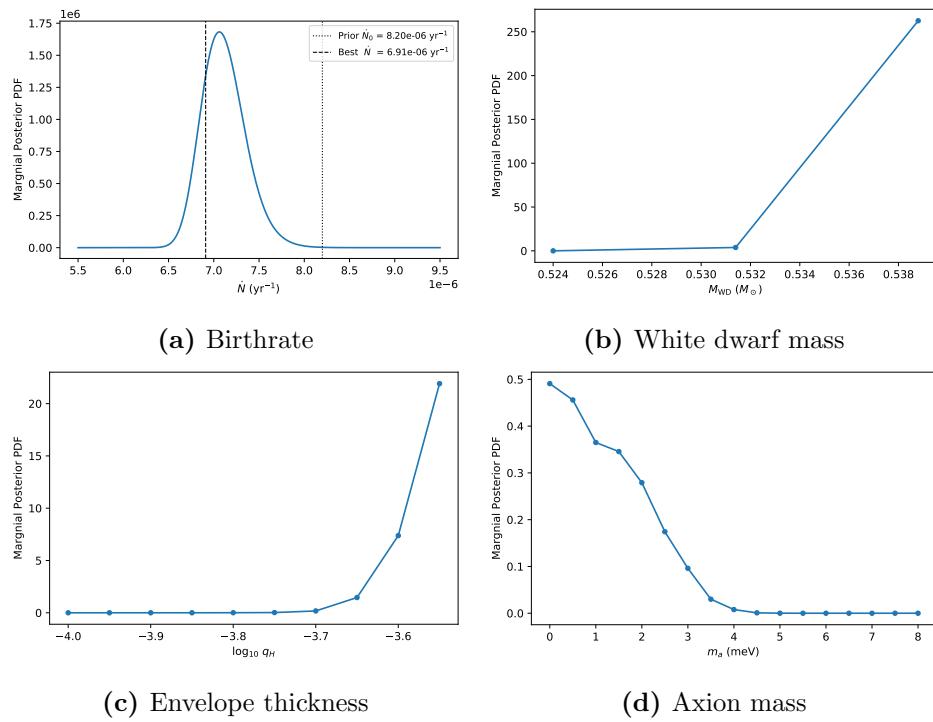


Figure B.9: WFC3, case 3: one-dimensional posterior density distributions for each parameter after marginalising over all other model parameters.

B.2.2 Best-Fitting Models Comparison

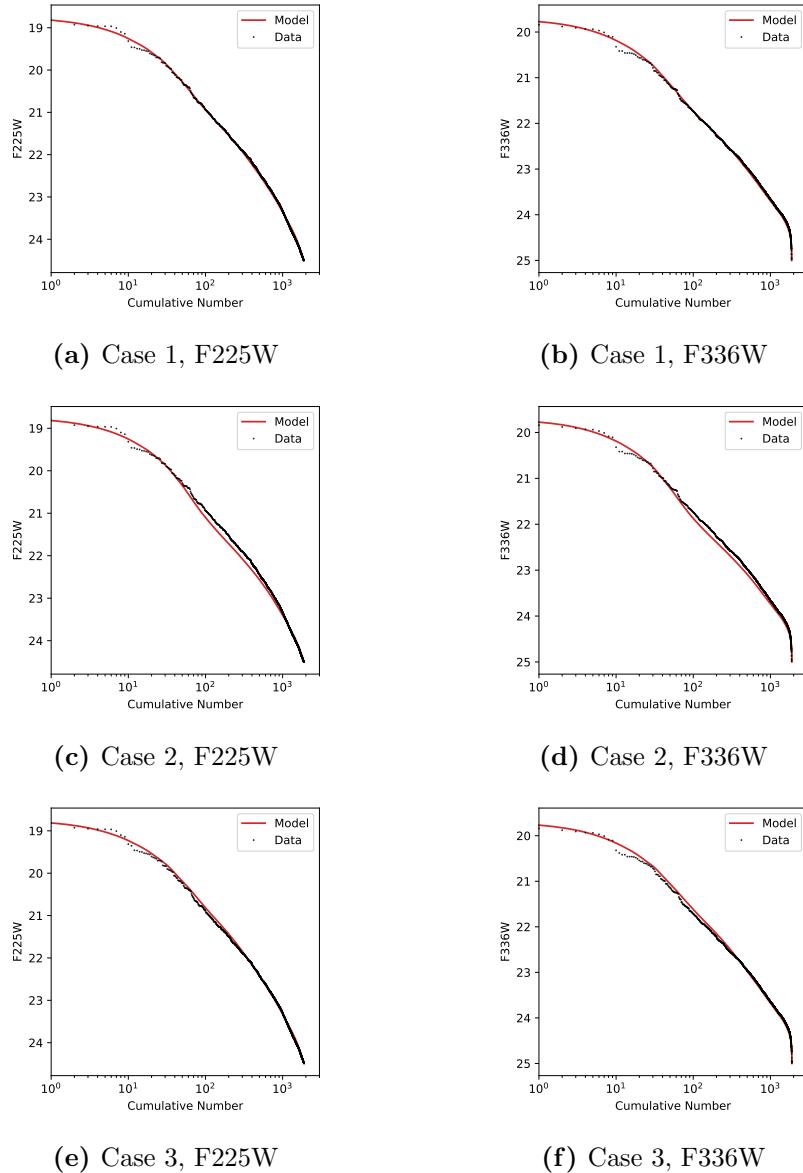
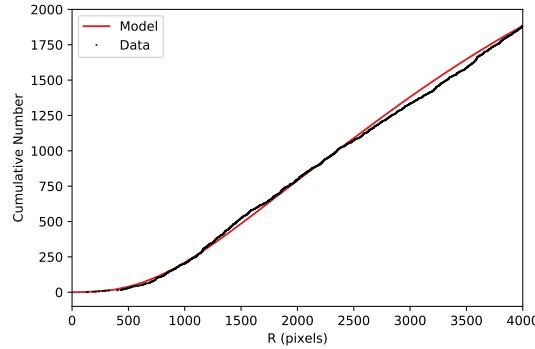
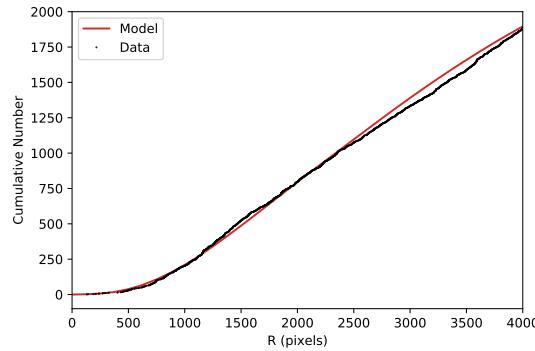


Figure B.10: WFC3, all cases: inverse cumulative luminosity function of optimal model (red curve) compared to the data (black points) for both F225W (left column) and F336W (right column). Each row shows a different case: case 1 (top), case 2 (middle), and case 3 (bottom).

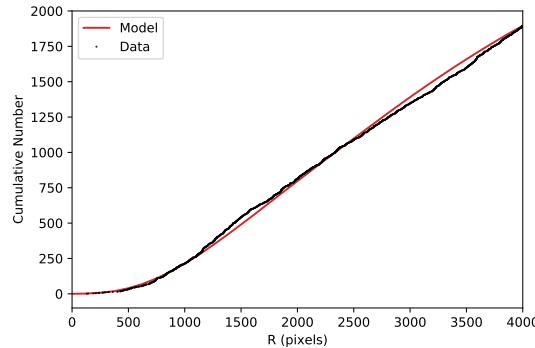
B.2. WFC3/UVIS Data Only



(a) Case 1



(b) Case 2



(c) Case 3

Figure B.11: WFC3, all cases: cumulative number distribution of radial distance (R) from cluster centre for optimal model (red curve) and WFC3/UVIS data (black points). Each row shows a different case: case 1 (top), case 2 (middle), and case 3 (bottom).

B.2. WFC3/UVIS Data Only

	\dot{N} (Myr $^{-1}$)	M_{WD} (M_{\odot})	$\log_{10} q_H$	m_a (meV)
Case 1	7.28	0.5388	-3.65	1.0
Case 2	8.06	0.5240	-5.15	0.0
Case 3	6.91	0.5388	-3.55	0.0

Table B.2: Parameter values of optimal model for the different WFC3 cases. In each case, this is the combination of parameter values that maximises the joint posterior distribution on the parameter grid.

	F225W	F336W	R
Case 1	0.0343	0.0124	0.0097
Case 2	0.0057	4×10^{-6}	0.0087
Case 3	0.0362	0.0132	0.0156

Table B.3: Results of KS tests for the different WFC3 cases. The p-values are reported for KS tests comparing the one-dimensional marginal cumulative probability distribution functions predicted by the optimal model for each case to the corresponding empirical distribution.

Table B.2 summarises the parameter values of the optimal model for each of the cases considered in the analysis of the WFC3 data alone. The optimal models for cases 1 and 3 (which both use uniform birthrate priors) correspond to very similar parameter values, both notably favouring large $\log_{10} q_H$ values. This is in contrast to case 2 (which used a Gaussian birthrate prior), which favours a much smaller $\log_{10} q_H$ value and larger \dot{N} . The case 2 result is pushed to larger \dot{N} values by the Gaussian birthrate prior used in this case, which seems to overestimate the birthrate. Forcing a birthrate value that is too large drives the fit to smaller $\log_{10} q_H$ values (for the data space considered in this work). Note that regardless of which priors are used, all of the cases favour m_a values near zero.

Table B.3 summarises the results of the one-sample KS tests for the WFC3 cases. Cases 1 and 3 have very similar p -values (when comparing the same data variable), and these p -values are reasonable for all of the data variables (i.e. F225W, F336W, and R). This is particularly the case for

the magnitudes (F225W and F336W); the p -value for R is somewhat better for case 3 than case 1, though both are reasonable. The KS test p -values are not a direct measure of goodness-of-fit, they simply provide a check as to whether it is reasonably likely that the data are drawn from the model distribution (and can be used to reject this null hypothesis in the case of a poor fit if the p -values are sufficiently small). The similarity of the KS test results for cases 1 and 3 is expected given the similarity of the best-fitting parameter values summarised in Table B.2. The smaller p -values for case 2 (which used a Gaussian instead of uniform birthrate prior) compared to case 1 (and case 3) are indicative of a worse fit. This is particularly notable in the p -values for the F336W magnitude (and to a lesser extent the F225W magnitude), with the case 2 value for F336W so small ($< 10^{-4}$) that we can reject the null hypothesis that the empirical F336W values are drawn from the corresponding distribution predicted by the best-fitting model of case 2. This suggests that the Gaussian WFC3 birthrate prior from Goldsbury *et al.* [215] is not an appropriate choice, at least for the data space used in this work. This may be due to the different definitions of the data space between this work and Goldsbury *et al.* [215]; axion emission can effect white dwarf cooling until later cooling times than the neutrino emission that was the focus of Goldsbury *et al.* [215], so we use a data space that extends to larger magnitudes (corresponding to older white dwarfs) than did Goldsbury *et al.* [215]. The older white dwarfs included in our data space could have a different birthrate than the younger white dwarfs.

B.3 ACS/WFC Data Only

B.3.1 Posterior Distributions

B.3.1.1 Case 1

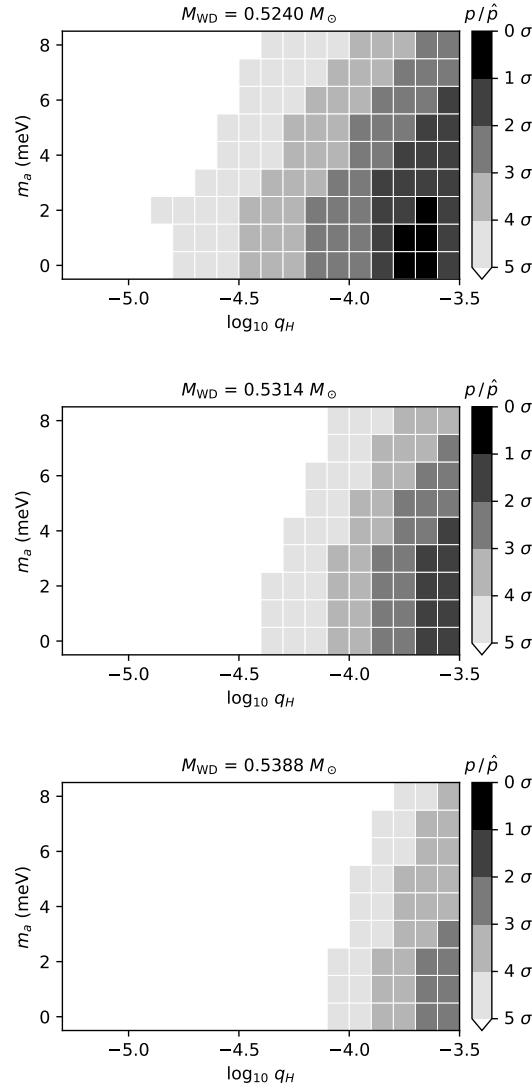


Figure B.12: ACS, case 1: joint posterior probability density distribution after marginalising over the birthrate.

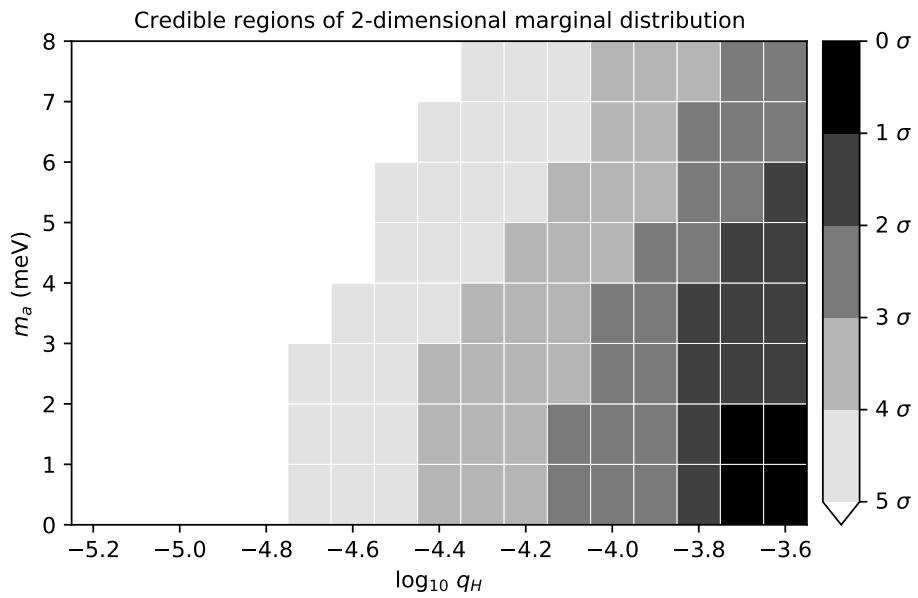


Figure B.13: ACS, case 1: two-dimensional joint credible regions of axion mass (m_a) and envelope thickness (q_H) after marginalising over the other parameters.

B.3. ACS/WFC Data Only

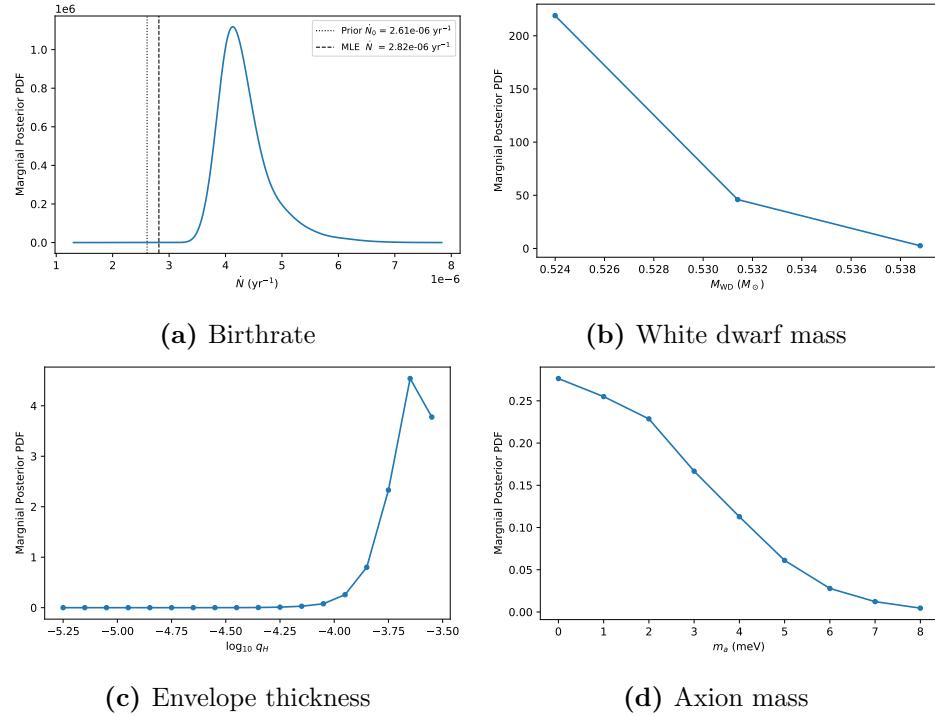


Figure B.14: ACS, case 1: one-dimensional posterior density distributions for each parameter after marginalising over all other model parameters.

B.3.1.2 Case 2

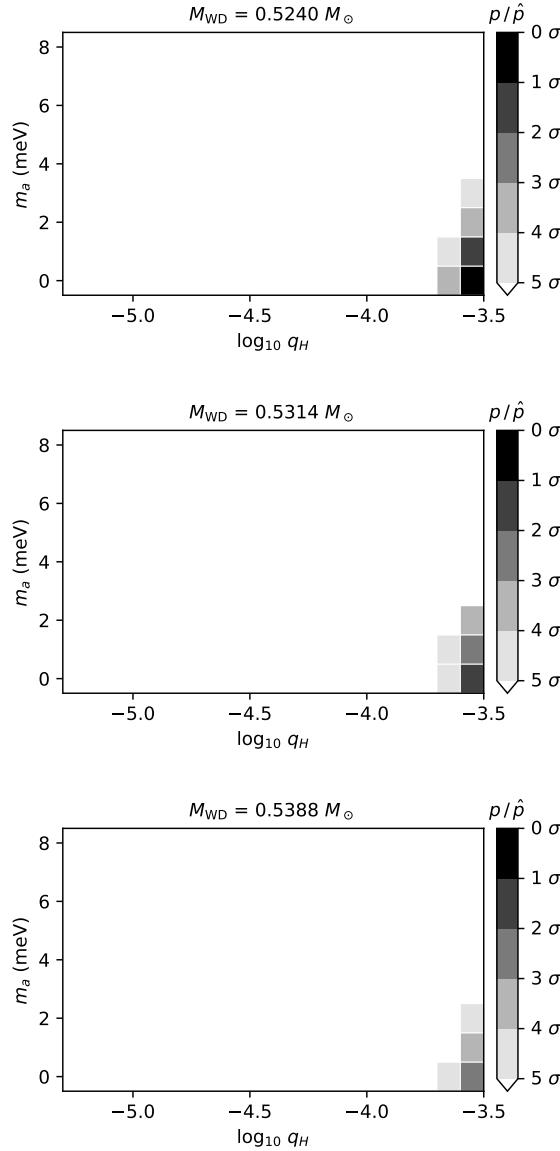


Figure B.15: ACS, case 2: joint posterior probability density distribution after marginalising over the birthrate.

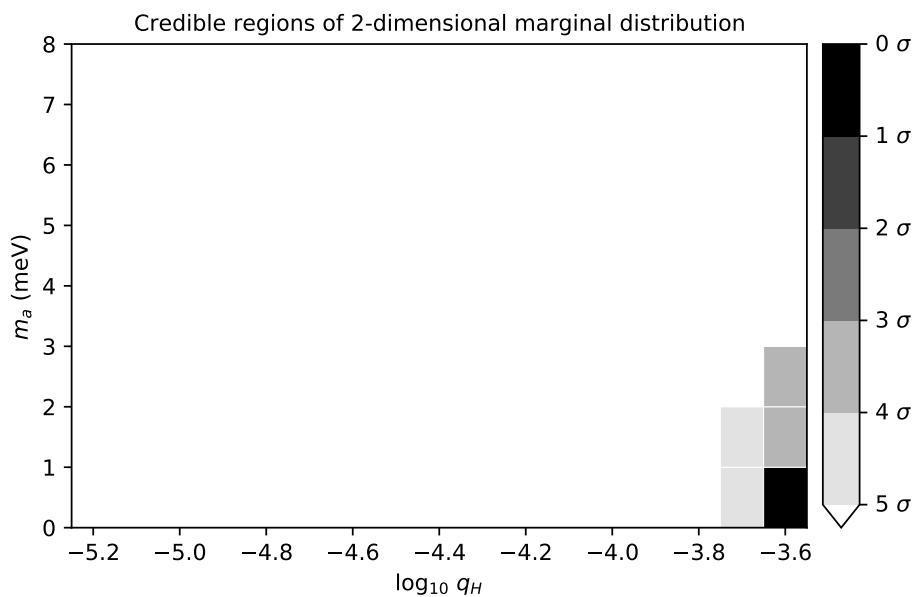


Figure B.16: ACS, case 2: two-dimensional joint credible regions of axion mass (m_a) and envelope thickness (q_H) after marginalising over the other parameters.

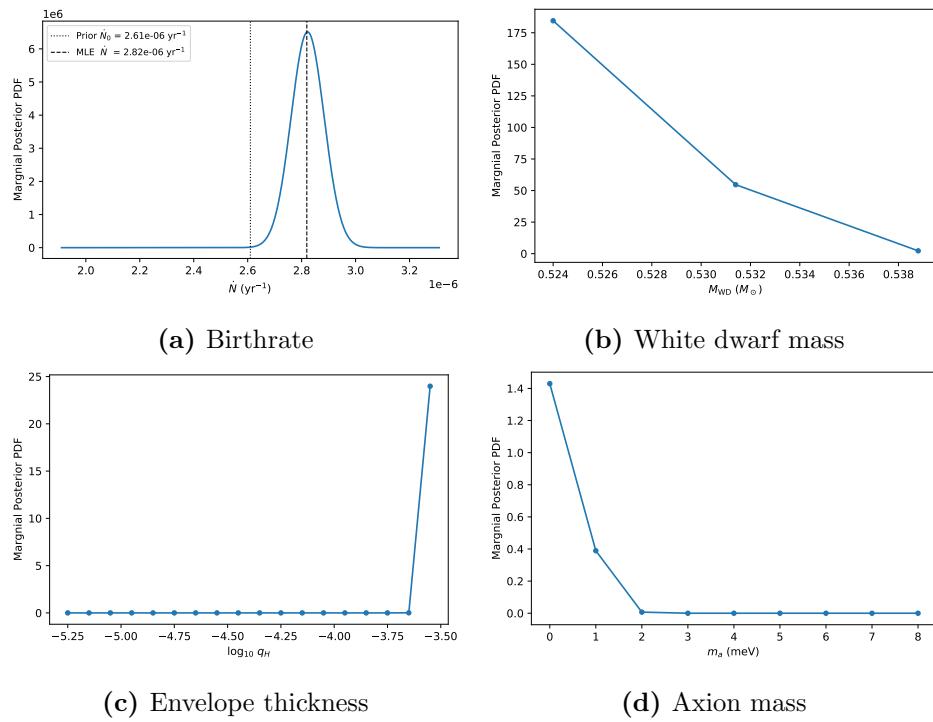


Figure B.17: ACS, case 2: one-dimensional posterior density distributions for each parameter after marginalising over all other model parameters.

B.3.1.3 Case 3

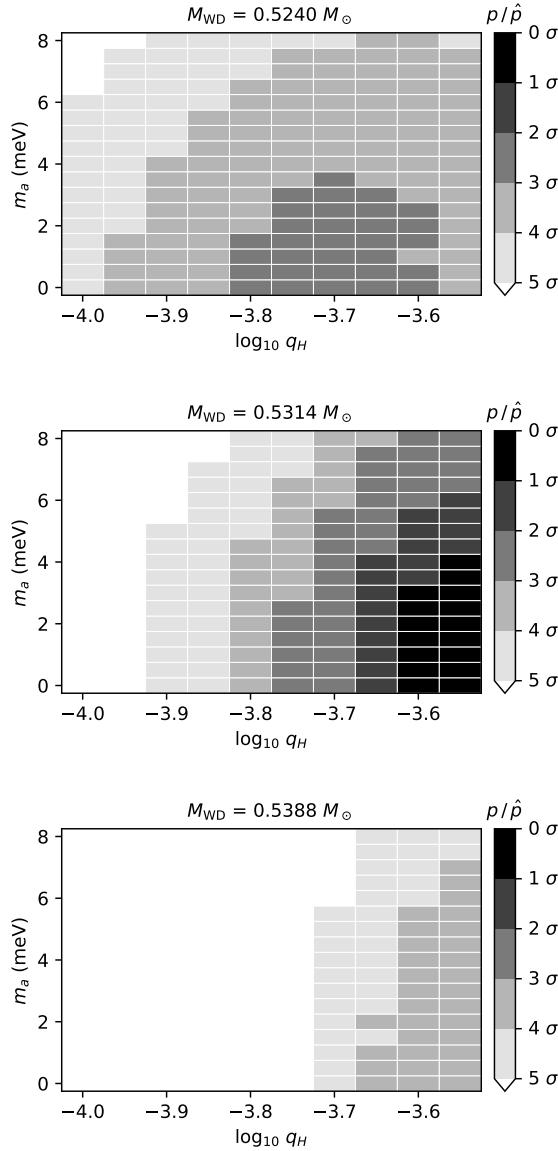


Figure B.18: ACS, case 3: joint posterior probability density distribution after marginalising over the birthrate.

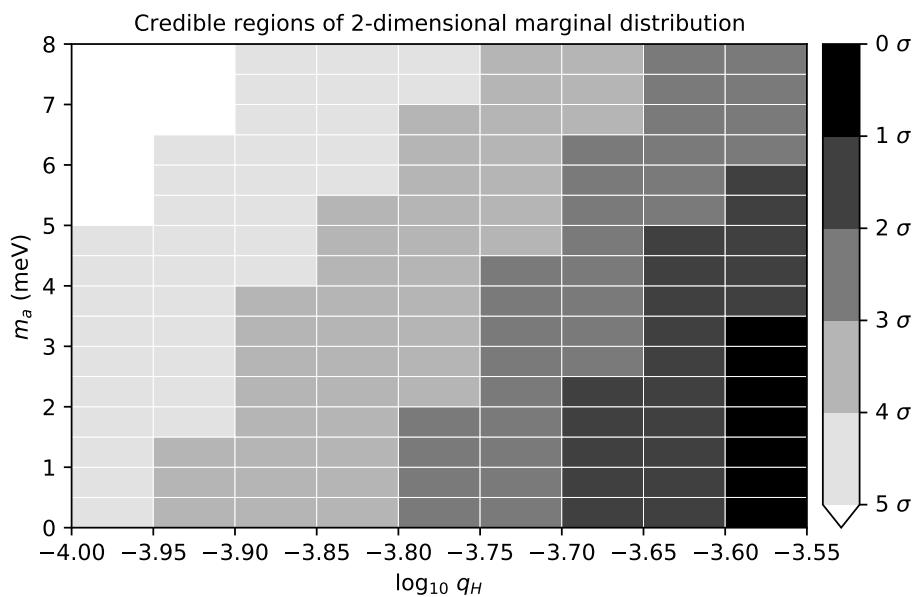


Figure B.19: ACS, case 3: two-dimensional joint credible regions of axion mass (m_a) and envelope thickness (q_H) after marginalising over the other parameters.

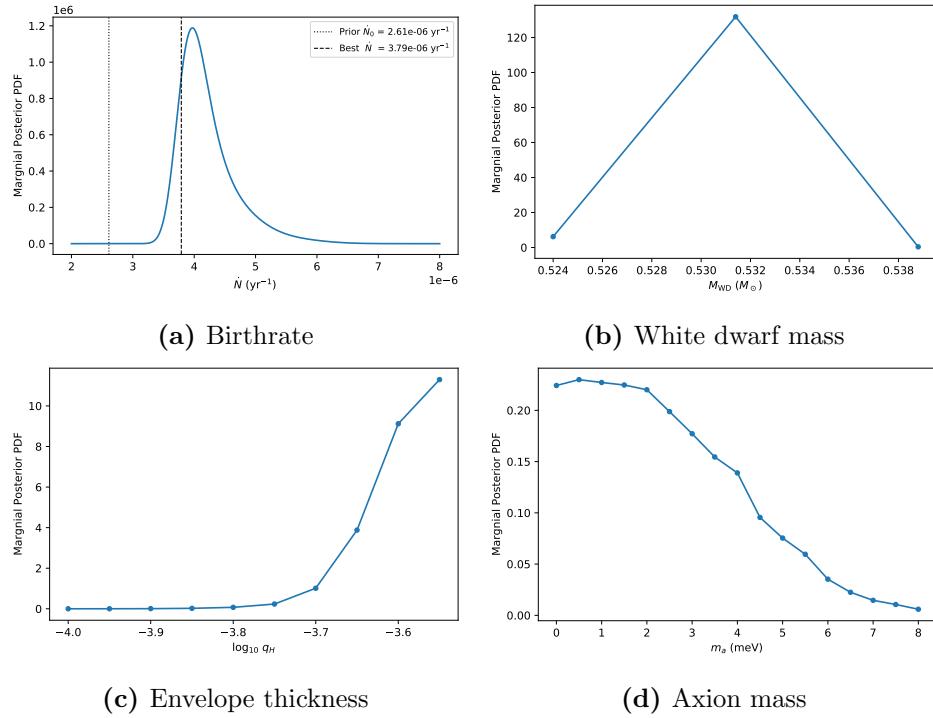


Figure B.20: ACS, case 3: one-dimensional posterior density distributions for each parameter after marginalising over all other model parameters.

B.3.2 Best-Fitting Models Comparison

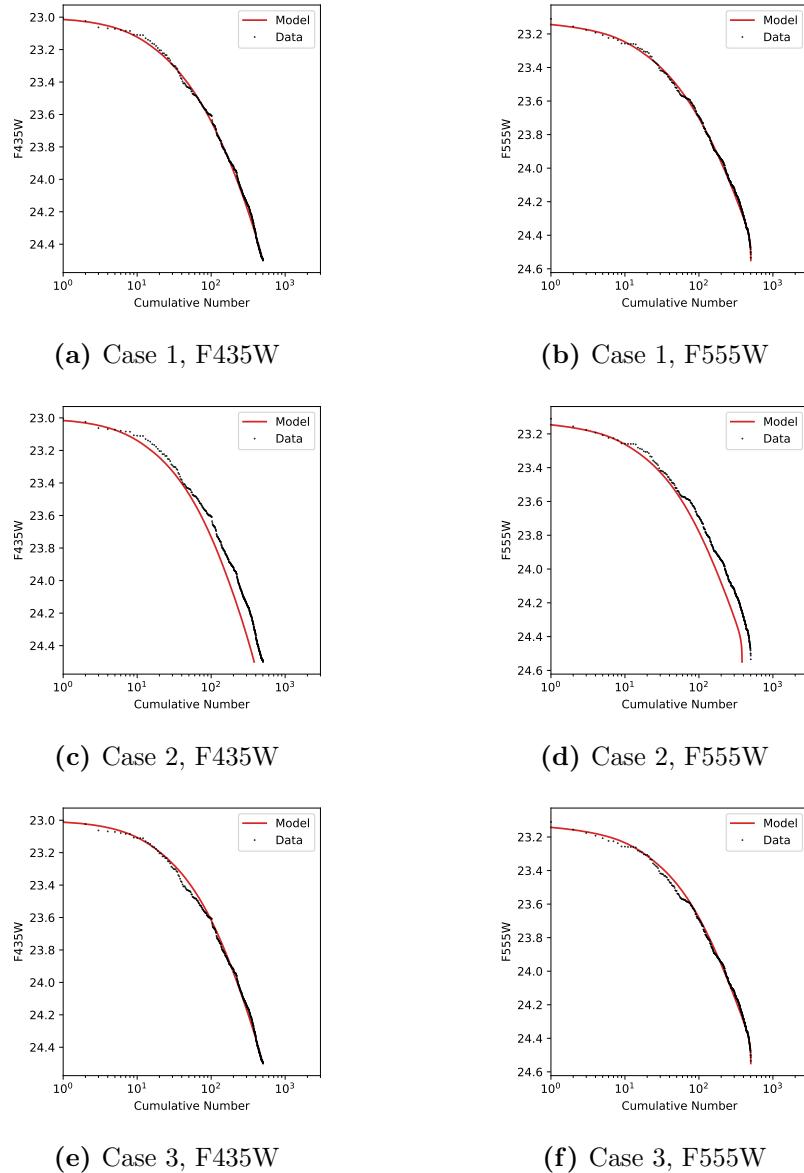


Figure B.21: ACS, all cases: inverse cumulative luminosity function of optimal model (red curve) compared to the data (black points) for both F435W (left column) and F555W (right column). Each row shows a different case: case 1 (top), case 2 (middle), and case 3 (bottom).

B.3. ACS/WFC Data Only

	\dot{N} (yrs $^{-1}$)	M_{WD} (M_{\odot})	$\log_{10} q_H$	m_a (meV)
Case 1	3.94×10^{-6}	0.5240	-3.65	0.0
Case 2	2.82×10^{-6}	0.5240	-3.55	0.0
Case 3	3.79×10^{-6}	0.5314	-3.55	1.5

Table B.4: Parameter values of optimal model for the different ACS cases. In each case, this is the combination of parameter values that maximises the joint posterior distribution on the parameter grid.

	F435W	F555W
Case 1	0.2098	0.2979
Case 2	0.3490	0.4693
Case 3	0.2660	0.2787

Table B.5: Results of KS tests for the different ACS cases. The p-values are reported for KS tests comparing the one-dimensional marginal cumulative probability distribution functions predicted by the optimal model for each case to the corresponding empirical distribution.

It can be seen from Fig. B.21 that the best-fitting models for both case 1 and case 3 fit the data well (for both magnitudes), as was also found for these cases when analysing the WFC3 data alone in Appendix B.2. This is expected because the optimal combination of parameter values is similar for case 1 and case 3. The benefit of case 3 (using the Chapter 6 priors for M_{WD} and $\log_{10} q_H$) over case 1 (using uniform priors for M_{WD} and $\log_{10} q_H$) is that case 3 provides tighter credible regions. To see this, compare Fig. B.13 and Fig. B.19 (and likewise Fig. B.2 and Fig. B.8 for the WFC3 only analysis) and note that Fig. B.19 (and Fig. B.8) are plotted over a smaller range of $\log_{10} q_H$ values (with the plot truncated at a lower limit of $\log_{10} q_H = -4.0$ instead of $\log_{10} q_H = -4.5$). The joint posterior and credible region plots for case 3 are shown over a smaller range of $\log_{10} q_H$ values than case 1 or case 2 for the sake of visualisation because the posterior distribution for case 3 drops below the 5- σ level at larger values of $\log_{10} q_H$ than the other cases; in other words, the posterior distribution (and credible regions) for case 1

(and case 2) are more extended in $\log_{10} q_H$ than for case 3.

It can also be seen from Fig. B.21 that the best-fitting model for case 2 (using the Goldsbury *et al.* [215] Gaussian birthrate prior) does not fit the data as well as the best-fitting models of case 1 or case 3 (using a uniform birthrate prior). The distributions predicted by the best-fitting model for case 2 in Fig. B.21 are shifted to the left (to smaller cumulative number values) compared to the corresponding empirical distributions, indicating that the birthrate is too small. Aside from the birthrate being too small (which just affects the overall normalisation of the distribution), the morphology of each of the case 2 model distributions appears reasonable. This is confirmed by the results of the KS tests (see Table B.5), which asses the morphology but not overall normalisation of the distributions. A reasonable morphology for the distributions for ACS case 2 should be expected because the best-fitting parameter values on the cooling model parameter grid are similar for all three cases in the ACS only analysis. In all three cases, the posterior probability is concentrated at the small m_a and large $\log_{10} q_H$ limits of the parameter grid (which are also physical limits, as discussed in Chapter 7). It can be seen from Fig. B.15 (and Fig. B.16) that imposing a prior that underestimates the value of the birthrate drives the most likely parameter values to smaller m_a and larger $\log_{10} q_H$ values, but as the posterior distributions for the uniform birthrate cases are already concentrated near the limiting values of m_a and $\log_{10} q_H$ (Fig. B.12 and Fig. B.18), the posterior distribution for case 2 simply ends up being very tightly concentrated in the low m_a , high $\log_{10} q_H$ corner of the plots while being optimised at a similar location in parameter space.

This is in contrast to what is seen for case 2 of the WFC3 only analysis, where the WFC3 birthrate prior from Goldsbury *et al.* [215] seems to overestimate the birthrate (instead of underestimating it like for the ACS birthrate). Imposing this likely overestimated value of the WFC3 birthrate prior drives the most likely parameter values to smaller $\log_{10} q_H$ values, with some degeneracy between $\log_{10} q_H$, M_{WD} , and m_a . This resulted in the $\log_{10} q_H$ value of the best-fitting model for WFC3 case 2 being much lower than for the other WFC3 (and ACS) cases, so the cumulative number

B.3. ACS/WFC Data Only

distribution functions for WFC3 case 2 had a notably different morphology than the other cases in addition to a notably different birthrate value (see Fig. B.10). This different morphology is why the p -values of the KS tests for WFC3 case 2 were much smaller than the values for the other cases (especially for the magnitude variables).

B.4 Comparison of WFC3, ACS, and Combined Analyses

Tables B.6 to B.8 summarise and compare the results of the combined analysis of the WFC3 and ACS data from Chapter 7 to the results of the individual analyses of the WFC3 data alone and the ACS data alone for the same choice of priors as were used in the combined analysis. This choice of priors consisted of uniform priors for the white dwarf birthrates and the prior from the Chapter 6 results for M_{WD} and $\log_{10} q_H$, which corresponds to case 3 for both the WFC3 only analysis in Appendix B.2 and the ACS only analysis in Appendix B.3.

The 95% credible regions are reported directly in Table B.6. These credible regions were calculated from the one-dimensional marginal posterior distribution of each parameter. In Table B.7, these same credible regions are given as errors on the best-fitting combination of parameters on the discrete parameter grid for M_{WD} , $\log_{10} q_H$, and m_a . These best-fitting values are the combination of values that optimise the full joint posterior distribution.

These results show that analysing the WFC3 and ACS data individually gives similar best-fitting values. The combined analyses likewise gives a similar result to analysing either of the data sets individually, but with tighter credible regions. Compared to the WFC3 only credible regions, the combined analysis gives only a minor improvement. The tightening of the credible region is much more dramatic when comparing the result of the ACS only analysis to the combined analysis. This is to be expected because there are less white dwarfs in the ACS data space than the WFC3 data space (for reasons discussed in Chapter 7), so the WFC3 data has more weight in the combined analysis. The KS test results to check the goodness-of-fit are given in Table B.8 and have reasonable values in all cases.

B.4. Comparison of WFC3, ACS, and Combined Analyses

	WFC3	ACS	Combined
$\dot{N}_{\text{WFC3}} (\text{Myr}^{-1})$	6.66 – 7.60	–	6.68 – 7.73
$\dot{N}_{\text{ACS}} (\text{Myr}^{-1})$	–	3.48 – 5.22	3.49 – 4.35
$M_{\text{WD}} (M_{\odot})$	0.5327 – 0.5388	0.5254 – 0.5371	0.5282 – 0.5388
$\log_{10} q_H$	–3.65 – 3.55	–3.69 – 3.55	–3.67 – 3.55
$m_a (\text{meV})$	0.00 – 2.93	0.00 – 5.64	0.00 – 2.85
$g_{aee} / 10^{-13}$	0.00 – 0.83	0.00 – 1.60	0.00 – 0.81

Table B.6: Comparison of 95% credible regions given by the analysis of only WFC3, only ACS, and combined WFC3 and ACS data. The same priors and parameter grid were used in all three cases (where applicable). These priors correspond to case 3 for the WFC3 only and ACS only analyses and correspond to the main results presented in Chapter 7 for the combined analysis.

	WFC3	ACS	Comb.
$\dot{N}_{\text{WFC3}} (\text{Myr}^{-1})$	$6.91^{+0.69}_{-0.25}$	–	$6.91^{+0.82}_{-0.23}$
$\dot{N}_{\text{ACS}} (\text{Myr}^{-1})$	–	$3.79^{+1.43}_{-0.31}$	$3.73^{+0.62}_{-0.24}$
$M_{\text{WD}} (M_{\odot})$	$0.5388^{+0.0000}_{-0.0061}$	$0.5314^{+0.0057}_{-0.0060}$	$0.5388^{+0.0000}_{-0.0106}$
$\log_{10} q_H$	$-3.55^{+0.00}_{-0.10}$	$-3.55^{+0.00}_{-0.14}$	$-3.55^{+0.00}_{-0.12}$
$m_a (\text{meV})$	$0.00^{+2.93}_{-0.00}$	$1.50^{+4.14}_{-1.50}$	$0.00^{+2.85}_{-0.00}$
$g_{aee} / 10^{-13}$	$0.00^{+0.83}_{-0.00}$	$0.42^{+1.17}_{-0.42}$	$0.00^{+0.81}_{-0.00}$

Table B.7: Comparison of best-fitting model parameters for the WFC3 only (case 3), ACS only (case 3), and combined (Chapter 7) analyses. The 95% credible regions from Table B.6 are reported as errors on the best-fitting combination of parameter values on the parameter grid.

B.4. Comparison of WFC3, ACS, and Combined Analyses

	WFC3	ACS	Comb.
R	0.0156	—	0.0156
F225W	0.0362	—	0.0362
F336W	0.0132	—	0.0132
F435W	—	0.2660	0.3527
F555W	—	0.2787	0.2739

Table B.8: Comparison of KS test results for WFC3 only (case 3), ACS only (case 3), and combined analyses.