

Library Declaration and Deposit Agreement

1. STUDENT DETAILS

Matthew John Green
1105284

2. THESIS DEPOSIT

- 2.1 I understand that under my registration at the University, I am required to deposit my thesis with the University in BOTH hard copy and in digital format. The digital version should normally be saved as a single pdf file.
- 2.2 The hard copy will be housed in the University Library. The digital version will be deposited in the University's Institutional Repository (WRAP). Unless otherwise indicated (see 2.3 below) this will be made openly accessible on the Internet and will be supplied to the British Library to be made available online via its Electronic Theses Online Service (EThOS) service. [At present, theses submitted for a Master's degree by Research (MA, MSc, LLM, MS or MMedSci) are not being deposited in WRAP and not being made available via EthOS. This may change in future.]
- 2.3 In exceptional circumstances, the Chair of the Board of Graduate Studies may grant permission for an embargo to be placed on public access to the hard copy thesis for a limited period. It is also possible to apply separately for an embargo on the digital version. (Further information is available in the Guide to Examinations for Higher Degrees by Research.)
- 2.4 (a) Hard Copy I hereby deposit a hard copy of my thesis in the University Library to be made publicly available to readers immediately.
I agree that my thesis may be photocopied.
- (b) Digital Copy I hereby deposit a digital copy of my thesis to be held in WRAP and made available via EThOS.
My thesis can be made publicly available online.

3. GRANTING OF NON-EXCLUSIVE RIGHTS

Whether I deposit my Work personally or through an assistant or other agent, I agree to the following: Rights granted to the University of Warwick and the British Library and the user of the thesis through this agreement are non-exclusive. I retain all rights in the thesis in its present version or future versions. I agree that the institutional repository administrators and the British Library or their agents may, without changing content, digitise and migrate the thesis to any medium or format for the purpose of future preservation and accessibility.

4. DECLARATIONS

(a) I DECLARE THAT:

- I am the author and owner of the copyright in the thesis and/or I have the authority of the authors and owners of the copyright in the thesis to make this agreement. Reproduction of any part of this thesis for teaching or in academic or other forms of publication is subject to the normal limitations on the use of copyrighted materials and to the proper and full acknowledgement of its source.
- The digital version of the thesis I am supplying is the same version as the final, hardbound copy submitted in completion of my degree, once any minor corrections have been completed.
- I have exercised reasonable care to ensure that the thesis is original, and does not to the best of my knowledge break any UK law or other Intellectual Property Right, or contain any confidential material.
- I understand that, through the medium of the Internet, files will be available to automated agents, and may be searched and copied by, for example, text mining and plagiarism detection software.

(b) IF I HAVE AGREED (in Section 2 above) TO MAKE MY THESIS PUBLICLY AVAILABLE DIGITALLY, I ALSO DECLARE THAT:

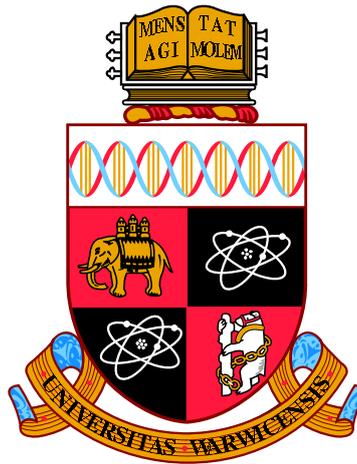
- I grant the University of Warwick and the British Library a licence to make available on the Internet the thesis in digitised format through the Institutional Repository and through the British Library via the EThOS service.
- If my thesis does include any substantial subsidiary material owned by third-party copyright holders, I have sought and obtained permission to include it in any version of my thesis available in digital format and that this permission encompasses the rights that I have granted to the University of Warwick and to the British Library.

5. LEGAL INFRINGEMENTS

I understand that neither the University of Warwick nor the British Library have any obligation to take legal action on behalf of myself, or other rights holders, in the event of infringement of intellectual property rights, breach of contract or of any other right, in the thesis.

Please sign this agreement and return it to the Graduate School Office when you submit your thesis.

Student's signature:Date:



The Evolution of AM CVn Binary Systems

by

Matthew John Green

Thesis

Submitted to the University of Warwick

for the degree of

Doctor of Philosophy

Department of Physics

August 2019

THE UNIVERSITY OF
WARWICK

Contents

Acknowledgments	iv
Declarations	vi
Abstract	viii
Abbreviations	ix
Chapter 1 Introduction	1
1.1 White dwarfs and single star evolution	2
1.2 Binary evolution and the processes involved	5
1.2.1 Common envelope evolution	5
1.2.2 Angular momentum loss	6
1.2.3 Mass transfer	7
1.3 AM CVn type binaries	12
1.4 The evolution of AM CVn binaries	15
1.5 Orbital Periods and AM CVn Classification	21
1.6 Overview	25
Chapter 2 Methodology	31
2.1 High-Speed Cameras	31
2.2 Photometric Reduction	35
2.2.1 Bias and Dark Subtraction	37
2.2.2 Flat Field Correction	38
2.2.3 Aperture Photometry	38
2.2.4 Sources of Noise in Photometric Data	40
2.3 Spectroscopic Reduction	41
2.3.1 Bias Subtraction	42
2.3.2 Flat Correction	42

2.3.3	Spectrum Tracing	43
2.3.4	Sky Subtraction	43
2.3.5	Spectrum Extraction	44
2.3.6	Wavelength and Flux Calibration	44
2.4	Light Curve Modelling	45
2.5	Doppler Maps	46
Chapter 3 An Eclipsing AM CVn that Challenges Formation Models		51
3.1	Introduction	51
3.2	Observations	53
3.3	Photometry	54
3.4	Light Curve Modelling	60
3.4.1	MCMC Modelling	61
3.4.2	Bootstrapping	63
3.4.3	Contact Phase Measurements	64
3.5	Modelling Results and Discussion	66
3.5.1	Ephemeris	72
3.5.2	White Dwarf Luminosity and Colour	73
3.5.3	Mass Ratio and Donor Mass	77
3.5.4	Comparison to Models	78
3.5.5	Implications For Gravitational Wave Emission	82
3.6	Conclusions	83
Chapter 4 Line Emission From Near the Central White Dwarf of Gaia14aae		84
4.1	Introduction	84
4.2	Observations	85
4.3	Results	88
4.3.1	Average Spectrum	88
4.3.2	In-Eclipse Spectrum and White Dwarf Spectrum	97
4.3.3	Trailed Spectra and Doppler Maps	103
4.4	Discussion	105
4.4.1	The Origin of the Central Spike	105
4.4.2	The Positions of the Bright Spots	106
4.4.3	Implications for the Formation of Gaia14aae	109
4.5	Conclusions	111

Chapter 5	A 15.7 Minute AM CVn Binary in a <i>K2</i> Field	113
5.1	Introduction	113
5.2	Observations	115
5.2.1	<i>K2</i> Photometry	117
5.2.2	McDonald/ProEM Photometry	117
5.2.3	SOAR/Goodman Spectroscopy	117
5.2.4	SOAR/Goodman Photometry	118
5.2.5	NTT/ULTRACAM Photometry	118
5.3	Analysis	120
5.3.1	Spectroscopy	120
5.3.2	<i>K2</i> Photometry	123
5.3.3	High-cadence photometry	126
5.4	Discussion	129
5.4.1	Classification as an AM CVn star	129
5.4.2	The Nature of the Photometric Periods	130
5.4.3	Mass Ratio	132
5.4.4	Space Density	135
5.5	Conclusions	135
Chapter 6	VLT Spectroscopy of Five Accreting Binaries Below the Period Minimum	137
6.1	Introduction	137
6.2	Observations	139
6.3	Results	140
6.3.1	SDSS J1505+0659	140
6.3.2	ASASSN-14ei	151
6.3.3	CRTS J2333–1557	157
6.3.4	CRTS J1028–0819	158
6.3.5	V418 Ser	160
6.4	Discussion	161
6.4.1	Evolutionary Natures	161
6.4.2	Prospects for Future Study	163
6.5	Conclusions	164
Chapter 7	Conclusions and Outlook	165
7.1	Implications for AM CVn Formation	167
7.2	Future Perspective	168

Acknowledgments

This thesis is written in first person singular, as is tradition, but it is the work of much more than one person. The papers collected here have several dozen authors between them, and there are dozens more people who have helped me, both scientifically and personally, over the last 3.5 years. I cannot hope to name all of them, and there will almost certainly be somebody important that I miss off. Still, I will try my best.

Firstly, I have to thank my two supervisors, Tom Marsh and Danny Steeghs. They taught me to be an astronomer; without them this thesis would be a sheaf of empty pages. I particularly need to thank Tom for being happy to patiently explain things several times until I understood. For similar reasons, I owe debts to Elmé Breedt and Mark Hollands, both of whom have always been happy to help whenever I've found myself stuck on something. I also want to thank Richard Ashley, who took me observing for the first time and taught me to love being at a telescope.

A PhD is more than just the work that you do, and I would have had a much harder time over the last few years if the astronomy group at Warwick was not such a welcoming and inclusive place. It is this way because of everybody who works there, and I love you all. In particular, I owe personal gratitude to George King, Ares Osborn, Matt Hoskin, Chris Manser, Elena Cukanovaite, James Jackman, and Tim Cunningham. Outside of the office, I also have to thank Jacques Dark, Jamie Bayne, and James Powell, for helping to keep me sane and happy. And, of course, I can go no further without mentioning Lewis, Emily and Isobel Bourne, who still put up with me even after 20 years, and Garion Swann, who hasn't done yet but I'm sure will do.

Lastly, I have to thank my family. My mum, for drawing diagrams to explain how eyes work, and my dad, for teaching me short multiplication – both of them, in short, for teaching me all my life to be curious, to love the world, and that there’s no shame in asking questions. And my brother James, for always being the person I can talk to about anything, no matter how personal or how obscure and nerdy it might be.

More formally, I need to acknowledge the funding that made this work possible. My studentship was funded by STFC via grant ST/N504506/1. Tom Marsh and Danny Steeghs were funded by STFC grants ST/L000733/1 and ST/P000495/1. ULTRACAM and ULTRASPEC, used heavily in this thesis, are funded by STFC grant ST/J001589.

Data in this thesis were obtained from numerous telescopes, including: the 4.5 m William Herschel Telescope (WHT), operated by the Isaac Newton Group (ING) at the Roque de los Muchachos Observatory, La Palma; the 2.4 m Thai National Telescope (TNT), operated by the National Astronomy Research Institute of Thailand (NARIT) at the Thai National Observatory on Doi Inthanon; the 200-inch Hale Telescope at Palomar observatory, operated by the California Institute of Technology; the 2.1 m Otto Struve Telescope at the McDonald Observatory, operated by the University of Texas at Austin; the 4.1 m Southern Astrophysical Research (SOAR) telescope on Cerro Pachón, operated by the Ministério da Ciência, Tecnologia, Inovações e Comunicações (MCTIC) do Brasil, the U.S. National Optical Astronomy Observatory (NOAO), the University of North Carolina at Chapel Hill (UNC), and Michigan State University (MSU); the 3.5 m New Technology Telescope (NTT) at La Silla observatory, operated by the European Southern Observatory (ESO); the 10.4 m Gran Telescopio Canarias (GTC) at the Roque de los Muchachos Observatory, La Palma; and the 8.2 m UT1 and UT2 telescopes of the Very Large Telescope (VLT) at Paranal observatory, operated by the European Southern Observatory (ESO). This thesis also contains data from the *K2* mission, which is funded by the NASA Science Mission directorate.

Declarations

I submit this thesis to the University of Warwick graduate school for the degree of Doctor of Philosophy. This thesis was written by me and has not been submitted to any other institution.

This thesis includes material from published papers written by me, including:

- Green et al. (2018b), MNRAS, 476, 1663, “*High-Speed Photometry of Gaia14aae: An Eclipsing AM CVn That Challenges Formation Models*”. Material from this paper is used in Chapter 3.
- Green et al. (2018a), MNRAS, 477, 5646, “*A 15.7-Minute AM CVn Binary Discovered in K2*”. Material from this paper is used in Chapter 5.
- Green et al. (2019), MNRAS, 485, 1947, “*Phase-Resolved Spectroscopy of Gaia14aae: Line Emission From Near the White Dwarf Surface*”. Material from this paper is used in Chapter 4.

Material from these papers also contributed to the Introduction and Methodology chapters.

The work presented in this thesis was performed by me, with the following exceptions:

- For Chapter 3, observations were carried out by members of the ULTRACAM and CHIMERA teams, under proposals led by Tom Marsh, Thomas Kupfer, Chris Copperwheat, and Steven Bloemen. Subsequent analysis was performed by me.

- For Chapter 4, observations were carried out by Elmé Breedt and Jan van Roestel, under proposals led by Pablo Rodríguez-Gil and Thomas Kupfer. Reductions were performed by Elmé Breedt. DB white dwarf model spectra were provided by Elena Cukanovaite. Subsequent analysis was performed by me.
- For Chapter 5, observations were carried out by members of the ULTRACAM team and various other observers, under proposals led by JJ Hermes and me. Some aspects of the photometric period analysis were performed by Keaton Bell. Subsequent analysis was performed by me.
- For Chapter 6, observations were carried out under an ESO service mode proposal (095.D-0888) led by Phil Carter. XHOOTER reductions were performed by Nicola Gentile Fusilo. Subsequent analysis, including FORS2 reduction, was performed by me.

Abstract

AMCVn-type binary systems are a class of ultracompact, accreting binary system. Each system consists of a white dwarf accreting helium-dominated material from a compact donor star, which may be degenerate or semi-degenerate. Due to the position of these binary systems within late binary evolution, they provide an opportunity to constrain poorly understood aspects of binary evolution such as common envelope evolution and the merger rate of double white dwarf binaries (the favoured progenitors of type Ia supernovae). Despite the growing number of known AMCVn systems, details of the prior evolution of these systems remain poorly constrained, with several competing models. The objective of this thesis is to use observations of AMCVn binaries to explore their prior evolution.

I present an in-depth study of Gaia14aae, the only known fully-eclipsing AMCVn system. High-speed photometry allows me to model the system and deduce its component stellar masses and radii. Phase-resolved spectroscopy provides an independent test of the photometric results, and allows me to probe the origin of the line emission seen in many AMCVn systems. The properties of Gaia14aae make it clear that the donor is not degenerate, and suggest an unusual evolutionary origin in which the system is descended from a short-period cataclysmic variable.

I present the discovery of SDSS J1351-0643, a 16 minute orbital period binary which was observed by the *K2* project. Several periods with physical origins are identified, allowing me to estimate the mass ratio of the component stars. From the literature, I compile all known mass ratio measurements and compare these with evolutionary models. The results suggest that AMCVn donor stars across the population are more inflated than has been previously expected, and that few, if any, donor stars are fully degenerate. Acknowledging the need for a greater number of well-characterised systems in order to fully explore this suggestion, I present spectroscopy of five AMCVn binaries and short-period cataclysmic variables. Four of these systems show potential to be well characterised with further observations.

The results in this thesis, in line with other recent results, suggest that the importance of the helium star donor channel and the evolved CV channel of AMCVn formation have historically been underestimated. The importance of the white dwarf donor channel has either been overestimated, or mechanisms that heat and inflate the donor have been overlooked. Future studies to find and characterise new AMCVn binaries will help to further explore these questions.

Abbreviations

4MOST	4-metre Multi-Object Spectroscopic Telescope
ADU	Analogue Digital Units
AGB	Asymptotic Giant Branch
AM CVn	AM Canum Venaticorum (type binary)
ASAS-SN	All-Sky Automated Search for SuperNovae
CCD	Charge Coupled Device
CRTS	Catalina Real-time Transient Survey
CV	Cataclysmic Variable
DESI	Dark Energy Survey Instrument
DWD	Double White Dwarf
EMCCD	Electron Multiplying Charge Coupled Device
ESO	European Southern Observatory
FWHM	Full Width at Half Maximum
GOTO	Gravitational wave Optical Transient Observer
GTC	Gran Telescopio Canarias
GW	Gravitational Wave

HWHM	Half Width at Half Maximum
INT	Isaac Newton Telescope
IR	Infrared
J-PAS	Javalambre Physics of the Accelerating Universe Astrophysical Survey
LISA	Laser Interferometer Space Antenna
LSST	Large Synoptic Survey Telescope
MCMC	Markov Chain Monte Carlo
NTT	New Technology Telescope
PCEB	Post Common Envelope Binary
PSF	Point Spread Function
PTF	Palomar Transient Factory
RGB	Red Giant Branch
RV	Radial Velocity
S/N	Signal-to-Noise Ratio
SDSS	Sloan Digital Sky Survey
SN Ia	Type Ia Supernova
TNT	Thai National Telescope
UV	Ultraviolet
VLT	Very Large Telescope
WEAVE	WHT Enhanced Area Velocity Explorer
WHT	William Herschel Telescope
ZTF	Zwicky Transient Facility

Chapter 1

Introduction

‘Astronomy is not terribly important,’ he said. He fell silent for a moment, admiring Arcturus. Then he said, ‘Although it is one of the more important things we do as a species.’ He did not see any contradiction there.

— Jim Gunn (as quoted by Richard Preston in *First Light*, 1995)

In approximately 1200 BCE, Ancient Egyptian astronomers knew of a star that seemed to disappear once every 2.85 days. The star became associated with misfortune, and superstitions about ‘lucky’ and ‘unlucky’ days can be found in Ancient Egyptian calendars on that 2.85 day period (Jetsu et al., 2013). This is perhaps the earliest evidence of humans studying stellar variability. The connection of the star with misfortune continued through millenia: the star’s English name, Algol, is derived from its Arabic name, *Ra’s Al-Ghūl*, which translates roughly as ‘the head of the ghoul’. The variability of Algol was explained almost 3000 years after those early observations: Algol is an unresolved system of multiple stars, and the dips in brightness occur when one star passes in front of another, blocking its light (Goodricke, 1783; Aldis, 1870; Pickering, 1881). The component stars of Algol were eventually resolved by Baron et al. (2012).

Binary systems, and indeed eclipsing systems like Algol, are far from unusual. More than half of the stars in the sky are part of multiple systems, either binaries or higher orders. Some of these are ‘resolved’ (both component stars are seen), but many are unresolved (appearing as single stars, but identifiable as binaries by either photometric variability or spectroscopic signatures).

The interactions that occur between close binaries affect their evolution in

a way that is important to consider if we are to understand stellar populations on a large scale. The ‘blue stragglers’ that are seen in clusters are products of binary evolution (Leonard, 1989; Knigge et al., 2009), and the spectral energy distributions (SEDs) of galaxies agree measurably better with models that include binary evolution (Eldridge et al., 2017; Stanway and Eldridge, 2018). Many classes of objects and phenomena result from binary evolution, including novae and dwarf novae (Warner, 1995), R CrB stars (Webbink, 1984; Saio and Jeffery, 2002), hot subdwarfs (sdO/sdB stars, Paczynski, 1976; Heber, 2016), and dwarf carbon stars (Green, 2013; Margon et al., 2018). Perhaps foremost among the products of binary evolution are Type Ia supernovae, thermonuclear explosions resulting from white dwarfs that have grown in mass until they reach the Chandrasekhar limit. Despite the importance of these phenomena to cosmologists as standard candles to allow distance measurements, the mechanism by which this mass limit is reached is still a matter of debate (Maoz et al., 2014).

This thesis collects work on a class of interacting binary known as AM Canum Venaticorum-type binaries, or AM CVns. AM CVn systems are helium-accreting, ultracompact binaries. Each system consists of a white dwarf accreting hydrogen-deficient matter from a donor star, which may be degenerate or semi-degenerate. As such, AM CVn binaries are a subset of cataclysmic variables (CVs), distinguished from the majority of CVs by the helium-dominated nature of the accreted material. AM CVn systems occur late in a binary’s evolution, and are only possible in a small minority of systems. The details of how these systems form are not clear. However, it is certain that they have evolved through many previous stages of evolution, including some of the least understood mechanisms in binary evolution (Marsh et al., 2004; Ivanova et al., 2013). By exploring and constraining the AM CVn population, we hope to come to a better understanding of the late-stage evolution of binary stars.

1.1 White dwarfs and single star evolution

Around 95% of all single stars will end their lives as white dwarfs – this includes all stars with masses $\lesssim 8 M_{\odot}$ (Althaus et al., 2010). A main sequence star such as the Sun will evolve into a white dwarf via several stages of giant evolution. Full descriptions of the process can be found in numerous astronomical textbooks (eg. Zeilik and Gregory, 1988; Carroll and Ostlie, 2013), but here I overview the relevant sections.

A main sequence star is defined as a star undergoing hydrogen fusion in

its core. Hydrogen is fused to helium via the proton-proton chain process and, in hotter stars, the carbon-nitrogen-oxygen (CNO) cycle. A star the mass of the Sun will spend $\approx 80\%$ of its life on the main sequence, or $\sim 10^{10}$ yrs. Once the hydrogen in the core is consumed, fusion in the core ceases. However, fusion continues in a shell around the core, where there is still fresh hydrogen to burn. This shell burning transfers heat to the surrounding envelope and causes the star to expand significantly; in the case of the Sun, the radius will increase by more than a factor of 100. Note that, although heat is transferred to the envelope, the result of expansion is that the temperature of the envelope decreases, causing the colour of the star to redden. Meanwhile, the helium core contracts and increases in density until its electrons become a degenerate gas, supported by electron degeneracy pressure. This phase of the star's evolution is known as the red giant branch (RGB).

For most stars, a second phase of core nuclear burning follows. Once the core temperature reaches $\approx 10^8$ K, helium burning by the triple- α process begins in the core, producing predominantly carbon and oxygen. Due to the speed with which ignition spreads throughout the degenerate core, this is known as the 'helium flash'. After some time the core helium is also exhausted. Fusion again moves to a shell surrounding the core, the star again expands, and the core again becomes degenerate. This phase of the star's evolution is the asymptotic giant branch (AGB). The rate of the triple- α is very sensitive to temperature, amplifying small fluctuations in stellar temperature and radius into pulsations known as 'thermal pulses'. A significant portion of the star's atmosphere is ejected by these thermal pulses. The core remains, as a small, hot, blue object surrounded by a nebula of the ejected material. As the core contracts, it evolves into a white dwarf.

The majority of white dwarfs are expected to consist of a carbon and oxygen core, surrounded by thin layers of helium and (in some cases) hydrogen (see Fig. 1.1). In some white dwarfs formed from the highest-mass progenitors, the core can instead be oxygen and neon. In white dwarfs with low-mass progenitors, or in which evolution is interrupted before the onset of helium fusion, the core can be helium. The main sequence lifetimes of progenitors to helium-core white dwarfs (stars with mass $\lesssim 0.7 M_{\odot}$) are exceedingly long, and so many helium-core white dwarfs are thought to have formed as a result of binary interaction (Althaus et al., 2010).

The radii of white dwarfs are generally on the order of Earth radii, and the mass distribution is tightly clustered around an average of $0.6 M_{\odot}$. An upper limit on the mass of a white dwarf can be found in the Chandrasekhar mass limit, $M_{\text{Ch}} \approx 1.4 M_{\odot}$. Beyond this mass, electron degeneracy pressure can no longer




Figure 1.1: Interior structures of white dwarfs with helium (*above*) and hydrogen (*below*) atmospheres. The x-axis shows the radial depth into the white dwarf in terms of enclosed mass, with the core being at 0, and the y-axis shows the mass fraction of each element. This is Figure 17 from Althaus et al. (2010).

support the white dwarf against its own gravity, and the white dwarf collapses. If the core is carbon and oxygen, this collapse triggers a runaway ignition of carbon-oxygen fusion which destroys the white dwarf. This thermonuclear explosion is a Type Ia supernova (SN Ia).

1.2 Binary evolution and the processes involved

For many binary systems, the separation of the two stars is wide enough that the stars evolve as single stars, unaffected by their companions. However, in close binary systems, the evolution of the component stars is affected by interactions between the two stars. The separation of the two stars is therefore a key parameter of the system. For the binaries this thesis is concerned with, an initial phase of interaction known as the ‘common envelope phase’ decreases the separation of the two stars. Following this, the orbital separation is driven by two key processes: the loss of angular momentum from the system, and any transfer of mass between the two stars.

1.2.1 Common envelope evolution

The large changes in radius that a star undergoes as it passes through the RGB and AGB phases can cause a star to expand until its radius is similar to, or greater than, the distance between itself and its companion. If this happens, a phase of unstable mass transfer occurs between the two stars, as the more evolved star attempts to transfer matter onto its less evolved companion. The rate of this mass transfer is higher than the rate at which the accretor can receive matter. A loose constraint on this limiting accretion rate can be found in the Eddington limit. However, tighter constraints might be found by considering what happens to the accreted matter. If the accretion is fast relative to the dynamical timescale of the accretor, the accretor will become surrounded by an expanding shell of the accreted material. If the shell grows to fill the Roche lobe of the accretor, no further accretion can occur. The material that cannot be accreted becomes an envelope of matter surrounding the two stars, beginning the so-called ‘common envelope phase’ of evolution (Ivanova et al., 2013).

The common envelope provides a source of friction for the two stars, extracting energy from the system. This causes the two stars to inspiral towards each other, which (by Kepler’s laws) reduces the orbital period. After a relatively short time (of order 1000 years, Ivanova et al., 2013), the material is ejected from the binary. The more evolved star is generally left as a white dwarf, while the less evolved star

remains on the main sequence. Binaries in which at least one star is a white dwarf are known as ‘compact binaries’.

The implications of this interaction are highly important for the future of the binary. The loss of matter from the more evolved star can alter its nature, in some cases terminating evolution before helium fusion can begin (resulting in a helium-core white dwarf), and in some cases stripping the atmosphere to reveal the semi-degenerate helium-burning core (resulting in a ‘hot subdwarf’ star; Paczynski, 1976; Heber, 2016). The loss of significant energy from the system allows for the formation of short-period binaries with small orbital separations, some of which will undergo further interactions, such as further phases of mass transfer. Cataclysmic variables (CVs), which are sources of nova and dwarf nova outbursts, are made possible by common envelope evolution (Warner, 1995). An additional class of objects known as post-common envelope binaries (PCEBs) are thought to form a stage of evolution between the common envelope phase and CVs (Rebassa-Mansergas et al., 2016).

While it is qualitatively understood, there are still numerous uncertainties surrounding common envelope evolution. These result from a combination of the difficulty modelling a common envelope binary (due to the range of distances and timescales involved) and a lack of empirical constraints.

1.2.2 Angular momentum loss

The loss of angular momentum is one of the driving forces in binary evolution. In a non-accreting binary, angular momentum loss decreases the separation between the two stars. In an accreting binary, the rate of angular momentum loss determines the rate of accretion, and hence whether accretion will be stable or unstable. For a system that has passed through a phase of common envelope evolution, there are two key mechanisms of angular momentum loss: magnetic braking and gravitational wave radiation.

Magnetic braking results from interactions between the magnetic field and stellar wind of the star. The star’s outflowing stellar wind attempts to move along the magnetic field lines of the star. As the star rotates, this outflowing material attempts to rotate with it. Thus, angular momentum is transferred from the star’s rotation to the outflowing material, and eventually lost from the star. If the star is tidally locked, as is common for main sequence stars in close binaries, this in turn transfers angular momentum from the orbit of the binary (Verbunt and Zwaan, 1981).

Magnetic braking becomes significantly weaker for main sequence donor stars at a mass of $0.20 \pm 0.02M_{\odot}$ (Knigge, 2006), a mass that coincides with an orbital

period of ≈ 3 hours and approximately coincides with the star becoming fully convective. It is notable that fully convective main sequence stars do still show strong magnetic fields, and it is therefore unclear why magnetic braking weakens at this point, though it may be that the structure of the magnetic field changes in such a way that magnetic braking becomes weaker. For binaries with $P_{\text{orb}} \lesssim 3$ hours, magnetic braking is weaker, though it may still be dominant for many CVs in this period range (Knigge et al., 2011, found that the angular momentum loss in CVs in this period range is approximately 2.5 times larger than would be predicted for gravitational wave radiation alone). Magnetic braking is also negligible for white dwarfs and hot subdwarfs, and for binaries consisting only of these objects, the dominant form of angular momentum loss is gravitational wave radiation.

Gravitational waves are emitted by all rotating, asymmetric bodies or systems. These waves consist of small-scale fluctuations in the metric of spacetime. Originally predicted by Einstein (1916), they were first measured directly a century later by Abbott et al. (2016).

The amplitude of gravitational wave emission from a binary can be found by the relation of Landau and Lifshitz (1971):

$$\frac{\dot{J}}{J} = -\frac{32}{5} \frac{G^3}{c^5} \frac{M_1 M_2 (M_1 + M_2)}{a^4}, \quad (1.1)$$

where J is the angular momentum of the system, M_1 and M_2 refer to the accretor and donor star masses, a is the orbital separation, and G and c refer respectively to the gravitational constant and the speed of light. Due to the large masses and small separations involved, compact binaries can emit a significant amount of their angular momentum as gravitational waves (Kraft et al., 1962). Period evolution measurements of compact binaries found indirect measurements of gravitational wave radiation many years before gravitational waves were discovered directly (Hulse and Taylor, 1975; Taylor and Weisberg, 1982).

1.2.3 Mass transfer

Following the end of its common envelope evolution phase, a PCEB consists of a white dwarf and a main sequence star with a reduced orbital separation. As the binary radiates angular momentum, the separation of the two stars decreases further. Gravitational distortion of the main sequence star become significant as the stars move closer together.

By considering the gravitational fields of the two stars, and a centrifugal force component, we can define a gravitational potential in the corotating reference




Figure 1.2: Schematic diagram of the evolution of a Roche lobe filling binary. *Top:* Both stars begin separate. As they inspiral, the main sequence star approaches the limits of its Roche lobe. *Middle:* The donor is now filling its Roche lobe and overflowing. The infalling material at first follows a ballistic trajectory, as described by a test particle in freefall. *Bottom:* The infalling material circularises. Viscosity causes the material to spread into an accretion disc.

frame (Kruszewski, 1966):

$$\Phi = -\frac{GM_1}{\sqrt{x^2 + y^2 + z^2}} - \frac{GM_2}{\sqrt{(x-a)^2 + y^2 + z^2}} - \frac{\omega^2}{2} \left[\left(x - \frac{aM_2}{M_1 + M_2} \right)^2 + y^2 \right] \quad (1.2)$$

where x , y , and z are corotating coordinates centred on the white dwarf, M_1 and M_2 are the masses of the white dwarf and main sequence star, a is the orbital separation of the two stars, and $\omega = 2\pi/P_{\text{orb}}$ is the angular velocity of the binary.

For each star, the largest enclosed surface of equal potential is its Roche lobe, which has potential equal to that of the inner Lagrange point between the two stars (see Fig. 1.2). The inner Lagrange point is the point on the line between the centres of the two stars at which a test particle will remain in the same position with respect to the two stars. A sphere whose volume is equal to the volume enclosed by the Roche lobe will have a radius (Paczynski, 1971; Eggleton, 1983)

$$R = 0.462a \left(\frac{q}{q+1} \right)^{1/3} \quad (1.3)$$

for $q < 0.3$, where $q = M_2/M_1$ is the ratio of component stellar masses.

As the separation of the stars decreases, the Roche lobes shrink. The main sequence star, being lower mass but larger radius, is deformed by the gravity of the white dwarf and eventually fills its Roche lobe. As the lobe continues to shrink, the low mass star begins to overflow its Roche lobe. Some fraction of the star's atmosphere will then be in the Roche lobe of the white dwarf, and feel a net gravitational pull towards the white dwarf. Hence, mass transfer between the two stars begins.

The transferred material begins its journey by falling towards the central white dwarf on a ballistic path (described by a test particle in freefall). However, without interaction, the material's angular momentum will prevent it from reaching the white dwarf. It is therefore necessary to form an accretion disc around the central white dwarf, which allows angular momentum to be redistributed among the material. The accretion disc can be thought of as a series of circular annuli, each rotating at a Keplerian velocity, with a weak interaction between each annulus and its neighbours termed 'viscosity'. This viscosity seeks to equalise the angular velocities of the annuli – in other words, to accelerate the outer annuli and decelerate the inner annuli. By this process, angular momentum is transported outwards through the disc, allowing matter in turn to be transported inwards (Shakura and Sunyaev, 1973).

As the donor loses mass, its radius will decrease. In order for mass transfer to be stable, this decrease in radius must be matched by a decrease in the radius of

its Roche lobe, which must be caused by a decrease in the orbital separation. If the Roche lobe shrinks too quickly, the mass transfer rate will increase and an unstable feedback loop develops. If it shrinks too slowly, mass transfer will cease. We can calculate the stability condition for mass transfer as follows. The rate of change of the separation of the binary can be derived as (Hellier, 2001)

$$J = M_1 M_2 \sqrt{\frac{Ga}{M}} \implies \frac{\dot{a}}{2a} = \frac{\dot{J}}{J} - \frac{\dot{M}_1}{M_1} - \frac{\dot{M}_2}{M_2} + \frac{\dot{M}}{2M} \quad (1.4)$$

where J is the angular momentum of the binary, $M = M_1 + M_2$ is the total mass of the binary, and G is the gravitational constant. Assuming mass transfer is conservative ($\dot{M} = 0$), we can combine this with Equation 1.3 to show that the Roche lobe radius (R_L) reacts as (Frank et al., 2002)

$$\frac{\dot{R}_L}{R_L} = 2\frac{\dot{J}}{J} + 2\frac{(-\dot{M}_2)}{M_2} \left(\frac{5}{6} - q \right). \quad (1.5)$$

There are then two cases, depending on whether q is larger or smaller than $5/6$. If $q > 5/6$, the Roche lobe around the donor will shrink in response to mass loss, causing runaway mass transfer. This high rate of mass transfer will continue until $q < 5/6$. Such episodes of ‘unstable’ mass transfer are a likely model for ‘supersoft X-ray binaries’, and are expected to be short-lived. Most accreting binaries must therefore be in the range $q < 5/6$. In this range, \dot{M} would cause the Roche lobe to grow; this effect must therefore be counteracted by some angular momentum loss, through the mechanisms discussed in the previous section.

Fig. 1.3 shows the orbital period distribution of known CVs. A deficiency of systems between 2 and 3 hours, known as the ‘period gap’, can be explained in terms of the above discussion. This gap coincides with the turning off of magnetic breaking, resulting in a significant decrease in \dot{J} . The remaining angular momentum loss from gravitational wave radiation is not sufficient to drive mass transfer for many systems in this period range, causing mass transfer to cease. Such systems continue to inspiral as non-accreting binaries, their separation decreasing and gravitational wave radiation increasing, until they evolve into contact and mass transfer begins again.

In Fig. 1.3 the ‘period minimum’ of hydrogen-accreting CVs can also be seen at ~ 78 min. When a CV reaches this limit, its donor has lost sufficient mass to become degenerate (though note that the radius of the donor is different to that of a degenerate star in isolation). As the mass-radius relation of the donor changes, the period evolution of the binary changes. By differentiating with respect to time




Figure 1.3: Orbital period distribution of known CVs. The period gap and period minimum are labelled. A small number of systems have periods below the period minimum, of which the majority are AM CVns. This figure is produced from data collected by Ritter and Kolb (2003, update RKCcat7.24, 2016), which includes systems known prior to 31 December 2015.

the well-known relation between orbital period and donor density, $P_{\text{orb}}^2 \propto MR^{-3}$ (Faulkner et al., 1972), we can show that the orbital period will grow if

$$\frac{\dot{M}}{M} - \frac{1}{3} \frac{\dot{R}}{R} > 0. \quad (1.6)$$

For an assumed mass-radius relation of the form $M \propto R^\alpha$, this will be true if $\alpha < 1/3$.

Considering Equation 1.5 and the above discussion, the period bounce can be considered as follows. As the donor becomes more degenerate, its response to mass loss changes until it reaches the critical $\alpha < 1/3$ limit. This causes an increase in \dot{M}_2 , which in turn causes the donor response to evolve more quickly and so the process feeds back. If \dot{M}_2 is able to increase sufficiently that the \dot{M}_2 term outweighs the \dot{J} term, the binary will begin to expand and remain stable. If not, the binary will merge. Most CVs that approach the period minimum are expected to survive by going through a period ‘bounce’, in which they reverse their migration through period space and begin to evolve to longer orbital periods (Faulkner, 1971; Rappaport et al., 1982). Such ‘post-bounce’ systems are rarely seen, despite models predicting that they should constitute the majority of CVs. Models such as the Consequential Angular Momentum Loss model suggest that an as-yet-unknown additional source of angular momentum loss may increase the rate of mergers (Schreiber et al., 2016; Belloni et al., 2018).

Despite this period minimum, Fig. 1.3 shows a number of systems at sub-minimum periods. The majority of these systems are the AM CVn type binaries, a class of ultracompact CVs in which the accreted material is dominated by helium. These objects are the subject of this thesis, and will be introduced in Section 1.3.

1.3 AM CVn type binaries

More than 60 years ago, Greenstein and Matthews (1957) obtained a spectrum of what they believed to be a helium-atmosphere white dwarf. The spectrum shows a series of helium absorption lines, but there are asymmetries in the lines unlike those standardly seen in an isolated white dwarf. Instead the line profiles are reminiscent of those seen in CV binaries, except that the lines are helium rather than hydrogen. Ten years later, Smak (1967) identified an 18 min photometric period from the object, which he suggested may be an orbital period. (A period that was confirmed to be orbital by Nelemans et al., 2001a). The system was given a variable star designation, AM Canum Venaticorum (AM CVn).

Several years after the discovery of AM CVn, GP Com was discovered (Warner, 1972). It appeared similar to AM CVn, but its spectrum was dominated by a series of helium emission lines rather than absorption lines. Following this, CR Boo became the third object of the class (Nather, 1985). CR Boo undergoes large changes in magnitude (‘outbursts’), alternating between a bright, absorption line state (similar to AM CVn) and a dim, emission line state (similar to GP Com)¹. It became clear that these objects all belong to one class of variable, which was given the designation ‘AM CVn-type binaries’².

The structure of an AM CVn binary is similar to the CV structure described in previous sections, with two key distinguishing features. Firstly, AM CVn systems have short orbital periods in the range of 5–65 min, all below the period minimum for typical CVs. Secondly, the accreted material in AM CVn binaries is dominated by helium, with most systems showing no detectable hydrogen. These properties combine to imply an accreting compact binary like a CV, but in which the donor is an evolved star that is both dense (to explain the short orbital periods) and dominated by helium. Either a white dwarf or a stripped, helium-burning star (akin to a hot subdwarf) are possible candidates. The nature of the donor star is closely tied to open questions about the evolutionary history of AM CVn binaries, which I introduce in Section 1.4 and discuss throughout this thesis.

Over the decades since their initial discovery, the number of known AM CVn binaries has climbed (Figure 1.4). Since the turn of the millennium, the advent of large-scale surveys has hastened the process. The Sloan Digital Sky Survey (SDSS) provided spectroscopy of a number of objects in the Northern sky, from which several AM CVn systems were detected (Anderson et al., 2005; Roelofs et al., 2005, 2007d,a). Targeted spectroscopic follow-up based on SDSS colour photometry discovered several more (Roelofs et al., 2009; Carter et al., 2013, 2014b,a). Outbursting AM CVn systems have been detected by a number of transient surveys, including the Palomar Transient Factory (PTF, Levitan et al., 2011, 2013, 2014), the Catalina Realtime Transient Survey (CRTS, Breedt et al., 2014), and the All-Sky Automated Search for SuperNovae (ASAS-SN, Kato et al., 2015). Most pertinent to this thesis, the system Gaia14aae was discovered in outburst by *Gaia* Science Alerts in 2014 (Campbell et al., 2015, see also Chapters 3 and 4 of this work). The most complete published list of AM CVns includes 56 systems (Ramsay et al., 2018). In Table 1.1 I present a list of 56 known AM CVns.

It is thought that the known sample of AM CVn systems may not be com-

¹Outbursts are discussed further in Section 1.5

²For the rest of this thesis I will refer to the original object, AM CVn, as ‘AM CVn itself’




Figure 1.4: Timeline of AM CVn binary discoveries. The discoveries of several key systems are labelled, as are the launch dates of two key surveys, SDSS and PTF, both of which significantly increased the rate of AM CVn discoveries.

plete, even to distances of 100 pc (Ramsay et al., 2018). Given the inhomogeneous methods by which AM CVn binaries have been discovered, the selection effects on the entire sample are complex. The subset with the best understood selection effects is found by selecting all known systems within the footprint of SDSS and a limiting magnitude of 19. It is believed that this sample is complete; however, only four systems fit these criteria. Using this small sample, Carter et al. (2013) estimated a space density for the AM CVn population of $(5 \pm 3) \times 10^{-7} \text{ pc}^{-3}$. Despite the large uncertainties, this is the current best constraint on the space density of AM CVn binaries. This density falls short of predictions based on population synthesis models (Nelemans et al., 2001b, 2004; Roelofs et al., 2007a) by a factor of ten, a discrepancy which is still not understood (Nissanke et al., 2012).

1.4 The evolution of AM CVn binaries

The focus of this thesis is on exploring empirical constraints on the prior evolution of AM CVn binaries. Three formation channels have been proposed for AM CVn binaries (see Fig. 1.5). The fractions of systems that form by each channel, and indeed whether all three channels are possible, are currently poorly constrained. It is entirely plausible that all three channels form AM CVn systems. However, the low space density of AM CVn systems relative to predictions can be explained more easily if one or more channels does not contribute to the AM CVn population.

From the empirical properties of AM CVn stars (their short orbital periods and hydrogen-deficient spectra) it is clear that the donor in the binary is an evolved object that is both dense and helium-dominated. The three evolutionary models are various attempts to explain the precise nature of the donor, and how it formed.

In the white dwarf donor channel (Paczynski, 1967; Deloye et al., 2007) and the helium star donor channel (Savonije et al., 1986; Iben and Tutukov, 1987; Yungelson, 2008), the binary passes through two phases of common envelope evolution as each star in the binary evolves from the main sequence. These phases resemble the common envelope phase described in Section 1.2.1. Each phase extracts angular momentum from the binary. The two channels differ in the nature of the lower-mass star (the future donor) following the second phase of common envelope evolution. In the white dwarf donor channel, the secondary star is a white dwarf. In the helium star donor channel, the secondary star is a helium-burning, helium-atmosphere star, which resembles the core of a giant star following the helium flash.

In these channels, AM CVn binaries begin mass transfer at short orbital periods (5–20 minutes, Deloye et al., 2007) and grow to longer orbital periods through-




Figure 1.5: The proposed evolutionary channels of AM CVn binaries. This figure is Fig. 2 from Solheim (2010), originally prepared by Lev Yungelson.




Figure 1.6: The period evolution of AM CVn binaries passing through each of the three formation channels. This is Fig. 1 from Nelemans et al. (2010).

out their lives (see Fig. 1.6). Extremely short-period AM CVns such as HM Cnc and V407 Vul can only have formed by the white dwarf donor channel. However, the long-term stability of these short period binaries is uncertain. The orbital periods of both systems are still decreasing, and it is unknown whether the binaries will merge (Strohmayer, 2002, 2004, 2005; Hakala et al., 2003, 2004; Israel et al., 2004; D’Antona et al., 2006; Deloye and Taam, 2006; Kaplan et al., 2012).

The third proposed channel is the evolved CV channel (Podsiadlowski et al., 2003; Goliash and Nelson, 2015). In this channel, the donor begins to evolve from the main sequence shortly before mass transfer starts. The donor has therefore already built up a core of helium, which is revealed as its hydrogen atmosphere is removed by mass transfer. Binaries forming by this channel evolve from long orbital periods, and undergo a period bounce at approximately 40 min (although the exact period of the bounce is a function of the donor properties; Nelemans et al., 2010). The channel therefore strongly favours the formation of longer period AM CVns.

Population synthesis models have found that either the white dwarf donor channel or the helium star donor channel should dominate the population, depending on the assumptions used regarding the stability of mass transfer in each channel (Nelemans et al., 2001b; Nisanke et al., 2012). The evolved CV channel has generally been disfavoured by evolutionary models (Nelemans et al., 2004; Goliash and Nelson, 2015). It should be noted that many population synthesis models have focused on short-period AM CVn binaries ($\lesssim 25$ min; Nelemans et al., 2001b, 2004; Nisanke et al., 2012), and that the fraction of results from the three formation tracks is likely to be different at longer periods.

The white dwarf donor channel is particularly prone to uncertainties about the fraction of inspiralling double white dwarfs (DWDs) that can reach a state of stable accretion and survive as AM CVns, versus the fraction that accrete unstably and merge. A key uncertainty concerns the efficiency with which tidal or magnetic forces in a young AM CVn binary can synchronise the spin of the accretor with the orbit of the system. As the radius of the central white dwarf is significant compared to the separation of the system, the angular momentum transferred from the donor’s orbit to the accretor’s spin via mass transfer is significant to the evolution of the system. Some synchronisation is required in order to counteract this angular momentum transfer. A synchronisation timescale of ≈ 1000 years is required if a significant number of white dwarf binaries are to survive; estimates have ranged from $\ll 500$ years to $\sim 10^{12}$ years (Marsh et al., 2004).

Another uncertainty concerns helium novae. AM CVn binaries may experience helium nova eruptions in which a significant amount of the accreted material is

ejected, similar to the nova eruptions seen in hydrogen CVs (eg. V445 Pup; Woudt et al., 2009). It was suggested by Shen (2015) that any DWDs which begin mass transfer may merge as a result of friction with their ejected material following their first helium nova.

Evidence found in the composition of the accreted material can give constraints on the nature of the donor star (Yungelson, 2008; Nelemans et al., 2010). In the white dwarf donor or evolved CV donor models, the donor is a helium-core object that has not undergone any helium burning. As such, the composition of carbon, nitrogen and oxygen (CNO) elements is expected to approximately equal the equilibrium ratio of these elements during the CNO cycle, a hydrogen-burning process seen in hot stellar cores. On the other hand, if the donor is a helium-burning star, then core helium-burning via the triple- α process will have produced an excess of the helium-burning products carbon and oxygen relative to nitrogen. The ratio of carbon to nitrogen can therefore constrain the nature of the donor. Many AM CVn systems show spectroscopically visible nitrogen (eg. Kupfer et al., 2016), while few show visible carbon. This has been interpreted as evidence against the helium star channel.

The evolved CV channel is thought to contribute only a minority of AM CVn binaries to the total population. The key difficulty for this channel is explaining the non-detection of hydrogen in almost all AM CVn binaries. For most sets of initial conditions that describe a binary, models show that the timescale required to remove all hydrogen is very long. Only a narrow range of initial conditions can produce AM CVn binaries with no detectable hydrogen within a Hubble time (Goliash and Nelson, 2015). As the range of initial parameters that can produce an AM CVn binary decreases, the probability of such an AM CVn binary forming in nature becomes low. Around ten CVs have been discovered with orbital periods of 50–70 min which accrete material that is helium-enriched while still containing hydrogen (Augusteijn et al., 1993; Breedt et al., 2012, 2014; Carter et al., 2013; Littlefield et al., 2013; Breedt, 2015). These helium-rich CVs must have evolved donors, and are evidence that some CVs do indeed traverse the evolved CV channel. We might expect the number of AM CVns forming by this channel to be of a similar order.

Studies of individual AM CVns, in order to determine certain key characteristics, can be used to place constraints on the prior evolution of each binary and thereby constrain the formation channels of the population. The key properties of interest are generally the binary orbital period (P_{orb}) and the mass ratio (q) of the two stars. Several methods exist for characterising AM CVn binaries, including




Figure 1.7: Orbital period distribution of AM CVn-type binaries with confirmed orbital periods. I include all systems in Table 1.1 for which the orbital period is known, but not superhump periods.

eclipse photometry, phase-resolved spectroscopy, and an empirical method known as the ‘superhump’ method. I will discuss these methods in more detail in the relevant chapters of this thesis.

The end stages of AM CVn evolution are uncertain. The vast majority of AM CVn accretor stars will not reach the Chandrasekhar mass limit, and are expected to contribute *at most* 1% to the SNe Ia population (Solheim and Yungelson, 2005). Edge-lit detonation triggered by helium accretion may increase this probability somewhat (see Maoz et al., 2014, for an overview). However, no AM CVn with mass measurements has a total mass even close to the Chandrasekhar mass limit, and it seems unlikely that more than a small fraction of AM CVn binaries will become supernovae. Instead, the majority of AM CVn binaries are likely to continue to accrete mass at lower and lower accretion rates, either indefinitely or until mass accretion eventually ceases. If mass transfer stops, the remaining system will be a white dwarf with a planetary-mass companion (termed ‘WD + Planet’ in Figure 1.5).

1.5 Orbital Periods and AM CVn Classification

One of the most important observable properties of an AM CVn binary is its orbital period, which plays a key role in determining the behaviour and future evolution of the binary. This is due largely to the steep relation between orbital period and the strength of gravitational wave radiation from the binary. Because the gravitational wave radiation drives both the mass transfer rate and the rate of orbital period evolution of the binary, orbital period correlates strongly with those properties. Fig. 1.7 shows the orbital period distribution of known AM CVn binaries.

The dependence of the gravitational wave radiation on orbital period can be understood as follows. From Equation 1.1, the change in angular momentum $\dot{J}/J \propto M_2^2/a^4$. Equation 1.3 gives $M_2/R_2^3 \propto P^{-2}$, for orbital period P . Kepler's law gives $a^3 \propto P^2$. Then, if the donor can be assumed to be degenerate, $R_2 \propto M_2^{-1/3}$ (Verbunt and Rappaport, 1988). Combining these gives

$$\frac{\dot{J}}{J} \propto P^{-14/3}. \quad (1.7)$$

As was discussed in Section 1.4, the assumption that the donor is degenerate is not necessarily valid for all systems, but it serves as an approximation.

As a result of this strong dependence, short period AM CV binaries are particularly strong sources of gravitational wave radiation. These gravitational waves will not be detected from the ground due to interference from seismic noise. However, to a space-based detector such as the Laser Interferometer Space Antenna (LISA), short-period AM CVn binaries will be among the brightest individual sources. Indeed, it has been suggested that known short-period AM CVn binaries, alongside short-period detached binaries, can be used to test the calibration of LISA when it launches (Nelemans, 2003; Korol et al., 2017; Breivik et al., 2018; Kupfer et al., 2018).

The gravitational wave radiation drives orbital period change in AM CVn binaries. The rate of orbital period change behaves as (Deloye et al., 2007)

$$\frac{\dot{P}_{\text{orb}}}{P_{\text{orb}}} = 3 \frac{\dot{J}}{J} \left(\frac{\xi - 1/3}{\xi + 5/3 - 2q} \right) \quad (1.8)$$

where the parameter $\xi = d \log(R)/d \log(M)$ characterises the donor star's response to mass loss. As a result of the strong period dependence of \dot{J} in Equation 1.7, short-period systems grow in orbital period much faster than longer period systems. It is therefore expected that short-period AM CVns have a much lower space density

than longer-period systems, as can be seen in Fig. 1.7; though it should be noted that the sample of known systems has selection effects in favour of short-period AM CVn binaries (due to their high luminosities) and mid-period systems (due to their outbursts).

Mass transfer rate is also driven by gravitational wave radiation, and behaves as

$$\frac{\dot{M}_2}{M_2} = \frac{\dot{J}}{J} \left(\frac{1}{\xi + 5/6 - q} \right). \quad (1.9)$$

The mass transfer rates of AM CVn binaries therefore vary widely across the period distribution, from $10^{-7} M_{\odot} \text{yr}^{-1}$ at the short period end to $10^{-11} M_{\odot} \text{yr}^{-1}$ at the long period end. This results in strong differences in the behaviour of AM CVn binaries as a function of period.

AM CVn binaries can be separated into four categories based on their accretion behaviour, as follows:

- In AM CVn binaries with periods $\lesssim 10$ min, accreted material impacts directly on the white dwarf with no accretion disc.
- AM CVn binaries with periods $10 \text{ min} \lesssim P_{\text{orb}} \lesssim 20$ min are ‘high state’ systems with hot, ionised accretion discs.
- AM CVn binaries with periods $20 \text{ min} \lesssim P_{\text{orb}} \lesssim 40$ min alternate between ‘outburst’ phases in which they resemble high state systems, and ‘quiescent’ phases in which they resemble low state systems.
- AM CVn binaries with period $\gtrsim 40$ min are in a permanent ‘low state’, with cold, neutral accretion discs.

Direct impact AM CVn binaries include only the two systems with the shortest orbital periods, HM Cnc (5.4 min) and V407 Vul (9.5 min). In these systems, the orbital separation of the two stars is so small that the accreting white dwarf takes up a significant fraction of its Roche lobe. The radius at which infalling material would normally circularise is within the accreting white dwarf, meaning that infalling material is able to fall directly onto the white dwarf and no accretion disc can form (Marsh and Steeghs, 2002; Ramsay et al., 2002; Roelofs et al., 2010).

High state systems have high accretion rates, driving the disc into a permanently ionised state. They are helium-accreting analogues to hydrogen-accreting nova-like CVs. Their flux is dominated by the accretion disc, and their spectra contain a series of helium absorption lines. AM CVn itself is an example of a high state system (see Fig. 1.8).




Figure 1.8: Spectrum of AM CVn itself, a high state binary. This is Fig. 1 from Roelofs et al. (2006b).

Conversely, low state systems have low accretion rates, leaving the accretion disc cold and neutral. Their flux is dominated by the central white dwarf, and their spectra contain helium emission lines. Fig. 1.9 shows example spectra of the low state binaries GP Com and V396 Hya.

Between the high and low state systems, outbursting systems such as CR Boo alternate between low state ‘quiescent’ periods and high state ‘outbursts’. Observationally, these outbursts are accompanied by changes in the brightness of the system of one or several magnitudes. Outbursting AM CVn binaries are analogous to dwarf nova CVs, and the outburst mechanism is similar (Kotko et al., 2012; Cannizzo and Nelemans, 2015). During quiescence, the rate of mass transfer from the donor into the accretion disc is higher than the accretion rate onto the white dwarf. Material therefore builds up in the accretion disc, leading to increasing density and temperature. At a certain point, the conditions cause the disc material to ionise. Ionisation increases the viscosity of the accretion disc, allowing for more efficient transfer of angular momentum towards the outer disc and increasing the rate of mass flow through the accretion disc and onto the central white dwarf. The accretion rate of the central white dwarf is now higher than the mass transfer rate of




Figure 1.9: Spectra of GP Com and V396 Hya, two low state binaries. This is Fig. 1 from Kupfer et al. (2016).

the donor, and the disc material begins to empty onto the white dwarf. Once the density has decreased sufficiently, the material once again becomes neutral, and the cycle repeats.

The above model of outbursts is complicated by the appearance of distinct ‘outbursts’ and ‘superoutbursts’ (Levitán et al., 2015). Superoutbursts are seen in AM CVn binaries and short-period CVs. They are distinguished from standard outbursts by having a larger increase in magnitude and a longer duration. A photometric signal known as a ‘superhump’ signal is often seen during superoutburst, occurring at the beat frequency between the orbital period and the period of precession of an eccentric accretion disc (Patterson et al., 1993; Kato et al., 2014). The mechanism for a superoutburst is believed to involve tidal heating of the eccentric disc by the donor star. If the mass ratio between the donor and accretor is sufficiently extreme ($q \lesssim 0.3$), the disc will contain an annulus orbiting three times for every orbital period. The 3:1 resonance between this accretion disc material and the orbit of the donor is sufficient to drive the disc to be eccentric and cause tidal heating. This tidal heating increases the magnitude of the outburst and sustains it beyond its standard duration. The superhump signal occurs at the frequency of interaction between the donor star and the accretion disc, and is a useful tool to probe the mass ratio of the binary, as I will discuss in Chapter 5.

Outbursts are an important observational property of mid-period AM CVns. As discussed in Section 1.3, a significant fraction of known AM CVns have been found in transient surveys due to this long-term variability. The recurrence rate of outbursts correlates strongly with orbital period (Levitan et al., 2015), with short period systems outbursting frequently (for instance, ASASSN-14cc is the only known AM CVn to spend the majority of its time in outburst), and long period systems outbursting much less frequently (no repeat outbursts have been detected for any AM CVn beyond 36.6 min).

The dependence of these classes on orbital period is evidence for the strong correlation between period and mass transfer predicted in Equations 1.7 and 1.9. However, there is overlap in terms of orbital period along the borders of these classes. The high state system CXOGBS J1751–2940 ($P_{\text{orb}} = 22.9$ min) is believed to have a longer orbital period than the outbursting system PTF J1919+4815 ($P_{\text{orb}} = 22.5$ min). Likewise, at the long period edge, outbursts have been seen in Gaia14aae ($P_{\text{orb}} = 49.7$ min) and ASASSN-17fp ($P_{\text{orb}} \approx 51$ min), while no outburst has been seen from GP Com ($P_{\text{orb}} = 46.5$ min) during the 45 years it has been known. The overlap in orbital period at the boundaries of these categories can be seen as evidence for inhomogeneity in the other AM CVn parameters in Equation 1.9, due to systems having either different donor stars natures (changing ξ) or different mass accretor stars (changing M_1).

1.6 Overview

The body of work brought together in this thesis consists of studies of several individual AM CVn binaries. The purpose of each study is to provide constraints on the formation channels of AM CVn binaries.

The thesis is organised as follows. In Chapter 2, I present several of the methodologies employed throughout the thesis, including instruments used to obtain high-speed photometry, modelling methods applied to that photometry, spectroscopic instruments, and Doppler tomography. More specific methodology details are found in the relevant chapters. In Chapter 3, I introduce the eclipsing AM CVn binary Gaia14aae, the first-known fully-eclipsing AM CVn, and present an analysis of the system using eclipse photometry. Chapter 4 presents a complementary analysis of Gaia14aae using phase-resolved spectroscopy. Chapter 5 presents the discovery of a new short-period AM CVn in a *K2* observing field, a period analysis of the binary, and parameters derived using the ‘superhump’ method. In Chapter 6, I present phase-resolved spectroscopy on five binaries, of which two are AM CVns

and three are evolved, helium-enriched CVs, and a preliminary analysis of these data.

Table 1.1: A summary of all known AM CVn binaries. For each period a method is specified, either spectroscopy (S), photometry (P), eclipse timing (E), or superhumps (sh). In the first three methods the period is orbital; in the fourth, it is the superhump period. State specifies the classification of the system, either direct impact (‘direct’), high state (‘high’), outburst detected (‘out’), or no outburst detected (‘low’). Magnitudes are during quiescence for outbursting systems. All distances are taken from Ramsay et al. (2018) and have 95% uncertainties.

Name	Coordinates	Period (min)	State	Magnitude	d (pc)	Refs
HM Cnc	08:06:22.84 +15:27:31.5	5.4 (S)	direct	21.1V	–	1,2,3
V407 Vul	19:14:26.09 +24:56:44.6	9.5 (S)	direct	>19.2R	1786 ± 667	4,5,6
ES Cet	02:00:52.17 –09:24:31.7	10.3 (P)	high	17V	1584 ± 291	7,8,9
SDSSJ1351–0643	13:51:54.46 –06:43:09.0	15.7 (S)	high	18.6g	1317 ± 531	10,11
AM CVn	12:34:54.60 +37:37:44.1	17.1 (S)	high	14.02V	299.1 ± 4.4	12,13
SDSSJ1908+3940	19:08:17.07 +39:40:36.4	18.2 (S)	high	16.08g	1044 ± 51	14,15
HP Lib	15:35:53.08 –14:13:12.2	18.4 (S)	high	13.59V	276 ± 4	16,17,18
PTF1J1919+4815	19:19:05.19 +48:15:06.2	22.5 (E)	out	20.16g	1339 ± 555	19
ASASSN-14cc	21:39:48.24 –59:59:32.4	22.5 (sh)	out	20	1019 ± 108	20
CXOGBS J1751–2940	17:51:07.6 –29:40:37	22.9 (P)	high	17.62r	971 ± 156	21
CR Boo	13:48:55.22 +07:57:35.8	24.5 (P)	out	14.5V	–	22,23,24,25
KL Dra	19:24:38.28 +59:41:46.7	25 (P)	out	20	956 ± 153	26,27
V803 Cen	13:23:44.54 –41:44:29.5	26.6 (S)	out	14V	–	28,18
PTF1J0719+4858	07:19:12.13 +48:58:34.0	26.8 (S)	out	19.4g	861 ± 304	29
ASASSN-15kf	15:38:38.24 –30:35:49.7	27.67 (sh)	out	19.4	–	30
YZ LMi	09:26:38.71 +36:24:02.4	28.3 (E)	out	19.0g	577 ± 324	31,32
CP Eri	03:10:32.76 –09:45:05.3	28.4 (P)	out	19.7	964 ± 615	33,34

Continued on next page

Table 1.1 – *continued*

Name	Coordinates	Period (min)	State	Magnitude	d (pc)	Refs
CRTS J0910–2008	09 10 17.45 –20 08 12.5	29.7 (sh)	out	20.4g	–	35
PTF1J0943+1029	09:43:29.59 +10:29:57.6	30.4 (S)	out	20.69g	–	36
CSS010550+190317 [§]	01:05:50.10 +19:03:17.2	31.6? (?)	out	19.6g	734 ± 374	
V406 Hya	09:05:54.79 –05:36:08.6	33.8 (S)	out	20.5	504 ± 493	37,9
PTF1J0435+0029	04:35:17.73 +00:29:40.7	34.3 (S)	out	21.12g	–	36
SDSSJ1730+5545	17:30:47.59 +55:45:18.5	35.2 (S)	low	20.14g	911 ± 420	38
NSV 1440	03:55:17.83 –82:26:11.5	36.3 (sh)	out	18.5G	677 ± 595	39
2QZ J1427–01	14:27:01.70 –01:23:10.0	36.6 (sh)	out	20.3	337 ± 17	40
SDSSJ1240–0159	12:40:58.03 –01:59:19.2	37.4 (S)	out	19.7V	577 ± 365	41
SDSSJ0129+3842	01:29:40.06 +38:42:10.5	37.6 (S)	out	19.8	508 ± 239	31,42
SDSSJ1721+2733 [†]	17:21:02.48 +27:33:01.2	38.1 (?)	out	20.1	995 ± 578	43,44
ASASSN-14mv [§]	07:13:27.28 +20:55:53.4	41.8 (sh)	out	17.3	247 ± 7	45,46
ASASSN-14ei	02:55:33.39 –47:50:42.0	43 (sh)	out	18	255 ± 4	47
SDSSJ1525+3600	15:25:09.58 +36:00:54.6	44.3 (S)	out	19.8g	524 ± 90	42,44
SDSSJ0804+1616	08:04:49.49 +16:16:24.8	44.5 (S)	low	18.2g	828 ± 173	48
SDSSJ1411+4812	14:11:18.31 +48:12:57.6	46 (S)	low	19.4g	429 ± 65	31,49
GP Com	13:05:42.43 +18:01:04.0	46.5 (S)	low	15.94	73 ± 0.4	50,51,52
CSS045020–093113	04:50:19.82 –09:31:12.8	47.3(sh)	out	20.5	–	53
SDSSJ0902+3819	09:02:21.36 +38:19:41.9	48.31 (S)	out	20.2g	461 ± 435	44
Gaia14aae	16:11:33.97 +63:08:31.8	49.7 (E)	out	18.0g	259 ± 11	10,54
SDSSJ1505+0659 [¶]	15:05:51.58 +06:59:48.7	50.6? (S)	low	19.11g	160 ± 12	55

Continued on next page

Table 1.1 – *continued*

Name	Coordinates	Period (min)	State	Magnitude	d (pc)	Refs
ASASSN-17fp	18:08:51.10 –73:04:04.2	51 (sh)	out	20+g	–	56
SDSSJ1208+3550	12:08:41.96 +35:50:25.2	52.96 (S)	low	18.8g	202 ± 18	57,42
SDSSJ1642+1934	16:42:28.08 +19:34:10.1	54.2 (S)	low	20.3	1044 ± 604	42,44
SDSSJ1552+3201	15:52:52.48 +32:01:50.9	56.3 (S)	low	20.2g	443 ± 227	31,58
SDSSJ1137+4054	11:37:32.32 +40:54:58.3	59.6 (S)	low	19.00g	209 ± 14	55
V396 Hya	13:12:46.93 –23:21:31.3	65.1 (S)	low	17.6	93.6 ± 1.4	59
SDSSJ1319+5915	13:19:54.47 +59:15:14.84	65.6 (S)	low	19.1g	205 ± 10	60
Gaia16all	06:27:20.53 –75:13:33.35	–	out	20.6G	956 ± 605	61
CRTSJ0744+3254 [‡]	07:44:19.70 +32:54:48.0	–	out	20.60g	–	62
SDSSJ0807+4852	08:07:10.33 +48:52:59.6	–	out	20.39g	883 ± 648	63
CRTSJ0844–0128	08:44:13.60 –01:28:07.0	–	out	20.00g	1474 ± 597	62
PTF1J0857+0729	08:57:24.27 +07:29:46.7	–	out	21.83g	–	36
SDSSJ1043+5632	10:43:25.08 +56:32:58.1	–	out	20.28g	979 ± 575	64
PTF1J1523+1845	15:23:10.71 +18:45:58.2	–	out	23.27g	–	36
PTF1J1632+3511*	16:32:39.39 +35:11:07.3	–	out	22.99g	–	36
MOA 2010-BLG-087	18:08:34.85 –26 29 22.75	–	out	20.0g	–	
SDSSJ2047+0008	20:47:39.40 +00:08:40.3	–	low	24g	–	57
PTF1J2219+3135	22:19:10.09 +31:35:23.1	–	out	20.38g	–	36
ASASSN-14fv	23:29:55.13 +44:56:14.4	–	out	20.5	–	65

Continued on next page

Table 1.1 – *continued*

Name	Coordinates	Period (min)	State	Magnitude	d (pc)	Refs
References: [1] Roelofs et al. (2010); [2] Israel et al. (1999); [3] Israel et al. (2002); [4] Cropper et al. (1998); [5] Steeghs et al. (2006); [6] Motch et al. (1996); [7] Copperwheat et al. (2011b); [8] Warner and Woudt (2002); [9] Woudt and Warner (2003); [10] Green et al. (2018b); [11] Green et al. (2018a); [12] Nelemans et al. (2001a); [13] Smak (1967); [14] Fontaine et al. (2011); [15] Kupfer et al. (2015); [16] Roelofs et al. (2007b); [17] O’Donoghue et al. (1994); [18] Roelofs et al. (2007c); [19] Levitan et al. (2014); [20] Kato et al. (2015); [21] Wevers et al. (2016); [22] Patterson et al. (1997); [23] Provencal (1994); [24] Nather (1985); [25] Provencal et al. (1991); [26] Jha et al. (1998); [27] Wood et al. (2002); [28] O’Donoghue et al. (1987); [29] Levitan et al. (2011); [30] vsnet-alert 18669; [31] Anderson et al. (2005); [32] Copperwheat et al. (2011a); [33] Abbott et al. (1992); [34] Armstrong et al. (2012); [35] Aungwerojwit et al. (in prep); [36] Levitan et al. (2013); [37] Roelofs et al. (2006b); [38] Carter et al. (2014a); [39] Isogai et al. (2019); [40] Woudt et al. (2005); [41] Roelofs et al. (2005); [42] Kupfer et al. (2013); [43] Levitan et al. (2015); [44] Rau et al. (2010); [45] vsnet-alert 18124; [46] Hardy et al. (2017); [47] Prieto et al. (2014a); [48] Roelofs et al. (2009); [49] Roelofs (2007); [50] Marsh (1999); [51] Nather et al. (1981); [52] Warner (1972); [53] Woudt et al. (2013); [54] Campbell et al. (2015); [55] Carter et al. (2014b); [56] Marsh et al. (2017); [57] Anderson et al. (2008); [58] Roelofs et al. (2007d); [59] Ruiz et al. (2001); [60] Kupfer (priv. comm.); [61] Breedt et al. (in prep); [62] Breedt et al. (2014); [63] Kong et al. (2018); [64] Carter et al. (2013); [65] Wagner et al. (2014).						
*	Classification is uncertain (Levitan et al., 2013)					
†	Levitan et al. (2015) cite Augustejn, priv. comm. for the orbital period					
‡	An orbital period was predicted by Levitan et al. (2015)					
§	The orbital periods of these systems were presented by Motsoaledi at a conference in 2016, but are not published.					
¶	The orbital period reported by Carter et al. (2014b) is uncertain, as discussed in Chapter 6.					

Chapter 2

Methodology

There are few natural philosophers as frustrated as astronomers in a fog.

— Philip Pullman, *The Subtle Knife* (1997)

The scientific techniques used in this thesis broadly fall into two classes: those related to photometry (applied in Chapters 3 and 5) and those related to spectroscopy (applied in Chapters 4, 5 and 6). In this chapter I will give a general overview of the instruments, reductions and techniques used, while more specific applications are detailed in the relevant following chapters.

2.1 High-Speed Cameras

Much of the photometric data in this thesis was collected using a family of high-speed photometers: ULTRACAM, ULTRASPEC and CHIMERA. These instruments are designed for the purpose of observing short period variables such as compact binaries.

ULTRACAM, the first built of these instruments, saw first light in 2002 (Dhillon et al., 2007). It is a triple-band photometer, which uses a series of dichroics to separate incoming light by wavelength and direct the light towards three charge-coupled devices (CCDs). The light path of incoming light in ULTRACAM is shown schematically in Figure 2.1. Before reaching the CCDs, the light is passed through one filter in each arm of the camera. The dichroics are designed such that Sloan $u'g'r'i'z'$ filters can be used with minimal additional loss of light (see Figure 2.2), with the standard deployment being u' in the blue arm, g' in the green arm, and




Figure 2.1: Ray diagram through ULTRACAM, showing how light is separated into the three cameras. Note that the ray paths are coloured according to their position in the original image (at the focal plane), not according to wavelength. This figure is reproduced from Figure 2 of Dhillon et al. (2007).

one of r' , i' or z' in the red arm.

In typical photometer designs, the CCD is exposed to light for some amount of time (the ‘exposure time’), after which a shutter is closed and the accumulated charge is converted to counts, which also takes some amount of time (the ‘readout time’). The cadence that can be achieved under this design is limited by the readout time of the CCDs. Exposure times similar to or shorter than the readout time are sub-optimal as the instrument will then be exposing for only a fraction of the total time (the ‘duty cycle’ will be low) and much of the incoming light will be wasted.

ULTRACAM avoids this issue by the use of frame-transfer CCDs. One half of each CCD is exposed to incoming light and the other is covered. Once the desired exposure time has passed, the accumulated charge is transferred to the covered half of the chip, from which it can be read out while the next image is being exposed (see Figure 2.3). The time taken to transfer charge across the chip is much shorter than the readout time. Therefore, as long as the exposure time is equal to or longer than the readout time, the instrument can maintain a duty cycle $> 99\%$. The standard readout time for a full frame image is 6 s. By windowing the CCD and only reading out the particular regions of interest, the readout time can be reduced and the instrument can achieve cadences on the order of 100 Hz – 1 kHz.




Figure 2.2: Transmission profiles of the SDSS $u'g'r'i'z'$ filters used with ULTRACAM are shown as coloured, solid lines. Also shown are the transmission profiles of the two dichroics (dashed and dash-dotted lines) and of the antireflective coating used on ULTRACAM lenses (dotted line). The dichroics are designed such that their cutoffs align with the edges of the filter profiles, and they have throughputs of 90–95%. This figure is reproduced from Figure 1 of Dhillon et al. (2007).




Figure 2.3: Diagram of the operation of the frame-transfer CCDs used in ULTRA-CAM. Images are exposed in the ‘image area’, then transferred to the covered ‘storage area’ to be read out while the next image is exposed. Two readout channels are used to increase readout speed. If desired, the image can be windowed as shown here, in which case only the windowed areas will be read out, decreasing the necessary readout time. This figure is reproduced from Figure 5 of Dhillon et al. (2007).

ULTRACAM has operated as a visitor instrument on various telescopes. For the observations presented in Chapter 3, it was mounted on the 4.5 m William Herschel Telescope (WHT) at the Roque de los Muchachos Observatory, La Palma. For those presented in Chapter 5, it was mounted on the 3.5 m New Technology Telescope (NTT) at La Silla Observatory, Chile. It has also been mounted on the Gran Telescopio Canarias (GTC) on La Palma.

ULTRASPEC was originally designed as a high-speed spectrograph. Following a redesign, it now functions as a single-beam, high-speed photometer (Dhillon et al., 2014). Like ULTRACAM, ULTRASPEC uses frame-transfer CCDs to maximise its duty cycle. The instrument is designed with electron-multiplying CCDs (EMCCDs), which offer a ‘normal’ mode or an ‘avalanche’ mode. The latter can be used to increase S/N for faint targets which are dominated by readout noise, but in practise is rarely used due to the limited region of parameter space in which it is advantageous. While ULTRASPEC can be used with any filter, the most commonly used filters are Sloan $u'g'r'i'z'$ filters or the custom filter *KG5*, a broad filter designed to maximise throughput for timing targets (Hardy et al., 2017). Since 2013, ULTRASPEC has been permanently mounted on the 2.4 m Thai National Telescope (TNT) on Doi Inthanon, Thailand.

A third high-speed photometer, CHIMERA, was also used for observations presented in Chapter 3. CHIMERA has a similar design to ULTRACAM, with incoming light divided into two beams (r' and g'), and frame-transfer CCDs used to minimise dead time (Harding et al., 2016). CHIMERA is permanently mounted on the 4.2 m Hale Telescope at the Palomar Observatory, USA.

A fourth instrument, HiPERCAM, is worth noting, although it was not directly used for observations in this thesis. HiPERCAM follows the same design principles as ULTRACAM, but with several innovations. It is a five-beam photometer, allowing for $u'g'r'i'z'$ bands to be observed simultaneously. HiPERCAM saw first light in 2018, and is currently mounted as a visitor instrument on the GTC.

2.2 Photometric Reduction

The ULTRACAM, ULTRASPEC, and CHIMERA photometric data presented in this thesis were reduced using the ULTRACAM reduction pipeline. This pipeline allows for the reduction steps explained below: bias and dark subtraction, flat correction, and aperture photometry with sky subtraction.

In imaging cameras such as those described here, light from a target field is directed and focussed onto a CCD. Photons incident on a pixel excite electrons into




Figure 2.4: ULTRACAM reduction images used for observations of Gaia14aae on 2015 May 23. *Top left*: Bias frame. The left and right CCD readouts have different bias levels. *Top right*: Flat field image. Several variations in throughput can be seen, notably vignetting around the edges of the image and dust on the chip. *Bottom left*: Dark frame. Low-level pixel-to-pixel variation results from readout noise. On top of this, a number of hot pixels can be seen. Pickup noise in the right-hand CCD also becomes pronounced. *Bottom right*: Science image with apertures overplotted. The target (Gaia14aae) is in aperture 1; other apertures are possible comparison stars. Aperture 2 (green) is selected for measurement of FWHM during reduction. Sky masks (purple) have been applied to several stars.

the valence band of the semiconductor. After the desired exposure time has passed, excited electrons are moved into a readout device, which for each pixel measures a voltage designed to be proportional to the number of electrons. The voltage is digitised into Analogue Digital Units (ADU), often referred to as counts (Howell, 2006). The ratio between the number of excited electrons and this voltage is termed the gain, G , measured in e^-/ADU . Once a digital count has been measured for each pixel, the image can be reconstructed. However, several steps remain in the process of converting a series of images into usable, reduced data. Examples of the types of calibration data discussed below are shown in Figure 2.4.

2.2.1 Bias and Dark Subtraction

Bias subtraction is necessary to correct for the bias voltage, which is a voltage applied to the readout chip such that some number of counts (generally $\gtrsim 1000$ ADU) are measured even with an exposure time of zero. In order to characterise this background level of counts, a series of bias frames (essentially images with zero-second exposure times) are recorded. Typically tens of bias frames are taken. These frames are then combined using a pixel-by-pixel median process, in order to produce a master bias image which describes the bias level, including any pixel-scale variation. Science images can then be corrected for the bias level by subtracting a corresponding master bias image from each science image.

When observing with ULTRACAM, bias frames are generally taken in the afternoon before the night begins. Each bias frame is taken in the dark (with the telescope dome and mirror covers closed, the focal plane slide blocking light from entering the camera, and all lights off inside the dome). As the bias level can drift over time, bias frames are generally taken on the day before or after the night on which the data were recorded. If any windowing was done on the scientific images, the bias frames must also be windowed as this can change the bias level.

The quality of science images can be further improved by dark subtraction. In a CCD, thermal excitation will cause some electrons to become excited, even without incident photons. The background that results from this is known as the ‘dark current’. The sensitivity of the chip to thermal excitation varies between pixels, with particularly sensitive pixels (generally due to defects) known as ‘hot pixels’. In order to correct for this thermal excitation, dark frames are taken. Like bias frames, dark frames are recorded without incident light. The exposure times used for dark images are generally long (of order tens of seconds) in order to increase the dark level, as dark current can be measured in counts/second. After a bias has been subtracted from the dark image, it can be scaled by the ratio between its

exposure time and the exposure time of the science image, before being subtracted from the science image.

With ULTRACAM, dark images are generally taken at night during poor weather conditions. It is therefore not possible to collect dark images every observing run. However, it is generally assumed that the dark current is relatively consistent over time, and so dark images from previous observing runs can be used during reduction. Additionally, when observing science images, care is generally taken to avoid placing targets of interest on known hot pixels. Dark subtraction is not always used during ULTRACAM reduction, but was used for all data presented in this thesis.

2.2.2 Flat Field Correction

For a star of a given flux, the counts measured will vary depending on where the star falls on the CCD, as a result of regional variations in pixel sensitivity, filter throughput (for instance due to dust on the filter), and possible vignetting. To correct for this, flat field images are used. A flat field image is obtained by observing a field of uniform brightness (such as a blank patch of sky during twilight). Any spatial variations in the measured brightness are therefore assumed to be instrumental. These variations can be corrected for by dividing every science image throughout by a flat field image.

ULTRACAM flat field images are generally collected at twilight using a set exposure time. There results a series of tens of frames, with count levels ranging from 40000 to 10000 ADU as the sky becomes fainter. In order to combine these frames into one master flat field image, the frames are grouped according to brightness, with generally seven frames per group. Each frame is normalised by its mean level and a pixel-by-pixel median is taken within each group. The resulting median frames are then combined using a weighted mean, so as to favour the frames with higher count levels.

2.2.3 Aperture Photometry

Once a science image has been bias subtracted, dark subtracted, and flat field corrected, the total counts of the target star can be extracted. This is done using aperture photometry.

Stars are unresolved, but they do not appear as point sources on CCD images. Each star in the image has a point spread function (PSF), which defines the distribution of its light across the CCD. The PSF is a combination of various atmo-

spheric effects (such as seeing) and instrumental effects (such as the focus). While some PSFs are complex, most can be modelled with fairly simple functions such as a Moffat function (Moffat, 1969),

$$f(x, y) = I_0 \left[1 + \left(\frac{x^2 + y^2}{\alpha^2} \right) \right]^{-\beta}, \quad (2.1)$$

for coordinates x and y , where I_0 is the peak intensity of the PSF and α and β are seeing-dependent parameters which will increase for poor seeing. For imaging photometers with small fields of view, such as ULTRACAM, it can generally be assumed that all stars in an image have comparable PSFs. The full-width at half-maximum (FWHM) of a star's PSF describes the scale of the seeing, and can vary on a timescale of seconds.

In the ULTRACAM pipeline, one places an aperture over each star that will be extracted. Each aperture consists of three concentric rings (see Figure 2.4). The inner ring defines the region that will be considered to contain the star. The counts within this ring are integrated to produce the total counts for the target. If desired, optimal weights (Naylor, 1998) can be used. From the total counts found in this way, it is necessary to subtract counts which are contributed by the sky background. The outer two rings of each aperture define a region of empty sky surrounding the target. Should any nearby stars fall into this region, they can be masked. The pipeline determines the mean sky brightness per pixel of all pixels in this region. It then multiplies this sky brightness by the number of pixels within the inner ring in order to find the overall sky contribution to the total counts of the target, which it subtracts from the total.

The ideal aperture size is large enough to include most of the flux from the target star, but not so large as to include too much sky background. As the seeing can vary on short timescales, the ideal aperture size changes between images. The ULTRACAM pipeline allows for a variable aperture size to be used, defined in terms of the seeing FWHM. In each image, a Moffat function (Moffat, 1969) is fitted to a selected bright star in the field, allowing the FWHM to be measured. The aperture radius is then set to be a certain multiple of the FWHM; a common aperture radius is $1.7 \times \text{FWHM}$. The fitting process also allows for small changes in the positions of the stars, which can result from telescope drift, seeing, or atmospheric diffraction.

Atmospheric effects, such as extinction or passing clouds, can cause spurious changes in the apparent flux from a star. In order to correct for these changes in transparency, a comparison star is used. This is a non-variable star in the same field as the target, observed simultaneously. Rather than directly considering the total

counts from the target, one considers the ratio of target counts over comparison star counts.

The final step in the reduction process is flux calibration of the target. If the magnitude of the comparison star is known, the magnitude of the target in each image can be calculated as

$$m_1 - m_2 = -2.5 \log_{10} (f_1/f_2), \quad (2.2)$$

where f and m are the counts measured from and the apparent magnitude of each star. If the magnitude of the comparison star is not known, a standard star that was observed separately must instead be used. Calibration in this case requires

$$m_1 - m_2 = -2.5 \log_{10} (f_1/f_2) - k_1 A_1 + k_2 A_2, \quad (2.3)$$

where the additional terms k and A refer respectively to the extinction coefficient of the given night, and the airmass of observations. The extinction coefficient of a night can be measured by plotting $\log_{10}(f)$ against A for a non-variable star.

2.2.4 Sources of Noise in Photometric Data

When reducing photometry for a single star, there are generally three important sources of noise: *shot noise* from the star, *sky background noise*, and *readout noise*.

The *shot noise* from the star is the Poisson noise resulting from counting photons from the star itself. For an image with N counts from the star, the shot noise is $\sqrt{N/G}$, where G is the gain of the CCD. Shot noise therefore decreases as a fraction of the total counts if the exposure time is increased.

Readout noise is a form of noise added when a pixel is read out from the CCD. It can be approximated by the standard deviation of counts in each pixel of a bias frame. Readout noise is an inevitable product of the readout process, though additional readout noise can be induced by electrical interference, for instance if there is poor shielding from nearby equipment. Instruments generally have slow and fast readout modes, where the slower modes add less readout noise. For a readout noise per pixel R , the total readout noise added to a star is $\sqrt{n}R$, where n is the total number of pixels in the aperture. For this reason, binning pixels together before readout (and thereby decreasing n) can increase the signal-to-noise (S/N) of the output data.

Sky background noise is another form of shot noise, resulting not from the star but from the counts in the sky background. This noise remains even after the sky background level is subtracted. For S sky counts per pixel, the uncertainty




Figure 2.5: Example of a raw spectrum image. The spectrum of the target star here runs left-to-right, but note that it is not perfectly aligned with the CCD pixels. Perpendicular lines are sky emission lines. Several bright spots result from cosmic rays. This is a 30 minute spectrum taken using the IDS instrument on the Isaac Newton Telescope (INT) on 2019 January 15, of the white dwarf WDJ092022.20-131531.59.

contributed by the sky background to the star goes as $\sqrt{nS/G}$. Poor seeing or a brighter sky background can therefore limit observations of faint targets. Unlike with readout noise, binning data does not decrease sky noise, as the decrease in the number of pixels is countered by an increase in sky counts per pixel.

The total uncertainty on the flux of the star measured in each image goes as

$$\sigma^2 = nR^2 + \frac{nS}{G} + \frac{N}{G}. \quad (2.4)$$

When a comparison star is being used, both stars have uncertainties of this form. It is therefore preferable to choose a bright comparison star if possible.

2.3 Spectroscopic Reduction

Spectroscopy is a method for measuring the spectrum of a light source, or its brightness as a function of wavelength. Many of the principles used in photometric reduction also apply in spectroscopic reduction, but the process of spectroscopic reduction involves several additional steps which I will describe below. As the majority of the spectra presented in this thesis were observed using long slit spectroscopy, I will

focus on this method. Note that some of the data presented in Chapter 6 were collected using XSHOOTER, an Echelle spectrograph, and that the reduction of Echelle spectroscopy requires still further steps.

In long slit spectroscopy, a slit is placed in the light path of the target star, in the focal plane, following the secondary mirror but before the detector. The purpose of this slit is to take a slice through what would otherwise be a two-dimensional image, allowing only one spatial dimension in the final image (and, in most cases, blocking any light from stars that are not of interest). The width of the slit can be chosen according to the observer's priorities: a slit narrower than the seeing will produce better spectral resolution, but at the expense of blocking some fraction of the target's light. Once light has passed through the slit, it is collimated and directed onto a dispersive element. This can be a prism, a diffraction grating, or a combination of the two (a 'grism'). This dispersive element separates the light according to wavelength, producing the spectrum. The dispersed light is then directed onto a CCD detector. The detector records a two-dimensional 'image', in which one dimension is spatial and the other is dispersive, as shown in Figure 2.5.¹

The process of reducing this two-dimensional image into a one-dimensional spectrum has numerous steps. For all long slit spectroscopy presented in this thesis, this reduction was carried out using STARLINK and its packages FIGARO, KAPPA and PAMELA. Wavelength- and flux-calibration were carried out using MOLLY.

2.3.1 Bias Subtraction

As with photometric reduction, all images contain a bias level of counts which must be subtracted at the start of the reduction process. This is done via essentially the same process as with photometry, using a series of bias frames taken during the day preceding observations. The frames are combined using a pixel-by-pixel median to produce a master bias image, which is subtracted on a pixel-by-pixel scale from all other science and calibration images.

2.3.2 Flat Correction

As with photometers, the pixels in a spectrograph are not all of equal sensitivity, and these variations in sensitivity must be corrected for with a flat field image. For this, a target with a smooth spectral response is needed, for which a tungsten lamp is generally used. Several tens of spectra are taken of the tungsten lamp, and the images combined into a master flat.

¹Note that spectra are more usually oriented with the dispersive direction vertically; I present this spectrum rotated by 90° to better fit on the page.

The raw flat spectra contain a combination of the pixel-to-pixel variations, the instrumental response as a function of wavelength (including any vignetting), and the spectrum of the lamp. The latter two must be removed before the flat can be used. This is done by collapsing the image to a one-dimensional spectrum and fitting with a low-order polynomial or spline. The master flat image is then divided by the resulting fit, producing an image that is mostly constant but for the pixel-to-pixel variations. Science images can then be divided by this master flat in order to correct for pixel sensitivity.

2.3.3 Spectrum Tracing

Once bias- and flat-corrections have been performed, one can begin to extract the spectrum itself. The first step in extracting the spectrum is to identify the path of the spectrum on the CCD. While one might consider the x -direction of the image to be spatial and the y -direction to be dispersive, the spectrum does not often fall entirely along the columns of the chip, and some amount of curvature is common. This is partly due to mis-alignment between the CCD and the grism, and partly due to a wavelength dependence in the spatial position of the target star (resulting from wavelength-dependent atmospheric dispersion). It is therefore necessary to fit a polynomial or spline through the path of the spectrum on the CCD, in order to characterise the region that will be extracted.

2.3.4 Sky Subtraction

In photometry, the sky background behind a target star could be treated as a constant value, making subtraction simple. In spectroscopy, sky subtraction requires more work, owing to the spectrum of the sky itself. As can be seen in Figure 2.5, the spectrum of a target star is superimposed on a series of emission lines resulting from the Earth's atmosphere, which must be removed before the extraction can be completed.

In order to remove these sky lines, the user defines a region of sky on either side of the target spectrum which is free from any other stars that may have fallen on the slit. The sky regions will generally run parallel with the trace of the target spectrum, and have a width of 50-100 pixels. Once these regions have been defined, a polynomial can be fitted to each CCD row in the sky regions. These polynomials are used to interpolate the sky background behind the target spectrum, which can then be subtracted.

2.3.5 Spectrum Extraction

Once the region of interest surrounding the target star has been defined and the sky background has been subtracted, one can go about converting the two-dimensional data into a one-dimensional spectrum. The simplest method to do so is, for every row of the CCD, to sum the flux present in all pixels in the row (this is known as ‘normal extraction’). However, for faint targets this is rarely the ideal method, as it adds a significant amount of noise from pixels in which there is little flux from the target.

The extraction can be improved by weighting the contributions according to the amount of flux present from the target. It was shown by Horne (1986) and Marsh (1989) that, for such extraction, the optimal weight for pixel i is given by

$$W_i = \frac{P_i/V_i}{\sum_i P_i^2/V_i} \quad (2.5)$$

where P_i is the fraction of the total flux of the row that is contained in pixel i , and V_i is the variance on pixel i . Extraction using these weights is known as ‘optimal extraction’. In order to determine the optimal weights, it is necessary to predict the values of P_i . Horne (1986) suggested a method of fitting polynomials along the columns of the CCD, in order to measure P for each pixel. Marsh (1989) expanded the method by instead fitting polynomials along the path of the spectrum (parallel to the measured trace of the spectrum), allowing for spectra which are significantly curved or tilted. Once the optimal weights are known, a weighted sum of flux can be performed along each row, producing the one-dimensional spectrum.

During this process, several outlying pixel values will be detected. Some of these will result from bad pixels; others from charged particles (such particles are typically referred to as cosmic rays, though the majority in fact originate from radioactive atoms in the surrounding instrumentation; Howell, 2006). If spectra were obtained with long exposure times, the number of cosmic rays may be particularly bad. During extraction, it is usual to ‘clip’ these outlying pixels. Removing values that fall more than 3.2 – 3.5σ from the polynomial fits is common.

2.3.6 Wavelength and Flux Calibration

With an extracted one-dimensional spectrum, all that remains is to calibrate the spectrum in terms of wavelength and flux.

Wavelength calibration is performed using arc lamps. An arc lamp produces emission lines at known wavelengths, dependent on the elements in the lamp

(CuNe+CuAr is a common combination). A spectrum taken of the arc lamp allows one to associate the measured pixel positions of these lines with their known wavelengths. A polynomial fit to these measurements allows one to define a conversion from a pixel scale to a wavelength scale, which can be applied to all spectra. Generally a reasonably low order polynomial (\sim four terms) is used.

For work that is particularly sensitive to wavelength, such as measuring velocities, it is common to bracket science observations with arc spectra observed before and after. Interpolating between the calibrations from these arcs allows one to compensate for small drifts in the wavelength calibration, for instance due to instrument flexure. Alternatively, sky emission lines can be used for this purpose. Often one fixes the higher order terms of the polynomial using long-exposure ‘master arc’ spectra taken during the day, and allows the lower order polynomial terms to be adjusted by night-time arcs or sky lines. This process allows wavelength calibration that is accurate on the level of ~ 0.1 pixel.

Flux calibration is done by observing a star with a known spectrum (a ‘spectroscopic standard star’), generally on the same night if possible. By comparing the extracted spectrum of the standard star with its known spectrum, one can produce a wavelength-dependent calibration curve. This calibration can then be applied to the science spectra observed on the night.

Optionally, the flux standard star can also be used to remove telluric features – features that appear in a spectrum due to absorption of the light as it passes through the Earth’s atmosphere. To do this, the spectrum of the standard star is continuum normalised to one, and all wavelengths outside of the telluric ranges are fixed to exactly one. Any noisy pixels which exceed one can also be clipped (as telluric features should be purely absorption). Science spectra can then be divided by the resulting spectrum, scaled to account for any change in airmass between the target and the standard star.

2.4 Light Curve Modelling

In Chapter 3, I describe the fitting of a model eclipse lightcurve to data using a Markov Chain Monte Carlo (MCMC) method. A detailed description of the specific application will be given in that chapter. Here, I will give a more general overview of MCMC methods.

MCMC is a numerical technique used to explore and draw samples from a distribution, where the distribution can be arbitrarily complicated. As the name implies, MCMC methods implement a chain wherein each step of iteration depends

on the previous step (a Markov chain). The chain evolves towards a steady state, which is reached when the chain is distributed according to the target distribution. The steps of the chain once it has converged (or ‘burnt in’) can then be used to draw samples from that target distribution (Cowles, 2013).

The MCMC performed in this thesis is implemented using the PYTHON package EMCEE (Foreman-Mackey et al., 2013), which uses an affine-invariant ensemble sampling algorithm (Goodman and Weare, 2010). The parameter space is explored by a cloud of ‘walkers’. At each iteration, new positions for each walker are proposed according to the positions of the other walkers. The newly proposed position has its likelihood evaluated according to a predefined function (in this case, the likelihood of a particular model given the data). If the likelihood is better than the previous position, the change is accepted. If the likelihood is worse, the probability that the change is accepted or rejected is set by the ratio of the likelihood of the new position to that of the old. In this way, the walkers explore the parameter space and arrive at the target distribution (the likelihood distribution of possible models given the data).

The advantages of ensemble sampler methods over other MCMC methods are that the walkers allow it to converge on the target distribution in fewer steps, and that few algorithm parameters are required (only the number of steps, the number of walkers, and the starting distribution of the walkers – the latter has little effect on the final results, only on the burn-in time). For comparison, a traditional Metropolis-Hastings algorithm generates jumps from a fixed covariance matrix that the user sets before beginning the run.

Once the walkers have converged, the distribution can be characterised by the properties of a sample taken from all walkers and all steps after the burn-in phase. Various ‘nuisance parameters’, which are necessary for the model but not of interest, can be ‘marginalised’ or integrated over by effectively ignoring these parameters. The marginal distribution of a single parameter includes implicitly the contributions in variance from all other parameters. For the parameter best-fit values I quote in Chapter 3, I take the median value of all chains and steps after the burn-in phase, with the 68% confidence interval taken as the range between the 16th and 84th percentiles.

2.5 Doppler Maps

In Chapters 5 and 6, phase-resolved spectra of a target star are used to produce Doppler tomograms. The method of Doppler tomography was developed by Marsh




Figure 2.6: Diagram of the reconstruction of a Doppler map based on the profile of a spectral feature observed at different phases. The map in this case presents an accretion disc including one bright spot. This figure is reproduced from Figure 2 of Marsh and Horne (1988).




Figure 2.7: Example of a fitting process to a two-pixel image. As before, p_i refers to the pixel values of a candidate map, and q_i refers to the corresponding pixels in the default image. Contours of constant χ^2 are shown as solid lines, and constant S are shown as dashed lines (centred on $p_i = q_i$ as expected). The trajectory of solutions which minimise Q^2 is plotted, spanning the range of α values from $\alpha = 0$ (minimum χ^2) to $\alpha = \infty$ (maximum S). Setting a target χ^2 limits the process to solutions contained within the corresponding χ^2 contour, and thereafter maximising S selects the solution which lies at the edge of that contour and along the plotted trajectory of Q^2 values. This figure is reproduced from Figure 2 of Horne (1994).

and Horne (1988) as a method to convert from spectroscopic data into a map of the emission sites within a rotating reference frame. The method has been used in AM CVn systems (Roelofs et al., 2005, 2006b; Kupfer et al., 2013, 2016), as well as in CVs (Marsh et al., 1990; Neustroev et al., 2016; Rutkowski et al., 2016; Hill et al., 2017), X-ray binary systems (Wang et al., 2018), and dust discs (Manser et al., 2016).

For a rotating source, such as an accretion disc, and a given spectral line, emission may be spread evenly throughout the body or it may be localised in certain regions. Emission from any given point on the disc will produce an ‘S-wave’ feature in phase-resolved spectroscopic data. The angular velocity and the azimuthal position of the point determine, respectively, the amplitude of the S-wave and its phase. Therefore, given spectra of the feature of interest with sufficient phase coverage, it is possible to reconstruct a map of the emission intensity of the rotating source, as shown in Figure 2.6. Such maps are generally constructed and used in velocity space, owing to the non-unique nature of conversions from velocity to spatial dimensions. All motions are assumed to be in the plane of the system, with no vertical motion being considered.

The typical implementation iteratively constructs candidate maps of emission in velocity space. Each map is used to calculate expected spectral line profiles, which are compared to the observed spectra using a χ^2 statistic. The map can then be adjusted so as to minimise χ^2 . Any number of maps can be produced to describe one set of data, and so a second statistic is required: the entropy of the map. In this case entropy is defined as

$$S = - \sum_i p_i \log \frac{p_i}{q_i}, \quad \text{where} \quad p_i = \frac{I_i}{\sum_j I_j} \quad \text{and} \quad q_i = \frac{J_i}{\sum_j J_j}. \quad (2.6)$$

In the above, I is the candidate image being tested, J is a predefined default image, and i and j refer to individual points in the two images. This entropy will be maximised for images that are equal to the predefined default image. In the case of a uniform default image, this is equivalent to the standard definition of entropy.

The entropy parameter allows one to select against overly complicated maps, avoiding over-fitting the data and reducing the number of artefacts in the selected image. The fitting procedure minimises the parameter

$$Q^2 = \chi^2 - \alpha S \quad (2.7)$$

where α is a constant that selects the desired balance between χ^2 and S , with

$\alpha = \infty$ being the solution that maximises entropy (the default image) and $\alpha = 0$ being the solution that minimises χ^2 (Horne, 1994). In practice the parameter α is not used directly by the user. Instead the user sets a target value of χ^2 . The first iterates until the desired χ^2 has been achieved, then iteratively increases S until a maximum is believed to have been found. The parameter α is implicitly varied during this process, and the final value of α (along with the chosen map) is that which has the maximum value of entropy while achieving the target value of χ^2 (see Figure 2.7).

It is also possible to choose a target value of entropy and minimise χ^2 for that value, which is done for certain situations: for bootstrapping Doppler maps in order to determine the significance of certain results, and for attempting to determine periods via the Doppler maps.

Chapter 3

An Eclipsing AM CVn that Challenges Formation Models

*You know your path in the sky, you hold your course and your aim,
And each in your season returns and returns, and is always the same.*

— “Stars” from *Les Miserables* (English lyrics by Herbert Kretzmer, 1980)

3.1 Introduction

In Section 1.4 I described the three proposed evolutionary channels for AM CVn binaries. Population estimates for these formation channels are poorly constrained by data due to the difficulty in distinguishing the products of these channels from one another. Products of the white dwarf and helium-star channels can be distinguished by the nature of their donor stars, as these channels produce donor stars with different levels of degeneracy which hence occupy different regions of a mass-radius diagram (Deloye et al., 2007; Yungelson, 2008). Due to the faintness of the donor star, it cannot be observed directly.

In hydrogen CVs, high-speed photometry of eclipsing systems has yielded the most precise measurements of component masses and radii while relying only on assumptions about the geometry of the system. (eg. Savoury et al., 2011; Littlefair et al., 2014; McAllister et al., 2015). Non-magnetic CVs have a geometry consisting of several key components: the donor and accretor stars, an accretion disc around




Figure 3.1: Long-term lightcurve of Gaia14aae, showing data from a range of surveys as detailed in the figure legend. The outburst of the system was first detected by ASAS-SN around MJD=56820, followed by one or two echo outbursts detected by ASAS-SN, PanSTARRS, and *Gaia*. I have included ASAS-SN upper limits at times that fall between outbursts (in the inset) and during the gap between *Gaia* and ZTF observations. Dashed vertical lines show the times of photometric observations presented in this chapter, and dotted lines the times of spectroscopic observations presented in the next chapter.

the accretor, a stream of matter passing from the donor star to the accretor star, and a ‘bright spot’ (sometimes referred to as a ‘hot spot’) on the edge of the accretion disc at the point of intersection with the infalling matter stream. During eclipse these components add characteristic features to the lightcurve which can be used to constrain the properties of the system (Cook and Warner, 1984; Wood et al., 1986).

There are only three eclipsing AM CVn binaries that have been discovered to date: PTF1J 1919+4815, in which just the edge of the disc and bright spot are eclipsed (Levitan et al., 2014); YZ LMi (also known as SDSS J092638.71+362402.4), in which the white dwarf is partially eclipsed (Anderson et al., 2005; Copperwheat et al., 2011a); and Gaia14aae (also known as ASASSN-14cn), the only known AM CVn system in which the white dwarf is fully eclipsed (Campbell et al., 2015). Due to their eclipsing nature, the latter two of these systems are ideal targets for parameter studies. Copperwheat et al. (2011a) used eclipse fitting with high time-resolution photometry to measure the donor mass of YZ LMi. This chapter does the same for Gaia14aae.

Gaia14aae has a g' magnitude of 18.3–18.6 outside of eclipse and an orbital period of 49.71 minutes, putting it at the long-period end of the AM CVn binary distribution. Gaia14aae experienced two outbursts in June and August 2014, but no outbursts have been recorded between then and the time of writing (June 2019). Following the discovery of Gaia14aae, Campbell et al. (2015) used time-series photometry to constrain the properties of the system. They were not able to measure the mass ratio, but were able to set a lower limit of $q > 0.019$, and found corresponding minimum masses for the primary and secondary of 0.78 and 0.015 M_{\odot} respectively.

This chapter presents follow-up photometry and analysis of Gaia14aae. In Section 3.2 I will describe how the data were taken and reduced. Section 3.3 will present the data. Section 3.4 will describe the process used to fit models to the eclipses and present the results. In Section 3.5 I will compare these results to modelled evolutionary tracks.

3.2 Observations

The observations presented in this chapter were taken using ULTRACAM, ULTRASPEC and CHIMERA (described in Chapter 2). 53 eclipses were observed, each in 1–3 colour bands, spanning a time period of 25 months. A summary of observations is presented in Table 3.1.

The images from all three instruments were reduced using the ULTRACAM

Table 3.1: Summary of the observations carried out for this chapter.

Observatory	Date at start (UT)	Filters	Exposure time (s) *	Num. Eclipses
WHT	2015 01 14	$u'g'i'$	3 (9)	3
+ ULTRACAM	2015 01 15	$u'g'r'$	3 (9)	4
	2015 01 16	$u'g'r'$	3 (9)	5
	2015 01 17	$u'g'r'$	3 (9)	3
	2015 05 23	$u'g'r'$	2.5 (7.5)	6
	2015 06 22	$u'g'r'$	3 (9)	4
TNT	2016 03 12	$KG5$	8	1
+ ULTRASPEC	2016 03 13	$KG5$	5	2
	2016 03 14	$KG5$	5	3
	2016 03 15	g'	5	2
	2017 02 21	$KG5$	4	1
Hale	2016 08 06	$g'r'$	4	7
+ CHIMERA	2016 08 07	$g'r'$	4	6
	2016 08 08	$g'r'$	4	6

* Brackets denote the u' exposure time, which was increased by a factor of 3 to compensate for the lower sensitivity in that band.

reduction pipeline, described in Dhillon et al. (2007). Each of the images were bias-subtracted and divided by twilight flat fields. The ULTRACAM and ULTRASPEC data were also dark-subtracted. The flux of Gaia14aae was then extracted using aperture photometry, dividing the flux in each frame by a comparison star (J2000 coordinates 16:11:30.53, +63:09:25.8) to remove any atmospheric transparency variations.

The $u'g'r'i'$ fluxes were then calibrated using the comparison star, the magnitude of which is available from SDSS ($m_u = 17.52$, $m_g = 15.45$, $m_r = 14.62$, $m_i = 14.34$). I ensured that this star did not show variability and tested this calibration using several nearby comparison stars, each of which resulted in consistent calibrations. Although the $KG5$ band can be flux-calibrated (see the Appendix of Hardy et al., 2017), I did not calibrate these data because they are used for timing purposes only.

3.3 Photometry

The complete set of photometry is shown in Figures 3.2-3.4, and Figure 3.5 shows an example set of individual eclipses. Figure 3.6 shows r' - and g' -band phase-folded eclipses. Figure 3.7 shows r' - and g' -band phase-folded data covering the two full orbits of the binary.




Figure 3.2: ULTRACAM photometry taken on the 4.2m WHT in 2015. Data shown were taken in u' (blue), g' (green), r' (red), and i' (magenta). The i' , g' , and u' data have been offset in the y-direction by -0.07 mJy, 0.07 mJy, and 0.15 mJy, respectively. No offset has been applied to the r' data. Data with error bars more than 3.5 times the mean (due to cloud) have been removed for clarity.




Figure 3.3: ULTRASPEC photometry taken on the 2.4m TNT in 2016. Data were taken in $KG5$ (black) and g' (green). Note that the scaling of the $KG5$ data is arbitrary as those data are not flux-calibrated. Data with error bars more than 3.5 times the mean (due to cloud) have been removed for clarity.




Figure 3.4: The CHIMERA data taken on the 5.1m Hale telescope in August 2016. These data were taken in g' (green) and r' (red). The g' data have been offset by 0.07 mJy. Data with error bars more than 3.5 times the mean (due to cloud) have been removed for clarity.

The flux from Gaia14aae is dominated by the central white dwarf, with the accretion disc and bright spot also making measurable contributions. No verifiable contribution from the donor star is seen (there is no ellipsoidal modulation or visible secondary eclipse), consistent with the low expected temperature of the donor star. Figure 3.6 shows a decomposition of the contributions of each component to the eclipse profile. The bright spot enters eclipse just as the white dwarf emerges, weakening the white dwarf’s egress feature and causing an asymmetry in its eclipse shape. Some lightcurves (Figure 3.2) also show a periodic ‘hump’ feature around phases -0.5 to 0, which is commonly seen in CVs and which indicates that the light from the bright spot is preferentially emitted in a certain direction rather than being isotropic.

Gaia14aae shows the short-term variability known as ‘flickering’ which is commonly observed in accreting systems. This arises from the variable nature of the accretion, and is generally localised to the bright spot and the inner disc. Flickering provides a source of correlated noise that is difficult to deal with analytically, particularly as the amplitude of the variability changes throughout the cycle with the eclipse of the bright spot and inner disc (McAllister et al., 2017). In Gaia14aae, the amplitude of the flickering is similar to the depth of the bright spot eclipse,




Figure 3.5: Example lightcurves of Gaia14aae, demonstrating the variability of the eclipse shape. Each lightcurve has been offset by 0.2 mJy from the lightcurve below it.




Figure 3.6: Phase-folded, rescaled, and averaged light curves combining both ULTRACAM and CHIMERA data, in g' and r' (bottom and top respectively, black points). The lightcurves show the consecutive eclipses of the accreting white dwarf and the bright spot. I show best-fit models to each (red line) and the residuals. I also show a breakdown of the lightcurve contributions of each component: central white dwarf (blue), accretion disc (green), and bright spot (cyan). The donor star's contribution was assumed to be negligible. There is a feature following the egress of the bright spot (phase 0.1) which is not described by the model, and may be due to flickering.




Figure 3.7: Phase-folded, binned and averaged light curves combining both ULTRACAM and CHIMERA data, in g' and r' (bottom and top respectively), showing two orbital cycles.

meaning that in some cycles the bright spot eclipse is completely hidden by the flickering (see some of the examples in Figure 3.5). The amplitude of flickering in Gaia14aae varies significantly between epochs, with the system being particularly variable in some epochs (eg. March 2016, Figure 3.3) and comparatively stable in others (eg. August 2016, Figure 3.4).

Alongside this flickering, Gaia14aae shows a significant amount of variability on longer timescales. The total flux from the system decreases by approximately 30% between the earliest and latest observations, with no corresponding colour change. Nights such as 23 May 2015 (Figure 3.2) show a long-term variation in flux. The depths of both the bright spot eclipse and the disc eclipse change significantly throughout the period of observations, both being particularly prominent in January 2015 (Figure 3.2) and barely visible by August 2016 (Figure 3.4). The phase of the bright spot ingress also appears later in the 2016 data than in the 2015 data, which may imply an increase in the accretion disc radius. This is somewhat unexpected as the disc radius is usually expected to decrease with time during periods of quiescence (eg. Wood et al., 1989). As the change in phase is marginal, no attempt was made to account for this in the phase-folding discussed later. The depth of the white dwarf eclipse shows a significant decrease over the time period of observation (see Section 3.5.2 for further analysis of the white dwarf eclipse depth).

3.4 Light Curve Modelling

The lightcurves contain the necessary information to derive the donor mass of the system (Cook and Warner, 1984; Wood et al., 1986), given an assumed M – R relation for the central white dwarf. If the donor star is assumed to be Roche-lobe filling, its radius as a fraction of orbital separation, R_2/a , is a function of the mass ratio, q , of the two stars. The phase width of the white dwarf eclipse $\Delta\phi$ therefore depends only on q and the orbital inclination i . For a given $\Delta\phi$, the radius of the donor star (and hence q) can be increased by moving the system towards lower orbital inclinations (more face-on). A lower limit of q can be found from $\Delta\phi$ by assuming $i = 90^\circ$. Using this method, Campbell et al. (2015) found $q > 0.019$ for Gaia14aae.

Lifting the degeneracy between q and i requires additional information, which can be found from the contact phases of the bright spot eclipse. The path of the infalling stream of matter relative to the two stars (on which the bright spot is assumed to lie) is a function of q . The bright spot ingress and egress phases therefore provide an additional constraint on q and i that serves to disentangle them.

Once q and i are known, the radius R_1 of the accreting white dwarf can be

found from the duration of the ingress or egress of the white dwarf eclipse Δw . An assumed M – R relation (eg. Verbunt and Rappaport, 1988) for the accretor can then give M_1 , and hence M_2 can be found from q .

3.4.1 MCMC Modelling

I modelled the eclipses of Gaia14aae using the package `lcurve`. A description of the model can be found in Copperwheat et al. (2010, Appendix), but I give a brief overview here. The code models each component of the system (white dwarf, accretion disc, bright spot and donor star) as a grid of elements, and for each timestep computes which elements are occulted by the donor star. The white dwarf is assumed to be spherical, the disc to be 2D and azimuthally symmetric, and the donor star to be Roche-lobe filling. The bright spot is modelled as a 1D line of points, of whose light some fraction is emitted isotropically while the remainder is ‘beamed’ at an angle from the system, allowing for the recreation of the hump feature in lightcurves. The temperature profile of the disc is described by $T \propto R^\lambda$ for a given value of λ . Following trial runs in which κ was allowed to vary, I found that the standard value $\lambda = -1$ seemed to be preferred by the data, and the variable was held fixed at this value for my final model fits. Key variables in the model are the mass ratio q , orbital inclination i , mid-time of the white dwarf eclipse t_0 , orbital period P_{orb} , the relative temperatures of all components, the radius of the accretor, the outer radius of the disc, and the orientation, beamed fraction and angle of beaming of the bright spot. All variables are listed in Table 3.2. Limb-darkening of the accretor was described according to Gianninas et al. (2013) for a 13000K, $\log(g)=8.5$ white dwarf.

In order to find the optimal values and uncertainties of these parameters, this model was converged on the data using MCMC code implemented by the Python package `emcee` (Foreman-Mackey et al., 2013), as described in Chapter 2. Flat priors were used with the exception of the spot length parameter: the prior on this parameter was flat up to a length of $0.2a$, after which the likelihood decreased sharply, in order to avoid unphysically large values (owing to the orientation of the spot with respect to the viewer at the time of eclipse, the spot length was not well constrained).

In order to measure changes in white dwarf flux and disc flux between eclipses, each eclipse in each colour band was first fitted separately using a shorter MCMC run. The ULTRACAM data were converged on using 50 walkers and a minimum of 20000 trials. The CHIMERA data were converged on using 200 walkers and a minimum of 5000 trials. This was sufficient to measure the relative flux of each

Table 3.2: Details on variables in `lcurve` and any constraints applied during the fit to the phase-folded lightcurves.

Variable	Description
t_0	Mid-time of the white dwarf eclipse
q	Mass ratio M_2/M_1
Δi	Departure of inclination i from that expected given q and the average eclipse width $\Delta\phi = 0.037$
R_1	Radius of the white dwarf
$R_{\text{bright spot}}$	Radial distance of the spot from the centre of the white dwarf
R_{disc}	Radial distance of the disc edge from the centre of the white dwarf
$T_{\text{white dwarf}},$ $T_{\text{bright spot}},$ T_{disc}	Effective temperatures of the components – these are really scaling factors for the flux contributions and do not reflect the true temperatures of these components
Spot length	Length of the line of elements that make up the bright spot
Spot angle	Angle between the line of elements of the bright spot and a tangent to the edge of the circle
Spot fraction	Fraction of light from the bright spot which is ‘beamed’ in a certain direction, producing the observed hump in the lightcurves
Spot yaw	Beaming angle of the light from the bright spot in the orbital plane

component. Values of q or i based on individual eclipses have large error bars, and some eclipses are biased by flickering, by the coincidence of white dwarf egress with bright spot ingress, or show too faint a bright spot eclipse for q to be well constrained. However, they are sufficient to provide a warning if the fit to the phase-folded lightcurve were a long way from the true value, and to compare properties of the white dwarf and disc which are independent of the bright spot parameters.

To measure q and i , I produced phase-folded, binned and averaged lightcurves using all ULTRACAM and CHIMERA data in the g' and r' bands. The ULTRASPEC data were excluded because these data have larger error bars, they were taken during poorer observing conditions, and the system showed more than the usual amount of flickering during that run. Data in i' and u' were excluded due to the smaller number of eclipses available for both, as well as the poorer cadence and noisier data in u' .

To deal with the variability in both white dwarf eclipse depth and baseline brightness, the individual eclipses were scaled based on their brightness both in-eclipse (phase 0) and after-eclipse (phase 0.1 – 0.5) so as to reduce the variability within each bin. In the model fitting described above, q and i are constrained primarily by the ingress and egress phases of the white dwarf and bright spot, and therefore should not be affected by changes in the flux from the system. The phases of the bright spot can change if the disc radius changes; however, as only a small change in the disc radius is seen, this is likely not to be significant compared to other sources of uncertainty (particularly the scatter in white dwarf radius discussed in Section 3.5.2). These lightcurves were converged on with 200 walkers and a minimum of 10 000 trials.

From each fitted model I found a value for M_2 based on the measured values of q and R_1/a , Kepler's laws, and a white dwarf M – R relation for the central white dwarf (that of P. Eggleton as quoted in Verbunt and Rappaport, 1988). The M – R relation was scaled by a factor of 0.9985, chosen as the ratio between the Eggleton radius and the stellar model radius for a DB white dwarf of temperature and mass similar to ours (Holberg and Bergeron, 2006; Bergeron et al., 2011). This scaling factor represents both the nonzero temperature of the white dwarf and the lack of hydrogen in its atmosphere. I assumed a white dwarf temperature of 12900K, matching that measured by Campbell et al. (2015).

3.4.2 Bootstrapping

The flickering of the system and the variability of the white dwarf discussed may induce systematic errors in the lightcurve fits. These errors are not taken into account

in my fit to the phase-folded data. In order to estimate the effect of these possible systematic errors on the mass ratio, I performed a bootstrap analysis. This involved selecting with replacement 44 from the 44 ULTRACAM and CHIMERA g' eclipses and 41 from the 41 r' eclipses, then phase-folding and binning the lightcurves. The procedure was repeated 1000 times in each band. Each output lightcurve was converged on with an `lcurve` model, using a combination of simplex and Levenberg-Marquardt algorithms. The results had a mean of $q = 0.0281 \pm 0.0007$, where the uncertainty is the standard deviation of individual results. This value is smaller than either the g - or r -band MCMC values, although it is within 1σ of the r -band value and 2σ of the g -band value specified in Table 3.3.

3.4.3 Contact Phase Measurements

Several peculiarities of Gaia14aae make it particularly difficult to model by the method discussed above. Firstly, the bright spot of the system is weak when compared to the flickering of the system. Therefore, an inopportune spike of flickering during the bright spot eclipse can cause the fitting routine to misfit the bright spot feature, sending the model towards inaccurate parameter values. Secondly, the ingress of the bright spot eclipse coincides with the egress of the white dwarf eclipse, causing some amount of degeneracy in fitting these features.

I therefore explored an independent method of identifying q , in order to test the veracity of the MCMC modelling results. I carried out a method of measuring bright spot contact phases based on that described in Wood et al. (1986). From these contact phases a unique value of q can be determined. Several assumptions, similar to those underlying the MCMC fitting method, are required for this: I assume that the bright spot is a point source, that it lies on the path of the infalling matter, and that the path of the matter can be described purely ballistically.

For each eclipse, I generated a model of best fit to the lightcurve. From this model I separated the white dwarf component and subtracted it from the original data, leaving residuals that should describe just the eclipses of the bright spot and accretion disc. I then performed a numerical differentiation on these residuals. In eclipses with clear bright spot features, the bright spot ingress is clearly visible in the differentiated data as a contiguous set of negative values, and the egress as positive values. Contact points 1, 2, 3, and 4 were defined as the start and end of ingress and the start and end of egress, respectively; in other words, the point at which the gradient departed from 0 or returned to 0. Mid-ingress was defined as halfway between contact points 1 and 2, and mid-egress as halfway between contact points 3 and 4. Uncertainties on each ingress and egress were judged by eye. Any eclipses




Figure 3.8: Bright spot ingress and egress phases measured by eye from the differential of individual ULTRACAM eclipse lightcurves, using the method described in Section 3.4.3. These measurements were made for u' (blue), g' (green) and r' (red) eclipses. Also shown are ballistic stream paths in ingress/egress space for the q values stated. In this manner q can be estimated, under the assumption that the bright spot is a point source that lies on the ballistic stream path. The value $q = 0.0185$ is the minimum mass ratio possible given the phase width of the white dwarf eclipse. *Inset:* The weighted mean of these measurements, compared to the same ballistic streams.

without clear ingress or egress features were skipped.

Several caveats must be stated regarding this method (see eg. Feline et al., 2004). Firstly, it is vulnerable to the same systematic biases as the MCMC method regarding flickering and the coincident white dwarf egress. My hope is that the human eye might be better at identifying these features correctly than the automated technique, but this is far from certain. Secondly, it is subject to a human bias in that the observer looks for eclipse features where they are expected to be. I therefore use this method primarily as a check on the reasonableness of the MCMC results, rather than as an alternative method.

With the aforementioned caveats in place, the measured bright spot mid-ingress and mid-egress phases are shown in Figure 3.8. There is a significant scatter in the results, which seems to be over-estimated by the uncertainties. All measurements imply $q > 0.0185$, the minimum q predicted from the phase width of the white dwarf eclipse (see Section 3.5). The $1/\sigma$ weighted mean of these measurements is plotted in Figure 3.8 (inset), along with uncertainties propagated from the error bars on individual measurements. These mean ingress and egress values and their uncertainties correspond to a value of $q = 0.0267 \pm 0.0012$. This will be compared to the MCMC results given in Section 3.5.3.

3.5 Modelling Results and Discussion

The best-fit models to the phase-folded g' and r' lightcurves are shown in Figure 3.6, and their parameters presented in Table 3.3. I show as an example the corner-plot of key fit parameters to the folded r' data in Figure 3.9, and a corner-plot of all fit parameters to the folded r' data in Figure 3.10.

Table 3.3 also includes several observables that describe the eclipses: the phase width of the white dwarf eclipse from mid-ingress to mid-egress ($\Delta\phi$), the duration of the white dwarf ingress or egress (Δw), and the ingress and egress phases for the bright spot eclipse ($\phi_{\text{spot},i}$, $\phi_{\text{spot},e}$). These parameters are calculated from the best fit models.

The geometry of the best-fit model for the r' phase-folded data is shown in Figure 3.11. The geometry of the disc and bright spot are particularly worthy of note. I find an accretion disc radius very close to the Roche lobe radius, and larger than the theoretical tidal limit of $0.58a$ (Paczynski, 1977, but note this approximation assumes no viscosity, and begins to deviate from the true limit for $q < 0.03$). The bright spot is not located on the edge of the disc, but closer to the circularisation radius. However, this is dependent on the modelled brightness profile of the

Table 3.3: Summary of the modelling results based on the phase-folded data. The means given in the final column here are produced using a weighting of $1/\sigma^2$. Uncertainties are the formal MCMC uncertainties except where marked.

Param.	r'	g'	Mean
$\Delta\phi$	0.0373 ± 0.0005	0.0373 ± 0.0004	0.0373 ± 0.0003
Δw	0.00787 ± 0.00014	0.00800 ± 0.00011	0.00795 ± 0.00009
$\phi_{\text{spot},i}$	0.0220 ± 0.0007	0.0234 ± 0.0005	0.0229 ± 0.0004
$\phi_{\text{spot},e}$	0.0863 ± 0.0004	0.0845 ± 0.0005	0.0856 ± 0.0003
q	0.0283 ± 0.0007	0.0290 ± 0.0006	0.0287 ± 0.0020 *
i ($^\circ$)	86.40 ± 0.12	86.27 ± 0.10	86.3 ± 0.3 *
M_1 (M_\odot)	0.870 ± 0.007	0.872 ± 0.007	0.87 ± 0.02 *
M_2 (M_\odot)	0.0246 ± 0.0008	0.0253 ± 0.0007	0.0250 ± 0.0013 *
R_1/a	0.0215 ± 0.0002	0.0215 ± 0.0002	0.0215 ± 0.0006 *
R_2/a	0.1393 ± 0.0011	0.1404 ± 0.0009	0.140 ± 0.002 *
R_{disc}/a	0.570 ± 0.008	0.640 ± 0.006	0.615 ± 0.005
R_{spot}/a	0.442 ± 0.004	0.421 ± 0.005	0.434 ± 0.003
a (R_\odot)	0.430 ± 0.001	0.430 ± 0.001	0.430 ± 0.003 *

* On the starred results in the final column, error bars are inflated to include the estimated systematic uncertainty in q , as discussed in Section 3.5.3.

disc, which may not accurately represent the physical disc. The bright spot in my best fit model also appears to be significantly extended along the direction of the infalling stream.




Figure 3.9: Corner plot of key parameters in the MCMC fit to the phase-folded r' data. I also include M_2 , which is not itself a parameter of the fit but is derived from q and R_1 .




Figure 3.10: Corner plot of all parameters in the MCMC fit to the phase-folded r' data. I also include M_2 , as in Figure 3.9. For a description of these parameters, see Table 3.2.




Figure 3.11: Schematic geometry of Gaia14aae calculated from the best-fit model to the phase-folded r' data, showing the central white dwarf (white), the accretion disc (blue) and the donor star (yellow). The position of the bright spot peak luminosity is shown by the red diamond, with the red line showing the extension to half maximum of the bright spot. The path of the infalling stream is shown as a dashed line, the edge of the primary Roche lobe as a dotted line, and the circularisation radius as a dashed-dotted line.




Figure 3.12: Mid-eclipse timing measurements of Gaia14aae, showing the difference between the measured times and those predicted from the ephemeris in Equation 3.1. The bottom panel is a zoomed-in version of the top. Measurements in multiple colour bands have been combined using a weighted mean for this diagram, although these were treated as separate eclipses for the calculation of the ephemeris.

3.5.1 Ephemeris

From the mid-eclipse times for individual eclipses, I calculated an ephemeris for Gaia14aae of

$$\text{BMJD(TDB)} = 57153.6890966(4) + 0.03451957084(8)E \quad (3.1)$$

where E is the cycle number and the time of zero phase corresponds to the centre of the white dwarf eclipse. The quoted zero phase was chosen so as to minimise correlation between the zeropoint and period, and the quoted uncertainties are 1σ .

This ephemeris differs significantly from that quoted in Campbell et al. (2015). This difference arises from an 11 s discrepancy in the BFOSC data used as part of that ephemeris calculation, possibly a systematic offset in the instrument which is not built for precise timings. Due to this and the relatively low time resolution of the lightcurves used by Campbell et al. (2015), I did not include their data in my ephemeris calculation.

Including all data in individual colour bands I have 119 measured eclipse times, giving us 117 degrees of freedom. The χ^2 of this linear ephemeris is 105.2.

A linear ephemeris agrees with expectations. The period evolution of this system is expected to be dominated by angular momentum loss due to gravitational wave radiation. Using the Landau and Lifshitz (1971) formula

$$\left(\frac{\dot{J}}{J}\right) = -\frac{32}{5} \frac{G^3 M_1 M_2 (M_1 + M_2)}{c^5 a^4} \quad (3.2)$$

one would expect $\dot{J}/J = -1.6 \times 10^{-17} \text{ s}^{-1}$. Under the assumptions that mass transfer is conservative and that angular momentum loss is completely dominated by gravitational wave radiation, \dot{P}_{orb} can then be found (Deloye et al., 2007) by

$$\left(\frac{\dot{P}_{\text{orb}}}{P_{\text{orb}}}\right) = 3 \left(\frac{\dot{J}}{J}\right) \left[\frac{\xi_{R_2} - 1/3}{\xi_{R_2} + 5/3 - 2q} \right] \quad (3.3)$$

where $\xi_{R_2} = d(\log R_2)/d(\log M_2)$. Taking $\xi_{R_2} \approx -0.2$ as the approximate gradient of the M - R tracks in Figure 3.16 gives $\dot{P}_{\text{orb}} \approx 1.7 \times 10^{-6} \text{ s yr}^{-1}$. Given these data, a quadratic term added to the ephemeris in Equation 3.1 would not be detected to 3σ unless its value were $\gtrsim 9 \times 10^{-6} \text{ s yr}^{-1}$. The predicted change is therefore not detectable with the data presented here, which have a baseline of two years. As the detectability scales with t^2 , one can expect the period change to become detectable by the early 2020s.

Copperwheat et al. (2011a) found some evidence of departures from a linear ephemeris in the other eclipsing AM CVn type binary, YZ LMi, which were ascribed to systematic errors induced by flickering or the superhump period. To search for similar systematic biases in Gaia14aae, I checked for a correlation between the central times of each eclipse and the brightness of the disc as measured during that eclipse. I find no strong correlation, but there are two eclipses which have both an unusually bright disc and an unusually late eclipse time. Both are i' measurements, out of only three eclipses which were measured in i' . This is the band in which the disc is brightest. It is therefore possible that the i' data may be biased in some way by the disc eclipse. In u' , g' and r' I find no evidence for a similar correlation.

I searched for significant periods in a Lomb-Scargle periodogram (Lomb, 1976a; Scargle, 1982a) of these data outside of eclipse. Other than the orbital period and its harmonics, I found no significant periods. By comparing with injected, sinusoidal signals, I estimate that one would detect a signal of approximately 1% strength or greater. After subtracting a sinusoid of period equal to the orbital period from the data and creating another Lomb-Scargle transform from the residuals, I still do not detect any other periods. Given the white dwarf temperature of 12900K (Campbell et al., 2015), one would not expect to see DBV pulsations, though this is close to the temperature range for DAV pulsations.

3.5.2 White Dwarf Luminosity and Colour

For each eclipse, I measured the flux contribution from the white dwarf in my best-fit model and calculated the $u'-g'$ and $g'-r'$ colour indices. I calculated these separately for January and May/June 2015, as well as the $g'-r'$ colour from August 2016. These are shown in Figure 3.13. These colours are corrected for interstellar extinction according to Schlegel et al. (1998) and Schlafly and Finkbeiner (2011), using $E(B-V)=0.018$. I also show for comparison the expected colours from DB and DA white dwarf atmosphere models (Holberg and Bergeron, 2006; Kowalski and Saumon, 2006; Tremblay et al., 2011; Bergeron et al., 2011). My colours are approximately 2σ different from those expected for any DB white dwarf, although the closest model at 13000K is similar to the expected temperature of 12900K, based on its GALEX UV flux (Campbell et al., 2015). It should be noted that the white dwarf may not have a pure DB atmosphere; the accreted material may have a significant fraction of heavier elements, particularly C, N, O, and Ne, depending on the prior evolution of the donor star (Yungelson, 2008; Nelemans et al., 2010). The white dwarf lies in a region of colour space that is occupied by some DA white dwarfs, but its spectrum shows no sign of hydrogen and a DA white dwarf of this




Figure 3.13: The position of the primary white dwarf in colour-colour space, as measured in January 2015 (cyan), May/June 2015 (orange), and August 2016 (red dashed line, $g'-r'$ constraint only). These fluxes have been corrected for reddening. Grey points show the positions of these measurements prior to reddening correction. Also shown are DB model atmospheres (red dots 13000K models, black dots other temperatures, solid lines connecting DB models of constant surface gravity) and DA model atmospheres (dotted lines for models of constant surface gravity), with temperatures and surface gravities labelled (Holberg and Bergeron, 2006; Kowalski and Saumon, 2006; Tremblay et al., 2011; Bergeron et al., 2011).




Figure 3.14: Evolution of the measured white dwarf flux over time, averaged between eclipse depth measurements from individual eclipses. These fluxes have been corrected for reddening. The g' measurement from March 2016 (MJD ~ 57500) is possibly biased by the large amount of flickering during these observations. The most recent outburst for this system occurred prior to the start of observations, at MJD = 56880.




Figure 3.15: White dwarf luminosity as a function of other fit parameters across all ULTRACAM and CHIMERA r' -band eclipses. The correlation coefficient quoted is Pearson's r . Colours represent the dates on which the data were taken: January 2015 (blue), May 2015 (red), June 2015 (magenta), and August 2016 (cyan).

colour would require a higher temperature. If the white dwarf eclipse depths are indeed biased by some region of the disc as discussed below, this could be evidence of a colour dependence to that bias.

Ramsay et al. (2018) used a *Gaia* distance measurement for Gaia14aae (259 ± 11 pc), combined with the white dwarf g' flux measurement described here, to estimate a white dwarf temperature of 17000K, significantly above the Campbell et al. (2015) measurement. The disparity may be evidence for metal pollution of the white dwarf, which can decrease the UV flux compared to a pure helium atmosphere white dwarf. In this case, the more reliable temperature estimate is likely to be 17000K.

The flux of the white dwarf over time is shown in Figure 3.14. Between January 2015 and August 2016, I measured a decrease of $(26 \pm 2)\%$ in both r' and g' bands. The ratio between r' and g' stayed approximately constant.

This could be due to genuine cooling of the white dwarf, which would have been heated by its outbursts in 2014. Assuming the white dwarf is approximately described by a blackbody spectrum at 13000K, and treating the filters as top-hat functions whose value is 1 inside their FWHM and 0 elsewhere, this flux decrease would imply a temperature change of $\gtrsim 1000$ K over the 18 month time interval (beginning approximately 150 days after the most recent outburst). This change is slightly less than the cooling seen in the short-period CV WZ Sge over the period 150-700 days after its outburst (Godon et al., 2004), and less than seen in GW Lib over the period 3-4 years after its outburst (Szkody et al., 2012). Such a temperature change would not produce a significant colour change, which is consistent with these observations.

It should be noted that the temperatures discussed above are higher than the ≈ 11000 K temperature prediction of Bildsten et al. (2006), which was based on compressional heating of the white dwarf by the accreted matter. Bildsten et al. (2006) predict that the temperature of the central white dwarf is driven by accretion-induced heating until the system reaches an orbital period of around 40 minutes. For periods longer than this the decreased mass accretion rate means that the temperature decouples from accretion-driven heating, and the white dwarf then follows standard cooling tracks. However, as will be discussed in the following sections, the donor star in Gaia14aae seems to have a higher mass and radius than was assumed by Bildsten et al. (2006), resulting in a higher mass transfer rate, which may explain this elevated temperature.

While a change in the temperature of the central white dwarf is one possible explanation of the change in eclipse depth, other possibilities must be considered.

Table 3.4: Summary of the q values found by the different methods described in Section 3.4.

Method	q
Phase-folded lightcurves MCMC	0.0287 ± 0.0005
Individual lightcurves MCMC	0.0268 ± 0.0010
Bootstrapping method	0.0281 ± 0.0007
Contact phase measurement	0.0267 ± 0.0012

This change in eclipse depth could be explained if there were a component of the system which is not included in these models and which is not resolved from the white dwarf eclipse, such that the apparent dimming of the white dwarf is in fact the dimming of this other component. Wood et al. (1986) and Spark and O’Donoghue (2015) discussed the optical visibility of the ‘boundary layer’ through which the white dwarf accretes from its disc. Though this boundary layer can theoretically be resolved from the white dwarf by the shape of the eclipse ingress and egress, doing so would require a higher S/N than is currently available for this system. If present, a bright equatorial boundary layer might bias measurements of the white dwarf radius towards smaller or larger values, and a variable boundary layer might therefore induce a correlation between the modelled white dwarf luminosity and radius. I checked for such a correlation in my best-fit models to individual eclipses, as shown in Figure 3.15. I did not find a significant correlation.

I do find a reasonably strong correlation between the temperature of the white dwarf and that of the disc. This could be evidence that some disc luminosity is not resolved from the white dwarf eclipse, biasing my measurement of white dwarf luminosity. Alternatively, there may be some causal link between the two, if both are heated by accretion or if both are still decreasing following the 2014 outburst.

The measurement of q described in the next section is constrained by the contact phases of white dwarf and bright spot eclipses. Therefore, q should not be affected by a bias in the apparent depth of the white dwarf eclipse. However, if the same mechanism were to bias the phase width of the white dwarf eclipse this would be a problem. I therefore searched for a correlation between the depth and phase width of the white dwarf eclipse as measured from my model fits, but found no evidence for a significant correlation (see Figure 3.15).

3.5.3 Mass Ratio and Donor Mass

From the phase width of the white dwarf eclipse, a lower limit on the mass ratio can be found by assuming $i = 90^\circ$. The phase width shown in Table 3.3 gives

$q_{min} = 0.0185 \pm 0.0005$. This is comparable to the lower limit found by Campbell et al. (2015), $q_{min} = 0.019$.

Thus far in this chapter I have described several methods of determining q : MCMC fitting to phase-folded lightcurves and to lightcurves of individual eclipses, fitting to a collection of bootstrapped phase-folded lightcurves, and measurement of bright spot contact phases by hand. The results found by each of these methods are summarised in Table 3.4.

The most rigorous method for calculating q is the phase-folded MCMC method, and hence I favour the result from this method. However, the scatter of the results in Table 3.4 suggests that the formal uncertainty on q quoted by the MCMC may be an underestimate. This is not unexpected, as the method does not take into account any systematic biases that may come from the variability of the source between eclipses. I therefore suggest a more conservative error bar of 0.0020, in an attempt to take this scatter into account. My canonical mass ratio is then $q = 0.0287 \pm 0.0020$.

To find the corresponding uncertainties on other stellar parameters, most importantly M_2 , I propagated this uncertainty through the derivation of the other parameters from the raw observables $\Delta\phi$ and Δw . The resulting uncertainties are shown in the marked entries in Table 3.3.

3.5.4 Comparison to Models

The evolutionary history of Gaia14aae can be explored by comparing the mass of the donor star with theoretical evolutionary tracks. From Kepler’s third law and the constraint that the radius of the donor star equals the radius of the Roche lobe comes the relation between donor mass and radius for a given orbital period (Faulkner et al., 1972)

$$P_{\text{orb}} = 101\text{s} \left(\frac{R_2}{0.01R_{\odot}} \right)^{\frac{3}{2}} \left(\frac{0.1M_{\odot}}{M_2} \right)^{\frac{1}{2}}. \quad (3.4)$$

This can alternately be expressed as a constraint on mean donor star density coming from the orbital period alone. Therefore, my measurement of M_2 combined with the orbital period of Gaia14aae ($P_{\text{orb}} = 49.71$ minutes) can locate the donor star in M – R space, allowing us to compare the donor star properties with the evolutionary tracks.

I take numerically calculated M – R evolution tracks for AMCVn binary donor stars that have been published for the white dwarf donor star (Deloye et al., 2007), helium-star donor (Yungelson, 2008), and evolved-CV (Goliasch and Nelson,




Figure 3.16: The measured donor masses and radii of Gaia14aae (red circle) and YZ LMi (blue diamond) and their uncertainties (solid lines of corresponding colours) compared to evolutionary tracks from Deloye et al. (2007, thin black lines), Yungelson (2008, dashed), Goliach and Nelson (2015, dashed-dotted), and Goliach and Nelson (priv. comm., dotted). The white dwarf and helium star donor tracks include multiple levels of degeneracy, with the most degenerate models at the bottom. The evolved-CV tracks are for initial accretor mass values of $0.6 M_\odot$ (dashed-dotted) and $0.85 M_\odot$ (dotted). They roughly show the edge of the parameter space this channel can easily explore; systems above and to the left of these tracks can form by this channel more easily than systems below and to the right. I also show the $M-R$ track for a zero-entropy donor star (thick solid line). The arrow shows the lower limit on M_2 corresponding to $q_{\min} = 0.0185$ for Gaia14aae. The diagonal uncertainties result from the strong constraints on the donor stars’ mean densities resulting (according to Equation 3.4) from their tightly constrained orbital periods.

2015) formation channels. I show these tracks in Figure 3.16.

I show white dwarf donor star models with total binary mass $M_{\text{tot}} = 0.825M_{\odot}$, similar to my measured $M_{\text{tot}} \approx 0.91M_{\odot}$ (Deloye et al., 2007). These tracks are for a range of initial donor star degeneracies, with the top tracks being the least degenerate. The highest of these denotes the approximate boundary between the least degenerate white dwarf donor tracks and the most degenerate helium-star donor tracks. The helium-star donor tracks are for a binary with initial conditions $M_{1,i} = 0.8M_{\odot}$, $M_{2,i} = 0.65M_{\odot}$ and are distinguished by the evolutionary status of the donor star, with the top tracks being the least evolved and the bottom tracks being the most evolved (Yungelson, 2008).

For the white dwarf and helium-star evolution tracks, a donor star would evolve into contact at the right-hand side of Figure 3.16 and evolve through mass transfer to a lower mass and longer orbital period (moving right-to-left in Figure 3.16). At a period of ≈ 40 minutes, both Deloye et al. (2007) and Yungelson (2008) predict that the thermal timescale of the donor star becomes shorter than the mass-loss timescale, allowing the donor star to cool in response to mass loss. The donor star therefore loses entropy and contracts towards complete degeneracy. For both evolutionary channels, donor stars with orbital periods significantly longer than this are expected to be completely degenerate.

The measured mass and radius lie outside the predicted parameter space for the white dwarf donor star evolutionary channel. Given the density constraint arising from its period, the true donor mass would have to be a factor of 2-3 times smaller to agree with the tracks shown here. Indeed, the density constraint coupled with the minimum donor mass found by Campbell et al. (2015) are sufficient to disagree with the predicted tracks, including the collapse due to cooling, without needing the mass estimate presented in this work. Irradiation of the donor star can delay the predicted collapse to longer periods. Even so, my mass estimate would lie significantly above the least degenerate white dwarf donor track. I therefore find that Gaia14aae is unlikely to have evolved through the white dwarf donor track.

The helium-star channel evolutionary tracks calculated by Yungelson (2008) end once the AM CVn systems reach periods of around 40 minutes, making it difficult to compare my measured mass and radius to these tracks. The predicted collapse towards degeneracy beyond this point (the beginning of which is just visible at the left-hand end of these tracks) would disagree with my measurements. However, if this collapse were delayed to a period of 50 minutes or longer, my measurements would appear to be close to agreement with the least degenerate of the tracks shown here. Delaying this collapse to longer periods may be possible by, for instance,

increasing the efficiency of irradiation of the donor star from the accretor (Deloye et al., 2007) or from outbursts.

I also show two evolved-CV tracks. The first, from Goliash and Nelson (2015, dashed-dotted line), has a $0.6 M_{\odot}$ accretor, $\dot{M} = 1.4 \times 10^{-11} M_{\odot} \text{yr}^{-1}$, and a hydrogen fraction of $\sim 5\%$ in the transferred gas. The second track (Goliash and Nelson, priv. comm., dotted line) has an accretor of $0.85 M_{\odot}$. These tracks are roughly representative of the edge of region of M – R space this channel explores; systems above and to the left of these tracks (including Gaia14aae) could easily have formed by this channel. However, AM CVn binaries that form by this channel are expected to retain a greater fraction of their hydrogen than AM CVn systems that form by the other two channels. Evolved-CV tracks with a smaller hydrogen fraction than $\sim 1\%$ in this region of M – R space are possible, but they can only form from progenitor CVs with a limited range of initial parameters, and require long timescales to evolve. One might therefore expect to see a number of short-period CVs with visible hydrogen and helium for each AM CVn binary that forms by this channel. At present few such CVs are known (see Augusteijn et al., 1996; Breedt et al., 2012, for examples).

In terms of its measured mass and radius, the donor star in Gaia14aae appears to fit best with tracks of the evolved-CV channel. However, the predicted hydrogen content of systems formed by these tracks is not seen in spectra of Gaia14aae (Campbell et al., 2015). Though no quantitative upper limit exists for Gaia14aae, a hydrogen abundance on the order predicted for this channel would be easily detectable: modelling of the AM CVn systems GP Com (with local thermal equilibrium models, Marsh et al., 1991) and V396 Hya (with non-local thermal equilibrium models, Nagel et al., 2009) find that any hydrogen abundances of $\gtrsim 10^{-5}$ should result in detectable Balmer emission. Similar upper limits have been found for DB white dwarfs in this temperature range (Koester and Kepler, 2015; Bergeron et al., 2011). Producing similar tracks without detectable hydrogen would require finely-tuned starting parameters that make them unlikely (Goliash & Nelson, priv. comm.). The absence of any visible hydrogen in the optical spectrum of Gaia14aae is therefore difficult to reconcile with predictions for this channel.

I conclude that Gaia14aae may have formed by either the evolved-CV channel or the helium-star donor channel. The results presented here are not in perfect agreement with either of the models as they currently stand, but one or both of these channels may have a region of parameter space that can explain the properties of Gaia14aae. The close agreement of my measured M_2 with that of the evolved-CV channel is particularly intriguing, given that this channel has generally been

considered the least probable formation channel. The clear disagreement of my result with the white dwarf donor formation channel, which has been widely considered to be the dominant channel, is also of interest in light of the prediction by Shen (2015) that double white dwarf binaries will not reach a state of stable accretion.

3.5.5 Implications For Gravitational Wave Emission

The distribution of AM CVn binary-originated gravitational wave strains detectable by the Laser Interferometer Space Antenna (LISA) depends on the distribution of both periods and stellar masses across AM CVn systems. The donor star mass of Gaia14aae presented above is greater than would be expected for a fully-degenerate donor star in a binary at this orbital period. Consequently, the gravitational wave emission will be greater.

Due to the high inclination of Gaia14aae, the h_+ component of its gravitational wave emission will dominate over the h_\times component. The amplitude of the h_+ component can be calculated (Korol et al., 2017) by

$$h_+ = \frac{2(G\mathcal{M})^{5/3}(2\pi/P_{\text{orb}})^{2/3}}{c^4 d} (1 + \cos^2 i) \quad (3.5)$$

where $\mathcal{M} = (M_1 M_2)^{3/5} (M_1 + M_2)^{-1/5}$ is the chirp mass of the system and d is its distance from the Earth. For the stellar masses shown above, I predict the strain from Gaia14aae to be

$$h_+ = 1.2 \times 10^{-20} \times \frac{1 \text{pc}}{d}. \quad (3.6)$$

For the distance of 259pc measured by Ramsay et al. (2018), this would give a strain of 4.6×10^{-23} . This is below the detection limit of LISA, unsurprisingly given the long orbital period of Gaia14aae. AM CVn systems at the short end of their period distribution are expected to be the brightest emitters of gravitational waves.

For comparison I consider a degenerate donor star with the same orbital period but a mass of $M_2 = 0.01M_\odot$. The strain for that system would be $h_+ = 1.7 \times 10^{-23}$. The gravitational wave emission of Gaia14aae is therefore a factor of 2.7 higher than would be expected in the degenerate case, and the volume of space in which the system would be detectable is increased by nearly a factor of 20. This emphasises the need to understand the nature of AM CVn donor stars in order to predict the distribution of strains LISA will detect, both as resolved sources and as unresolved background. The results from both Gaia14aae and YZ LMi suggest that non-degenerate and partially degenerate donor stars are more common among AM CVn binaries than previously believed. If this is the case, one would expect

AM CVn systems as a population to be brighter emitters of gravitational waves for a given orbital period than previously predicted.

3.6 Conclusions

Gaia14aae is the third known eclipsing AM CVn-type binary and the first in which the central white dwarf is fully eclipsed. As such, it provides an unprecedented opportunity to measure the properties of the component stars in one of these systems and thereby constrain the system’s prior evolution. The results are difficult to reconcile with existing models.

Any measurement of the properties of Gaia14aae is complicated by several unfortunate peculiarities of the system. In particular, the weakness of a key feature of the system (its ‘bright spot’) means that lightcurve fitting can be biased by the intrinsic red noise of the system. These difficulties increase the systematic uncertainty in my measurements. I have attempted to take this uncertainty into account by inflating the quoted error bars on my results.

I measured a mass ratio $q = 0.0287 \pm 0.0020$ and a donor star mass $M_2 = 0.0250 \pm 0.0013M_\odot$. Combined with the donor density constraint arising from the orbital period, this mass shows that the donor star is not degenerate and that the system did not evolve from a double degenerate binary. The system therefore must either have a non-degenerate helium star as its donor star or be descended from a hydrogen CV with an evolved donor star. In both cases there are unexplained questions: the donor star in the former is expected to have collapsed towards degeneracy before reaching this orbital period, and the latter is expected to show traces of spectroscopic hydrogen. Neither of these predictions are observed, but it may be possible to tweak the models in order to explain Gaia14aae’s evolution.

The evolutionary channel of Gaia14aae will be discussed further in Chapter 4, in which I present phase-resolved spectroscopy of the system.

Chapter 4

Line Emission From Near the Central White Dwarf of Gaia14aae

I have loved the stars too truly to be fearful of the night.

— Sarah (Sadie) Williams, *The Old Astronomer* (1868)

4.1 Introduction

The results in the previous chapter constrain the masses and radii of the component stars in Gaia14aae. Further opportunities to constrain the system present themselves which are only possible for an eclipsing binary. For instance, Ramsay et al. (2018) used the white dwarf and bright spot eclipse depths of the system (measured from the data in Chapter 3) and its *Gaia* parallax to estimate the temperature of the accreting white dwarf and the mass transfer rate of the binary. In this chapter I will discuss phase-resolved spectroscopy of Gaia14aae and further constraints on the nature of the binary that arise from these data, as well as one of the ways in which the eclipsing nature of Gaia14aae allows us to test an assumption that has been used for other AM CVn binaries.

Phase-resolved spectroscopy has been used to measure the mass ratios of a number of AM CVn systems. The spectroscopic method relies on resolving the radial

velocity of the central white dwarf as a function of orbital phase, thereby determining the dynamics of the binary. It therefore requires a spectroscopic signature of the central white dwarf (Roelofs et al., 2005; Kupfer et al., 2016).

Spectral lines of AM CVn-type binaries have orbital phase-dependent profiles arising from the presence of multiple components, which I describe in the following two paragraphs. The accretion disc contributes a broad, double-peaked line profile, as the approaching limb of the disc is blueshifted and the receding limb is redshifted. The point of contact between the stream of infalling matter and the accretion disc, a region known as the ‘bright spot’, is the origin of a spectral component whose Doppler shift varies with orbital phase (known as an ‘S-wave’). Many AM CVns show a ‘second bright spot’ spectral feature which is Doppler shifted with a similar amplitude, but offset in orbital phase by 120° from the first bright spot (Roelofs et al., 2005, 2006b; Kupfer et al., 2013, 2016). The physical cause of this feature is uncertain. Similar features have been seen in some short-period hydrogen CVs, though it is unclear whether the mechanism is the same (eg. EX Dra, Joergens et al., 2000).

Lastly, some AM CVn spectral lines show a sharp emission feature known as the ‘central spike’ (Nather et al., 1981; Marsh, 1999). The central spike is redshifted by an amount approximately consistent with a gravitational potential expected on or near the white dwarf surface. Wavelength-dependent blue-shifts are also seen in the central spike of some He I lines, resulting from forbidden transitions excited in regions of high electron density (Morales-Rueda et al., 2003). On top of these systemic velocity offsets relative to the rest wavelength, the feature undergoes low-amplitude velocity changes which appear to coincide with the phase and amplitude of the central white dwarf (Kupfer et al., 2016). The central spike has therefore been assumed to originate at the central white dwarf, and these phase-dependent Doppler shifts have regularly been used to measure AM CVn mass ratios. However, there has been no direct test of this assumption.

4.2 Observations

Phase-resolved spectroscopy of Gaia14aae was obtained using the ISIS spectrograph on the 4.2m William Herschel Telescope (WHT), and the OSIRIS spectrograph (Sánchez et al., 2012) on the 10.4m Gran Telescopio Canarias (GTC), both at the Roque de los Muchachos observatory on La Palma. The observations undertaken for this project are summarised in Table 4.1. Conditions for all observations were clear with good seeing ($0.4\text{--}1''$). All images were bias-subtracted and flat field-

Table 4.1: Observations undertaken for this work.

Date	Instrument	Exposures
2015 July 16	WHT+ISIS	31×240 s
2015 July 17	WHT+ISIS	51×240 s
2016 May 14	GTC+OSIRIS	136×70 s

corrected using tungsten lamps. Extraction and calibration were carried out with optimal weights, using the software packages PAMELA and MOLLY as described in Chapter 2.

The WHT spectra were collected using a 1" slit and exposure times of 240s. The ISIS spectrograph has two arms, allowing it to simultaneously collect data at blue and red wavelengths. In the blue arm the R300B grating was used, giving a mean dispersion of 1.7 Å/pixel across wavelengths 3500–5400 Å. In the red arm the R316R grating and a GG495 order-sorter filter were used, giving a mean dispersion of 1.8 Å/pixel across wavelengths 5400–8000 Å. The resolutions were 3.5 Å in the blue arm and 3.2 Å in the red arm, both measured using arc lines close to the centre of the spectral range.

Wavelength calibration arcs were obtained once per hour throughout the night. The red arm was calibrated against 32 arc lines using a five-term polynomial, achieving a root-mean-square (RMS) deviation from the polynomial of 0.017 Å. The blue arm was calibrated against 30 arc lines using a seven-term polynomial, achieving an RMS deviation of 0.175 Å. Due to a lack of viable arc lines at the blue end of the spectrum, the calibration for wavelengths 3500–3900 Å is unreliable. Flux calibration was carried out against the standard star BD+33 2642, observed on both nights.

The GTC spectra were collected using a 0.8" slit and exposure times of 70s. The R2500V grating was used, giving a dispersion 0.8 Å/pixel across wavelengths 4400–6000 Å. The resolution was 2.1 Å based on a fit to the arc lines, or 2.3 Å based on the 5577 Å sky emission line. Wavelength calibration was carried out against 27 arc lines using a five-term polynomial. The resulting RMS was 0.016 Å. The flux was calibrated against the standard star Feige 66, observed on the same night.

Table 4.2: Features identified from WHT and GTC spectra of Gaia14aae, along with their measured EWs. The EW of some features could not be measured, due either to a blend or insufficient signal. The quoted uncertainties are scaled to account for any scatter in the local continuum. Regardless, I note that the uncertainties are statistical only, and do not represent the false alarm probability of a particular line. I class the He I lines as emission features, but note that these features also have prominent absorption components which in some cases result in positive EWs.

Line (wavelength in Å)	WHT EW (Å)	GTC EW (Å)
<i>Emission features</i>		
He I 3888.6	-9.2 ± 0.5	–
He I 3964.7	1.45 ± 0.28	–
He I 4026.2	-4.39 ± 0.27	–
He I 4120.8/43.8	-3.25 ± 0.27	–
He I 4387.9	-2.09 ± 0.16	–
He I 4471.5 *	-13.42 ± 0.31	-7.39 ± 0.26
He II 4685.7 *	X ^a	X ^a
He I 4713.1 *	-12.37 ± 0.32	-6.30 ± 0.22
He I 4921.9	-9.72 ± 0.19	-4.61 ± 0.14
He I 5015.6/47.7	-35.8 ± 0.5	-17.77 ± 0.32
Si II 5640–5710 †	-3.19 ± 0.15	-1.30 ± 0.16
He I 5875.6 *	-46.4 ± 1.0	-33.2 ± 0.9
Si II 5957.6/78.9 *	-0.38 ± 0.11	-0.49 ± 0.16
Si II 6347.1/71.4	-15.7 ± 0.5	–
He II 6560.1	-0.82 ± 0.16	–
He I 6678.2	-42.1 ± 1.2	–
He I 7065.2	-51.4 ± 1.5	–
He I 7281.4	-26.7 ± 1.2	–
<i>Absorption features</i>		
Mg I 3829.4/32.3/38.3	4.47 ± 0.08	–
Si II 3856.0 *	X	–
Ca II 3933.7 *‡	X	–

Continued on next page

Table 4.2 – *continued*

Line (wavelength in Å)	WHT EW (Å)	GTC EW (Å)
Mg II 4481.2 *	X	X
Mg I 5167.3/72.7/83.6	0.82 ± 0.12	1.38 ± 0.11
Na I 5890.0 *	X	X
Na I 5895.2 *	X	X
Mg I 6145.1	0.53 ± 0.09	–
Ne I 6402.2 *	X	–
N I 6483.7	0.67 ± 0.08	–
Ne I 6506.5	0.58 ± 0.08	–
N I 7423.6/42.3/68.3	X	–
O I 7774.2	1.5 ± 0.5	–

– Wavelength not covered.

* Marked features are blended with their neighbours.

† Blended feature consistent with several Si II emission lines.

‡ Ca II 3968 Å was excluded due to a blend with He I 3965 Å.

^a EW for He I 4713 Å includes blended He II 4686 Å line.

X Line present but EW could not be measured, due to either a blend with another line or significant continuum noise.

4.3 Results

4.3.1 Average Spectrum

Helium lines

The average spectrum of Gaia14aae is shown in Figure 4.1. It consists of a blue continuum with a series of strong helium emission lines, and some weaker metal lines which are in a mixture of emission and absorption. The spectral lines identified and their measured equivalent widths (EWs) are listed in Table 4.2. The majority of these spectral lines originate from He I, and have a double-peaked profile characteristic of accretion disc emission, along with deep absorption cores. Also clearly visible is an emission line from He II at 4685.7 Å. Unlike the He I emission lines, He II appears single-peaked in the average spectrum. I identify this emission as the ‘central spike’ feature. There is no evidence for the central spike feature in any He I line.

A weak emission line is seen at 6560 Å which may originate from the H α line (6562.8 Å) or from a line in the He II Pickering series (6560.1 Å). As discussed later,




Figure 4.1: *Above:* The average WHT spectrum, obtained in 2015 July, created by combining 78 spectra with 240s exposure times that did not coincide with eclipses. The join between the red and blue arms of the spectrum is around 5400 Å. The positions of various identified spectral features are marked. For wavelengths < 3900 Å, the wavelength calibration is unreliable and I therefore do not identify these features. *Below:* Close-ups of selected spectral features. As can be seen, detections of N I and O I are marginal relative to the noise. Several absorption lines are still unidentified, including the feature next to the Ca II *K* line.

I believe this to be He II.

Absorption cores of the type seen in Gaia14aae are commonly seen in CV and AM CVn systems when the system is viewed edge-on relative to the orbital plane. These cores result from obscuration of the central white dwarf by an accretion disc which is optically thick at those wavelengths and optically thin elsewhere. The narrow profile of the absorption cores in Gaia14aae suggests that they are dominated by this form of absorption rather than photospheric absorption, which, when seen in helium-atmosphere white dwarfs, typically produces broader spectral lines (Bergeron et al., 2011). I discuss this further in Section 4.3.2.

Between the WHT observations in 2015 July and the GTC observations in 2016 May, the EW of He I lines decreased by a mean factor of 0.65. Gaia14aae underwent an outburst in 2014 (Campbell et al., 2015) and, at the time of these observations, was still decreasing in brightness towards quiescence (see the photometry in Chapter 3). A comparison of WHT and GTC average spectra is shown in Fig. 4.2. There is a clear decrease in both the continuum level and the line strengths by the time of the GTC spectra in 2016 May. This may be due to slit losses or, if intrinsic to the system, may be a result of the system continuing to cool following its outburst. The decrease in EW may be explained by the line strengths decreasing more significantly than the continuum level and the absorption cores. Similar changes in EW have been seen in dwarf novae as the system cools following an outburst (eg. Fig. 1 of Szkody et al., 2012).

The measured He I EWs can be compared with other AM CVns. For a sample of 23 AM CVns with emission line spectra, Carter et al. (2013, their Fig. 10) found a correlation between orbital period and the EW of the He I 5875 Å line. Based on this correlation, the 5875 Å line of Gaia14aae is weaker than expected, though within the scatter. This relative weakness is likely due to the absorption cores that result from the edge-on viewing angle of Gaia14aae. Note, however, that the partially-eclipsing AM CVn binary YZLMi (which is necessarily viewed edge-on) does not show such strong absorption cores (Anderson et al., 2005).

Metal lines

In Gaia14aae I detect metal lines from a variety of elements listed in Table 4.2, including N, O, Ne, Na, Mg, Si, and Ca. Of these, the only prominent *emission* lines are Si II lines at 6347 and 6371 Å, along with a blended Si II feature around 5640–5710 Å. Marsh et al. (1991) predicted that Si II 6347 and 6371 Å should be among the strongest metal emission lines in helium-dominated, optically thin accretion discs.

All other metal features in the spectrum of Gaia14aae are absorption lines.




Figure 4.2: Comparison of WHT and GTC average spectra for a selected wavelength range. No flux offsets or flux scaling have been applied, apart from the original flux calibration described in Section 4.2. The difference in continuum level may be genuine or may be a result of slit losses. However, there is a visible decrease in the strength of the emission lines relative to the continuum by the time of the GTC spectrum.

Metal absorption lines are relatively rare in quiescent AM CVn binaries, though they have been previously detected in SDSS J1552+3201 (Roelofs et al., 2007d), SDSS J1208+3550, and SDSS J1642+1934 (Kupfer et al., 2013). All three of these systems show Mg I absorption, and the latter two additionally show Si I/II absorption. N I and N II absorption has been seen in GP Com and SDSS 1908+3940 (Morales-Rueda et al., 2003; Kupfer et al., 2015, 2016), with GP Com also showing significant N emission lines. O and Ne have not been previously detected in absorption, but are present in emission in GP Com and V396 Hya (Kupfer et al., 2016). Similarly, Ca has previously been detected in emission in systems such as V406 Hya (Roelofs et al., 2006a), but never before in absorption. The Ca II *K* line is seen in this spectrum, while the nearby Ca II *H* line at 3968 Å would be blended with a He I line at 3965 Å and is therefore not detected. The detection of Na appears to be unprecedented in AM CVn systems. However, note that the Na absorption lines in these spectra may be residuals of the sky subtraction, as the Na emission from the sky was strong on the nights of these observations.

The detection of a wide variety of metals in absorption is unusual in an AM CVn, but not necessarily unexpected for Gaia14aae. Under the assumption that these absorption lines result from obscuration of the central white dwarf, the high inclination of Gaia14aae is likely to make such otherwise weak absorption features more prominent.

In Table 4.2 I have not identified any lines of Fe. This is something of a surprise; Fe is a common element and has been detected in a number of AM CVns (eg. Roelofs et al., 2005, 2006a, 2007c; Kupfer et al., 2013). Alongside Si II, Marsh et al. (1991) predicted that Fe II at 5169 Å should be among the strongest metal emission lines. The presence of Si II and absence of Fe II is consistent with some previously studied AM CVn systems, while others have shown both Si II and Fe II emission or no detectable metal emission (eg. Kupfer et al., 2013, 2016). Note that I cannot rule out that some spectral features may result from Fe. The absorption series in the 3500–3900 Å range (which I was unable to precisely wavelength calibrate) is most likely to be He I, but some lines may be explained by Fe. The absorption lines that I identify as Mg I 5167 and 5173 Å may in fact be Fe I 5167 and 5172 Å, although there is no Fe line corresponding to Mg I 5183 Å. The unknown feature around 3927 Å (Fig. 4.1, lower left panel) could be Fe I 3928 Å, though it would have to be blueshifted where most features are slightly redshifted, and I do not see the nearby Fe I transitions at 3923 and 3930 Å despite their similar predicted strength. I also note non-detections of Fe I at 4064, 4272, 4282, 4308 4325, 4958, 5227, and 5270 Å, and of Fe II at 6248 and 6456 Å, all of which are expected to be reasonably

strong and occur at regions of clear baseline in the spectrum.

The presence of N and O is of particular interest, as the abundances of CNO elements can give insights into the evolutionary history of the donor and hence of the binary (Yungelson, 2008; Nelemans et al., 2010). If the donor has been through a phase of helium burning, C and O will be significantly enriched relative to N. I do not detect C in Gaia14aae, but given the marginal nature of the detections of N and O this is not a significant non-detection.

Radial Velocity

I measured the average radial velocity of Gaia14aae from the He I emission lines in the WHT spectra in several different ways. First, I selected He I lines that were not blended with other lines, and avoided lines that were extremely asymmetric. This left 5 suitable He I lines, shown in Fig. 4.3. For measuring the radial velocity one is most interested in the outer edges of the line profile. The outer edges correspond to the inner edge of the accretion disc and are therefore less likely to be affected by asymmetries in the disc that may occur due to the bright spots or eccentricity of the disc. In Gaia14aae the centre of each line is also affected by the absorption cores. I therefore masked velocities < 500 km/s in each line. A double-peaked Gaussian fit to these lines was performed, with the widths of the Gaussians held equal but their heights allowed to vary, from which the systemic radial velocity, γ , was found as the midpoint of the centres of the two Gaussians. This process found radial velocities consistent with each other, though some have large error bars. Their average systemic velocity $\gamma = 15.5 \pm 2.7$ km/s.

I note the possibility for the mean spectrum to be biased if the sampled spectra are not uniformly distributed across orbital phase. I therefore produced ten phase-binned and averaged spectra. The RV of each spectrum was measured using a double-Gaussian fit as before. In Fig. 4.4 I show sinusoidal fits to these RV measurements. The 4388 Å and 4921 Å lines showed a significant number of anomalous RV measurements and were excluded (though note that, due to the large uncertainties on the sinusoidal fits to these lines, including or excluding these lines does not significantly change the weighted mean). The weighted mean γ from these fits was found to be 14.7 ± 1.2 km/s, consistent with the above.

In Section 4.3.3 I perform a sinusoidal fit to the RV of the bright spot of He II 4686 Å, which is found to be 16.0 ± 5.7 km/s, which is again consistent with the values discussed previously.

However, note that there are several sources of uncertainty in γ which these methods do not counteract. Phase-dependent variations in the bright spot flux




Figure 4.3: Emission lines of He I with best-fitting double-peaked Gaussian profiles. Lines which are blended with neighbours or strongly asymmetric were omitted from the fitting process. The central absorption region of each line was masked (masked points are indicated in red). The measured systemic velocity of each line is quoted in the top right of the panel.




Figure 4.4: Radial velocities measured from phase-binned and averaged WHT spectra for three He I lines. RV measurements are colour-coded according to the spectral line to which they correspond. Sinusoids were fitted to measurements from each spectral line, and are also shown. Note that the phase bins centred on phases 0.05 and 0.95 are affected by the eclipse of the disc, and were therefore not considered during the sinusoidal fit.

are possible, for instance due to partial obscuration of the bright spot by the disc, and could potentially distort measured RVs. The AM CVn binary GP Com shows measurably non-sinusoidal variations in the RV of its bright spot (Marsh, 1999); if the RV curve of the bright spot in Gaia14aae is similarly non-sinusoidal, this may again bias these measurements of γ which are based on sinusoidal fits.

Upper Limit on Hydrogen

The spectral feature near 6560 Å has two possible origins, a He II line (6560.1 Å) or H α (6562.8 Å). I measured the radial velocity shift required to produce this feature in each interpretation, assuming a systemic redshift γ of 16 km/s to match that measured in the previous section and Section 4.3.3. In the He II case, a redshift of 46 ± 17 km/s is required, which agrees well with the redshift measured for the strong He II 4686 Å central spike (Section 4.3.3). In the H α case, a blue-shift of -78 ± 18 km/s would be required. As I do not know of any mechanism that would produce such a blue-shift, I feel confident in identifying this feature as He II emission.

As the hydrogen content of Gaia14aae is of interest regarding its evolutionary history, I attempted to measure upper limits on the EW of H α emission. I did this for two assumed line profiles of H α : firstly, a single-peaked Gaussian with the same width and the same baseline redshift as the He II 6560 Å line, and secondly a double-peaked Gaussian profile with properties held at the average of the He I velocity profiles shown in Fig. 4.3. For this process, the He II 6560 Å was first subtracted from the spectrum using a single Gaussian fit. The uncertainties on each data point in the spectrum were then scaled such that they were consistent with the scatter on data points in the region around 6562 Å. Assuming Gaussian uncertainties, I derived 3σ limits on the H α EW of -0.44 Å in the single-peaked case, and -1.14 Å in the double-peaked case.

Quantifying an upper limit on the hydrogen fraction in the accreted material would require modelling of the accretion disc. Such modelling has been done for several AM CVn binaries, most recently non-LTE modelling of V396 Hya (aka CE 315, Nagel et al., 2009). As Gaia14aae and V396 Hya are similar in some ways (most importantly orbital period and mass transfer rate), I compare my upper limits to the predicted spectra of that work (Nagel et al., 2009, their Fig. 9). Based on this comparison, one cannot rule out a hydrogen fraction of 10^{-5} (predicted EW ≈ 0.6 Å), and it seems unlikely one would be able to rule out 10^{-4} either. This represents an easing of the upper limit of 10^{-5} assumed in the previous chapter.

4.3.2 In-Eclipse Spectrum and White Dwarf Spectrum

As the first known fully-eclipsing AM CVn binary, Gaia14aae presents the opportunity to disentangle the spatial components of the system. Because the eclipse duration is slightly longer than the exposure time of the GTC observations (70 s), it was possible to contain a spectrum that was contained entirely within the eclipse of the central white dwarf. The difference between this spectrum and the out-of-eclipse average spectrum then gives the spectrum of the eclipsed region of Gaia14aae, a region which includes the central white dwarf and part of the inner disc which is eclipsed simultaneously. In Fig. 4.5 I show the spectrum during eclipse, the average spectrum outside of eclipse, and the difference spectrum (the spectrum of the eclipsed light).

The strong central spike emission feature from the He II 4686 Å line is clearly visible in the difference spectrum in Figure 4.5, and is not visible at all during eclipse. This feature must therefore arise from a region which is occulted by the donor at this phase. Likely suggestions are the central white dwarf itself, a region near the central white dwarf, or the inner accretion disc.

The majority of He I emission is still visible during eclipse, as expected because at any one time the majority of the disc is still visible. However, some amount of He I emission is seen in the difference spectrum in Figure 4.5, most clearly around 5875 Å. There are two possible sources for this emission: this may be emission from the regions of the disc which are eclipsed, especially the inner disc; or it may be a residual artefact of the emission in the average spectrum.

The absorption cores of He I lines are visible in the difference spectrum of Figure 4.5 as deep, sharp absorption lines. As was mentioned in Section 4.3.1, these lines are narrower than would be expected from photospheric absorption. In Figure 4.6 I compare the difference spectrum to synthetic spectra of 17000K and 13000K white dwarfs, both with $\log(g)=8.5$ and pure-helium (DB spectroscopic class) atmospheres. These models are identical to those presented by Bergeron et al. (2011).

The absorption seen from the white dwarf in Gaia14aae is both narrower and deeper than either of these model spectra. This is consistent with an understanding that the disc is less transparent at the frequencies of these lines than it is in the continuum. If the accretion disc scale height is negligible compared to the white dwarf radius, as is expected, then slightly less than 50% of the white dwarf light will pass through the accretion disc. For some He I lines, in particular at 5015 and 5875 Å, the depth of the absorption core is over 50% of the continuum, therefore requiring a contribution from photospheric absorption as well.




Figure 4.5: *Above*: GTC spectra of Gaia14aae. The top line shows the average of all out-of-eclipse spectra, offset by +0.2 mJy. The bottom line shows the spectrum taken during eclipse. The middle line shows the difference between the two spectra, offset by +0.07 mJy. The difference spectrum combines the spectra of the white dwarf and of a region of the inner disc that is eclipsed at the same time. It is characterised by He I absorption lines (the cores of the He I features seen in the average spectrum) and He II emission at 4686 Å. Most He I features show little or no emission in the eclipse spectrum, though some emission is seen at the He I 5875 Å line. *Below*: Close-ups of selected He I and He II features: 4686 and 4713 Å (left), 5015 and 5048 Å (centre), and 5875 Å (right).

Table 4.3: Measured positions in velocity space of the central spike and bright spot in the GTC He II 4686 Å line of Gaia14aae, as measured using the method of shifting Gaussians. The systemic velocity offset of each component is denoted by γ .

Feature	K_x (km/s)	K_y (km/s)	γ (km/s)
Central Spike	0.9 ± 3.8	-13.8 ± 3.1	60.7 ± 2.4
Bright Spot	-548.9 ± 12.4	662.1 ± 5.8	16.0 ± 5.7




Figure 4.6: Spectrum of the eclipsed light of Gaia14aae from Fig. 4.5, compared with model spectra for pure helium atmosphere white dwarfs at the two proposed temperatures for the central white dwarf, 13000 K and 17000 K. The 13000 K and 17000 K models were scaled such that the continuum matched the continuum of Gaia14aae, then offset by +0.1 and +0.2 mJy respectively. It is clear that the absorption lines in the observed spectrum are narrower and deeper than the lines in either model, suggesting a significant contribution from absorption of the white dwarf light by the accretion disc.




Figure 4.7: *Above*: Phase-folded, trailed spectra of He I and Si II lines using data from the GTC (left panel) and WHT (centre and right panels). Absorption is shown in blue (dark for greyscale prints), and emission in yellow (light). The features show central absorption, a strong S-wave (maximum blue excursion at phase 0.7), and disc emission (either side of the central absorption). *Below*: Doppler maps produced from the individual spectra. In addition to the central absorption and disc emission, two bright spots can be seen. The second spot is weaker than the first, though by how much varies depending on the emission line. On the Doppler maps I also plot the predicted velocities of the Roche lobes (solid lines), the central white dwarf (cross), the stream of infalling matter (solid lines), and the velocity corresponding to the photometrically-measured bright spot separation radius (dotted line). The rotation of the map is such that the predicted position of the central white dwarf is below the origin of the map.




Figure 4.8: *Above:* Phase-folded, trailed spectra of He II lines using data from the GTC (left panels) and WHT (centre and right panels). Absorption is shown in blue (dark for greyscale prints), and emission in yellow (light). The 4686 Å line shows a prominent, almost motionless central spike emission feature and an S-wave resulting from the bright spot (maximum blue excursion at phase 0.7). The 6560 Å line shows another central spike, albeit much weaker, while the bright spot is not visible. *Below:* Doppler maps produced from the individual spectra. As in the trailed spectra, the 4686 Å Doppler maps show a central spike and one bright spot, and the 6560 Å Doppler map shows just the central spike. On the Doppler maps I also plot the predicted velocities of the Roche lobes (solid lines), the central white dwarf (cross), the stream of infalling matter (solid lines), and the velocity corresponding to the photometrically-measured bright spot separation radius (dotted line). The rotation of the map is such that the predicted position of the central white dwarf is below the origin of the map.

Table 4.4: Coordinates in velocity space of the bright spots in the He I emission lines of Gaia14aae, measured from Doppler maps based on GTC spectra using a two-dimensional Gaussian fit. I also quote total velocity K_{tot} and phase ϕ . The systemic velocity in each map was assumed to be 16 km/s.

Line (\AA)	Bright spot 1				Bright spot 2			
	K_x (km/s)	K_y (km/s)	K_{tot} (km/s)	ϕ ($^\circ$)	K_x (km/s)	K_y (km/s)	K_{tot} (km/s)	ϕ ($^\circ$)
4471.5	-860 ± 30	550 ± 40	1020 ± 40	-123 ± 2	660 ± 70	-560 ± 80	860 ± 70	49 ± 5
4713.1	-1280 ± 80	670 ± 70	1450 ± 80	-117 ± 3	0 ± 200	-1390 ± 60	1390 ± 60	0 ± 8
4921.9	-860 ± 40	680 ± 30	1100 ± 40	-129 ± 2	430 ± 240	-930 ± 140	1020 ± 160	25 ± 13
5015.7	-790 ± 60	580 ± 50	980 ± 50	-126 ± 3	170 ± 80	-950 ± 20	970 ± 30	10 ± 5
5047.7	-930 ± 30	640 ± 40	1130 ± 30	-124 ± 2	-180 ± 80	-1170 ± 30	1180 ± 30	-9 ± 4
5875.6	-720 ± 30	420 ± 80	830 ± 50	-120 ± 5	340 ± 170	-690 ± 100	770 ± 120	26 ± 12

4.3.3 Trailed Spectra and Doppler Maps

In Figures 4.7 and 4.8 I show trailed, phase-folded spectra using data from the GTC and WHT. I show selected spectral lines in each case, including He I, He II, and Si II.

For each spectral line I show a Doppler map, as discussed Chapter 2. These maps were produced from the spectra, after excluding all spectra that were affected by eclipses. On the Doppler maps, I plot theoretical predictions for the Roche lobes, the central white dwarf, and the disc radius at which the bright spot is expected based on eclipse photometry. I also plot two versions of the stream of infalling matter: the ballistic velocities of the stream, and the Keplerian velocities corresponding to positions along the stream. These Doppler maps are oriented such that the predicted velocity of the central white dwarf is in the negative y-direction.

The trailed spectra and Doppler maps of He I and Si II lines are qualitatively similar to each other. Each shows a prominent bright spot which appears as an S-wave in the trailed spectra, and as an emission feature in the top left of the Doppler map. Both He I and Si II lines also show a weaker second bright spot, offset by approximately 120° relative to the first bright spot and appearing at the bottom of the Doppler maps. Finally, the absorption core of all He I lines is visible. In the trailed spectra, the absorption core appears to undergo radial velocity shifts. However, the phase of these velocity shifts is difficult to interpret as following the central white dwarf or the accretion disc. Instead, I interpret these apparent shifts as corresponding to interference from the two bright spots.

The trailed spectra and Doppler maps of the He II 4686 Å line are different to He I. The first bright spot is still present, but a central spike feature is prominent while the second bright spot is not visible. Small velocity shifts can be seen in the central spike, though at a level much lower than the width of the feature itself (the velocity shifts are of order 10s of km/s, where the width of the feature is of order 100s km/s). In the Doppler maps, the central spike feature is slightly offset in a downwards direction with respect to 0, consistent with the predicted velocity of the central white dwarf.

Two methods have been previously used to measure the velocity shifts of such features: a method of shifting Gaussians applied to the individual spectra (Marsh, 1999; Kupfer et al., 2016), and a method in which the Doppler map is fitted with a two-dimensional Gaussian (Roelofs et al., 2005, 2006b; Kupfer et al., 2013, 2016). For GP Com and V396 Hya, using high resolution data, the two methods were found to give consistent results (Kupfer et al., 2016). For the data presented here I found that optimum results could be obtained by using the method of shifting Gaussians

to fit the He II 4686 Å line, as the features are present with sufficient signal strength in individual spectra for the method to converge, and it returns results at a higher precision than the Doppler map fitting method. In the He I lines the bright spots (in particular the second bright spots) are weaker, and the method of shifting Gaussians was unable to converge on these features in individual spectra. I instead used the Doppler map fitting method for all He I lines.

In the method of shifting Gaussians, the region of interest in each spectrum is simultaneously fitted. Each feature in each spectrum is fitted with a single Gaussian whose velocity is offset from the rest wavelength by

$$K_x \sin(2\pi\phi) + K_y \cos(2\pi\phi) + \gamma \quad (4.1)$$

where K_x and K_y are the position of the feature in velocity space, ϕ is the orbital phase at which the spectrum was taken (defined such that eclipses occur at integer values of ϕ), and γ is any systemic Doppler shift of the feature from the rest wavelength.

In order to estimate uncertainties a bootstrapping process was used. 2000 sets of bootstrapped spectra were produced from the parent set. Within each set, each spectrum was produced from its parent spectrum in the original set by selecting with replacement N of the N pixels in the parent spectrum. The method of shifting Gaussians was then applied to each of these bootstrapped sets of spectra. The results from this bootstrapping process were examined by eye, and found to be approximately normally distributed for each parameter. The final uncertainty was chosen as half the difference between the 16th and 84th percentile of these results.

For the Doppler map fitting method, I measured the positions of the two bright spots in velocity space by producing Doppler maps of the spectral lines and fitting the features with a two-dimensional Gaussian. Uncertainties on these values were derived using the bootstrapping method described by Wang et al. (2017, 2018). 2000 sets of bootstrapped spectra were produced as described above. After this, a Doppler map was produced from each set of bootstrapped spectra. Features of interest were fitted using a two-dimensional Gaussian in each bootstrapped Doppler map. The uncertainty on each parameter was again chosen as half the difference between the 16th and 84th percentile of the bootstrapped results.

The results for the He II 4686 Å line are shown in Table 4.3. The total projected velocity for the central spike (CS) is $K_{\text{CS}} = \sqrt{K_x^2 + K_y^2} = 13.8 \pm 3.2$ km/s. This can be compared to the prediction of $K_{\text{WD}} = 17.6 \pm 1.0$ km/s based on the photometry-derived masses and inclination in Chapter 3. The two values agree at the

level of 1.1σ . One can also measure the orbital phase offset between the central spike and the predicted period of the central white dwarf by $\phi_{\text{CS}} = \arctan(-K_x/K_y) = 4 \pm 15^\circ$, which agrees with the expected 0° offset.

The systemic velocity of the bright spot in the He II 4686 Å line, $\gamma = 16.0 \pm 5.7$ km/s, agrees well with the velocity of 15.5 ± 2.7 km/s measured from the wings of the He I lines in Section 4.3.1, and is likely due to the overall velocity of the binary. Relative to this velocity, the central spike is redshifted by a further 44.7 ± 6.2 km/s. Similar redshifts have been seen in the central spikes of other AM CVns and can be attributed to a gravitational redshift. I discuss this further in Section 4.4.1.

The positions of the He I bright spots are shown in Table 4.4. The measured positions of each bright spot is consistent between different emission lines, and the first spot is consistent with the bright spot in He II. It can be seen from the Doppler maps that both spots appear to lie on or close to either the velocity of the ballistic stream or the corresponding Keplerian velocity of its position. I discuss this further in Section 4.4.2.

4.4 Discussion

4.4.1 The Origin of the Central Spike

The central spike appears to follow the central white dwarf in both K -velocity and orbital phase, and disappears when the central white dwarf is eclipsed. This confirms the widely used assumption that the spike originates on or near the surface of the white dwarf, and traces its motion (Marsh, 1999; Kupfer et al., 2016). This agreement can also be taken as an independent test of the stellar properties measured in the previous chapter, producing a reasonably secure confirmation.

As discussed in Section 4.3.3, the spike has a baseline redshift of 44.7 ± 6.2 km/s relative to the velocity of the binary. Such redshifts have been previously seen in AM CVn binaries. The redshift of the central spike is a combination of two possible sources: gravitational redshift, and the Stark effect which causes blueshifts in He I features. As the Stark effect does not affect He II, my measured redshift is likely to be entirely gravitational in origin.

This gravitational redshift can be used to further constrain the origin site of the central spike. Given the mass and radius of the central white dwarf measured photometrically, a redshift of 60 ± 2 km/s would be expected at the white dwarf surface. The redshift I measure for the central spike disagrees with this. Instead my redshift implies that the origin site is a distance of $0.34_{-0.17}^{+0.26} R_{\text{WD}}$ above the surface of the white dwarf. Line emission from this height is difficult to explain. The emitting

material cannot be rotating with the Keplerian velocity as this would significantly broaden the emission line. Infalling material with a significant velocity would also produce Doppler broadening, and hydrostatic suspension at this height is difficult to believe for material cool enough to be dominated by singly ionised helium. It is possible that the mass of the central white dwarf is significantly lower than the eclipse-based measurement; surface emission would suggest $M_1 = 0.75 \pm 0.05 M_\odot$ rather than the $0.87 \pm 0.02 M_\odot$ measured photometrically in Chapter 3. However, it is more likely that the redshift measurement is unreliable. The most likely source of error is the measured RV of the binary, which (as discussed in Section 4.3.1) is subject to several possible complications. Finally, I note the large and non-Gaussian uncertainties on this measured height, which do not allow us to conclusively say that the emission site is inconsistent with the white dwarf surface.

The best-fit Gaussians to the central spike have a FWHM of $4.03(7) \text{ \AA}$, slightly more than the resolution of 2.3 \AA . This FWHM would correspond to a velocity difference of 260 km/s . If the emission comes from the white dwarf atmosphere, this velocity difference would imply a $310 \pm 9 \text{ s}$ the spin period of the white dwarf based on the photometrically measured radius. Given that a significant portion of the broadening will be from instrumental resolution, this is a lower limit on the spin period.

In many other AM CVn binaries, the central spike is visible in both He I and He II. It may be that for these systems, either the overall disc is cooler or the emission site of the central spike is further from the white dwarf, resulting in a cooler emission site that allows both He I and He II to be seen. A cooler disc may be a result of a lower mass transfer rate or a lower mass white dwarf. The mass transfer rate in Gaia14aae is known to be high for its orbital period (Campbell et al., 2015; Ramsay et al., 2018), but note that it is still significantly lower than the mass transfer rate of AM CVn itself, in which He I central spike emission is seen. Alternatively, the absence of central spike emission from the He I lines in Gaia14aae may be related to the edge-on disc absorption seen in this system.

4.4.2 The Positions of the Bright Spots

As can be seen in Figs. 4.7 and 4.8, both bright spots in the He I lines of Gaia14aae lie either on or close to the path of a ballistic stream of matter. In Fig. 4.9, I plot the absolute velocity amplitude (K) and binary phase angle (ϕ) of the bright spots in each He line in the GTC data. I do not use the He I 4713 \AA measurements, which are clearly discrepant with measurements from other lines, perhaps due to blending between He I 4713 \AA and He II 4686 \AA . I define ϕ from the Doppler maps as




Figure 4.9: The measured velocities of the two bright spots, transformed from the values in Table 4.4 to absolute velocity K and orbital phase ϕ , are plotted as black points. I show for comparison the ballistic velocity of the infalling matter stream, and the Keplerian velocity at the position of each point in the stream. The latter may be seen as an approximation for the disc velocity at that point. Lastly, I show as horizontal lines the Keplerian velocities of three radii within the disc: the photometrically measured separation of the first bright spot and outer radius of the disc, and the circularisation radius. Note that larger radii will have a lower Keplerian velocity. All measurements for both bright spots seem to be consistent with the ballistic or Keplerian velocities of the stream, and are at a radius roughly consistent with the photometric measurement of the separation of the first bright spot.

the rotation of the component anticlockwise around the centre, with the accreting white dwarf at phase zero. In Fig. 4.9 I additionally plot lines corresponding to the velocity and phase of a ballistic stream of matter (Lubow and Shu, 1975), and the Keplerian velocity that the disc would be expected to have at that position.

The first bright spot originates from the point at which the infalling stream of matter is decelerated to match the velocity of the accretion disc. As such its velocity is expected to be either equal to that of the ballistic stream, equal to the Keplerian disc velocity at the point of collision, or between the two. The scatter in these measurements appears to be centred between the two velocities, but my most precise measurement, the He II bright spot, lies exactly on the velocity of the ballistic stream. This is of interest for AM CVn mass ratio derivations. In order to convert from the projected velocity of the white dwarf (K_{WD}) to the true velocity of the white dwarf, from which the mass ratio can be found, one must first estimate the inclination of the system. The inclination can be measured from the velocity of the bright spot, if it is assumed to agree with either the ballistic stream or the Keplerian disc velocity. Past studies have generally either assumed the ballistic case (Roelofs et al., 2006b) or used the ballistic and Keplerian cases to derive lower and upper limits (Marsh, 1999; Kupfer et al., 2016). Based on the scatter in these measurements, the latter assumption is safer.

While a second bright spot is seen in many AM CVn binaries (Roelofs et al., 2005, 2006b; Kupfer et al., 2013, 2016), its origin is still unclear. Past studies have suggested it may result from some fraction of the infalling stream continuing past its first contact with the disc, and instead interacting with the disc at a later point in its path (eg. Kupfer et al., 2016). Roelofs et al. (2005) suggested an explanation in which this ‘overflowing’ fraction of the stream passes through the disc along its ballistic trajectory (remaining within the plane of the disc). The second bright spot occurs at the point where the material re-impacts on the disc during its second approach, once the stream has exited the disc once and is falling back towards the white dwarf. An alternate proposal has been modelled by several numerical simulations (Lubow, 1989; Armitage and Livio, 1998, Wood, priv. comm.). In this proposal, the impact of the stream at the edge of the accretion disc causes some fraction of the infalling material to deflect above and below the disc. The deflected material continues to follow an approximately ballistic path, and terminates at a second bright spot where the material reimpacts the disc. The key observational difference between the two models therefore concerns the radius at which the bright spot forms – at the outer edge of the disc in the Roelofs case, and at the point of closest approach of the ballistic path in the Lubow/Armitage/Livio case. Other

suggested models include those in which a standing wave forms in the outer disc, potentially explaining the approximate 120° separation of the two spots.

In Gaia14aae, the agreement of the second spot velocity with the path of the ballistic stream is suggestive. However, unlike in the mechanism suggested by Roelofs et al. (2005), the second bright spot in Gaia14aae appears at a region in which the ballistic stream would be travelling *outward* through the disc, not inward. As such, the second bright spot in Gaia14aae may favour the mechanism of upward and downward deflection suggested by Lubow (1989) and Armitage and Livio (1998), in which the path of the material is approximately ballistic. However, I note that the location of the second bright spot in Gaia14aae does not exactly match the predictions of Lubow (1989) and Armitage and Livio (1998), in that they predict a location near the closest approach of the stream to the central white dwarf. Indeed, it is unclear that the ballistic stream path remains physically meaningful beyond the point of closest approach. It is interesting to note that the radius corresponding to the velocity of the second spot in Gaia14aae is approximately equal to the radius of the first bright spot, but it is unclear whether there is a physical reason for this.

If the second bright spot can be assumed to lie on the ballistic stream path in all AMCVns, this would provide an additional constraint on spectroscopic mass ratio measurements. However, given the uncertainty that remains regarding the formation mechanism of the second bright spot, extreme caution should be exercised before adopting such assumptions. It is possible that the agreement of the measured velocity and phase with the ballistic stream is entirely coincidental.

4.4.3 Implications for the Formation of Gaia14aae

A key purpose of characterising AMCVn binaries is to identify their formation channel. As discussed in the previous chapter, the donor mass and radius of Gaia14aae agree well with models for the evolved CV channel. They also lie in a region of mass-radius space which the helium star donor channel may be able to explain following tweaks to the model.

The biggest discrepancy with the evolved CV channel is that these models predict an amount of hydrogen in the transferred material generally on the order of 1% (Nelson, private communication, 2016; Goliash and Nelson, 2015). Smaller amounts of hydrogen are possible, but the initial conditions required become increasingly improbable as the hydrogen fraction decreases. As the region of possible initial conditions becomes smaller, one would expect to find an increasing number of CVs below the period minimum with spectroscopically visible hydrogen and helium for every AMCVn that is found.

The upper limit on H α EW measured in Section 4.3.1 is less constraining than I had expected. Allowing the accreted material to contain a higher fraction of hydrogen makes the evolved cataclysmic variable channel more favourable for Gaia14aae.

The abundances of CNO elements can also discriminate between evolutionary channels. In the white dwarf donor and evolved CV channels, the relative abundances of these elements will roughly match equilibrium abundances for the CNO cycle (Nelemans et al., 2010). In the helium star donor channel, helium burning increases the abundance of triple- α products C and O relative to N (Yungelson, 2008). The extent of this abundance change depends on the length of time for which burning is able to occur before mass transfer starts. In Gaia14aae I have detected N and O absorption lines, but no C. This may be evidence against the helium star donor channel and in favour of the evolved CV channel. However, given the marginal strength of the N and O detections, the non-detection of C cannot be taken as truly significant. Further observations in ultraviolet or infrared may allow for a tighter constraint on the CNO composition.

In Ramsay et al. (2018), a direct measurement of the mass transfer rate of Gaia14aae was made based on the absolute magnitude of the bright spot. Mass transfer rate can be used to explore the donor’s nature, as it constrains the donor’s response to mass loss $\xi = d \log(R_2/R_\odot)/d \log(M_2/M_\odot)$. The measured mass transfer rate $\log(\dot{M}/[M_\odot \text{yr}^{-1}]) = -10.74 \pm 0.07$ implies $\xi = -0.31_{-0.20}^{+0.23}$ (by Equation 1.9 and assuming gravitational wave radiation is the only form of angular momentum loss). This value agrees better with the prediction based on the helium star and white dwarf donor channels (≈ -0.2) than the prediction based evolved CV channel (≈ 0.0). However, the measurement is not precise enough to rule out any channel. A more precise measurement of mass transfer rate may be possible by eclipse timing. AM CVn binaries show measurable increases to their orbital period over a timescale of years (Copperwheat et al., 2011b; de Miguel et al., 2018), governed by a combination of gravitational wave emission and mass transfer (Deloye et al., 2007). If the stellar masses are known, the rate of orbital period change may be used to measure mass transfer rate. Timing the eclipses of Gaia14aae will make it possible to measure the mass transfer rate to enough precision to discriminate between, or at least further constrain, the evolutionary channels.

There is growing evidence that Gaia14aae may have an unusual formation channel. When compared to other AM CVns, the donor of Gaia14aae appears to be unusually inflated for its orbital period, and indeed has the largest radius of all AM CVn donors with measured radii (see Fig. 5.10). It remains the second-longest

period AM CVn to undergo a dwarf nova outburst (Campbell et al., 2015), and its absolute magnitude is unusually bright for systems of its orbital period (Ramsay et al., 2018, Fig. 3), both suggesting an unusually high mass transfer rate for its orbital period. The mass transfer rate measured in Ramsay et al. (2018) agrees with this, though with large uncertainties.

If Gaia14aae is atypical for a long-period AM CVn, this strengthens the case for its having formed by the evolved CV channel, which is expected to contribute a minority of AM CVns at long orbital periods. Binaries formed by the evolved CV channel are expected to have more inflated donors and higher mass transfer rates than AM CVns formed by the other two channels. A number of CVs with visible hydrogen have been detected that are below the typical CV period minimum and show both hydrogen and helium. Ten such systems are known (as of Breedt, 2015; see Augusteijn et al., 1993, 1996; Breedt et al., 2012, 2014; Carter et al., 2013; Littlefield et al., 2013 for individual systems). As such, one might expect a small number of AM CVns to have formed by the same channel.

4.5 Conclusions

I have presented phase-resolved spectroscopy of Gaia14aae. The system shows a series of He I and He II emission lines, as well as absorption and emission from several metals. The He I lines have a line profile consistent with emission from a disc and two bright spots, while the He II lines show one bright spot and the ‘central spike’ emission feature seen in many AM CVn binaries.

By comparing in-eclipse and out-of-eclipse spectra, I confirm that the central spike is eclipsed at the same time as the central white dwarf. Using the Doppler shifts of the central spike, I derive a projected white dwarf velocity amplitude which is in agreement with my prediction based on the stellar properties derived from eclipse photometry, thus confirming the widely used but never directly tested assumption that this emission traces the velocity of the white dwarf.

I tested two methods to measure the radial velocity of the central spike. The widely used method of fitting to Doppler tomography was found to be too imprecise to measure the white dwarf velocity for Gaia14aae, but the method of sinusoidally shifting Gaussians produced a higher precision measurement.

I measured the systemic velocities of both the central spike and the bright spot, using the difference to compute the gravitational redshift of the central spike. I found a gravitational redshift of 44.7 ± 6.2 km/s, corresponding to an emission site $0.34^{+0.26}_{-0.17} R_{\text{WD}}$ above the white dwarf surface.

The positions of both bright spots coincide with predicted velocities of the stream of matter. In the case of the first bright spot, this matches with expectations. In the case of the second bright spot, this is somewhat surprising as the spot appears on the outward trajectory of the accretion stream. This may suggest that the second spot, like the first, results from an interaction between stream and disc. However, the mechanism of such an interaction, and why it appears at this particular point on the stream's path, is unknown, and the association may be down to chance.

In Chapter 3, I attempted to answer the question of how the binary formed, but was unable to distinguish between two possible formation channels, the helium star donor channel and the evolved cataclysmic variable channel. In this chapter I have presented marginal detections of nitrogen and oxygen compared to a non-detection of carbon, and an upper limit on hydrogen that was less restrictive than expected. I also discuss evidence from Campbell et al. (2015) and Ramsay et al. (2018) that Gaia14aae may have an unusually high mass transfer rate for its orbital period. Combining these points, I tentatively suggest that Gaia14aae may be an unusual example of an AM CVn that has formed by the evolved cataclysmic variable channel. However, these arguments are far from conclusive. Further observations may help to constrain the formation channel of Gaia14aae, including ultraviolet or infrared spectroscopy to search for carbon, and long-term timing of its eclipses in order to constrain the gravitational wave radiation and hence its mass transfer rate.

Chapter 5

A 15.7 Minute AM CVn Binary in a *K2* Field

Our understanding of accretion disks is comparable to astronomers' understanding of stars before the discovery of nuclear fusion.

— Bohdan Paczyński (as quoted in *First Light* by Richard Preston, 1995)

5.1 Introduction

The previous two chapters have focused on one AM CVn binary, Gaia14aae. In this chapter I present the discovery of a new AM CVn system, SDSS J135154.46–064309.0 (henceforth J1351). The binary has an orbital period of 15.7 min, putting it at the short-period end of the AM CVn period distribution (see Fig. 1.7). As discussed in Chapter 1, this makes J1351 a high-state system, whose accretion disc is in a permanently hot and ionised state. Such short-period systems are expected to make up a minority of AM CVn binaries, due to their high rate of period evolution. However, they are of interest from an evolutionary standpoint as the youngest AM CVn systems (according to the white dwarf and helium star donor models). Only five disc-accreting, high-state AM CVn-type systems were previously known (including ES Cet), and there was a large gap at short periods between ES Cet (10.3 minutes) and AM CVn itself (17.1 minutes) which is filled in by J1351.

J1351 was discovered as a variable blue object which was observed by the

Kepler satellite as part of *K2* Campaign 6. This gives an exciting opportunity to characterise the binary. Space-based photometry can be a powerful tool for resolving photometric signals by providing continuous, long-baseline coverage of a target. Fontaine et al. (2011) reported the discovery of SDSS J1908+3940, a high-state AM CVn found in the *Kepler* field. The full 1052-day *Kepler* lightcurve on that system was presented in Kupfer et al. (2015), in which the long baseline allowed for exquisite constraints on the system’s periods and their long-term phase evolution.

The analysis of this binary will make use of the ‘superhump’ technique for characterising AM CVn binaries. In the previous two chapters I described the application of two techniques used for characterising AM CVn binaries, using eclipse photometry and phase-resolved spectroscopy. Both techniques can be applied only to a limited number of systems: only two eclipsing AM CVn systems are known, and only certain systems show the central spike that traces the motion of the primary star. The superhump method can be applied to a much larger number of AM CVn binaries, although certain caveats must be borne in mind when interpreting results derived by this method. This technique uses photometric periodicities measured in the binary.

While AM CVn stars often show variability on a multitude of timescales (Fontaine et al., 2011; Kupfer et al., 2015), there are three characteristic timescales that have physical motivation (Skillman et al., 1999). Firstly, the orbital period can be measured spectroscopically, and in some systems has a photometric equivalent as well (eg. Copperwheat et al., 2011b). Secondly, if the disc of the AM CVn is eccentric (as is possible due to their mass ratios, Whitehurst, 1988), the disc will precess under the tidal field of the donor. This precession period is occasionally visible in either spectroscopy or photometry of AM CVn systems (eg. Patterson et al., 1993; Skillman et al., 1999). Thirdly, a photometric signal at a period known as the ‘superhump’ period is visible in many AM CVn stars, especially in high-state systems or systems in outburst. This signal originates from a tidal interaction between the disc and the donor star, and is found at the beat frequency between the orbital and disc precession periods (Patterson et al., 1993)

$$f_{\text{sh}} = f_{\text{orb}} \pm f_{\text{prec}}. \quad (5.1)$$

A superhump period which is longer than the orbital period (‘−’ in Equation 5.1) indicates that the disc precession is apsidal (precession within the plane of the system). A superhump period which is shorter than the orbital period indicates that the disc precession is nodal (precession of the axis of rotation of a tilted disc).

In both AMCVn binaries and SU UMa binaries (a class of CVs which exhibit the same phenomenon), apsidal precession is found to be more common. The orbital period and superhump period are similar in length (generally within a few percent). Therefore, even in systems which show photometric signatures on both timescales, photometry over a long baseline is often required to separate the two signals (eg. Armstrong et al., 2012).

An empirical relation exists relating the orbital period, superhump period, and the mass ratio (q) of the two stars (Patterson et al., 2005; Knigge, 2006; Kato et al., 2014). As a large fraction of AMCVn systems show superhumps, this relation provides a means to determine q for systems for which the spectroscopic and eclipse methods are not viable. However, it should be noted that there is some uncertainty in applying this method to AMCVn systems. The relationship was determined and calibrated using hydrogen-accreting CVs, and has not been well tested for helium-accreting systems (Pearson, 2007). There are two AMCVn binaries for which results from the superhump method can be compared to results from other methods: AMCVn itself has been studied using both spectroscopic and superhump methods, producing discrepant measurements of $q = 0.18 \pm 0.01$ and 0.101 ± 0.005 respectively (Skillman et al., 1999; Roelofs et al., 2006b), while the partially-eclipsing binary YZLMi was studied with both eclipse and superhump methods, producing consistent values of $q = 0.041 \pm 0.002$ and 0.049 ± 0.008 (Copperwheat et al., 2011a)¹. As long as this uncertainty is considered, the superhump method represents an opportunity to widen the number of AMCVn binaries with measured q , a vital step in any statistical discussion of AMCVn formation channels.

In Section 5.2, I will describe the original *K2* observations as well as follow-up observations undertaken to characterise the system. In Section 5.3, I present the data obtained during these observations. Finally in Section 5.4, I justify the AMCVn classification and describe my interpretation of these data in the context of that classification.

5.2 Observations

A summary of the observations obtained for this work is given in Table 5.1.

¹Several versions of the empirical superhump relationship are available. For both superhump-derived mass ratios quoted in this paragraph I have used the relationship of Knigge (2006).

Table 5.1: A summary of the follow-up observations presented in this work. Spectra are in the wavelength range 3600–5200 Å except where marked with a †.

Instrument	Date	Filters / Slit Width (")	Exposure time (s)
McDonald	2017-03-02	BG40	241×30
McDonald	2017-03-03	BG40	101×30
McDonald	2017-03-05	BG40	324×30
McDonald	2017-03-06	BG40	595×30
SOAR spectra	2017-04-20	3.21	6×300
SOAR spectra	2017-04-20	1.19	8×300
SOAR phot.	2017-04-20	S8612	125×20
SOAR spectra	2017-04-21	3.21	4×300
SOAR phot.	2017-04-21	S8612	230×20
ULTRACAM	2017-05-03	$u'g'r'$	780×8 *
ULTRACAM	2017-05-04	$u'g'r'$	865×8 *
ULTRACAM	2017-05-14	$u_s g_s r_s$	1516×7 *
SOAR spectra†	2017-05-29	1.19	4×300
SOAR spectra	2017-05-29	1.19	6×300
SOAR spectra	2017-05-30	1.19	24×300

* Exposure times for u' and u_s were increased by a factor of 3 to compensate for the lower sensitivity in that band.

† Wavelength range 5200–6700Å.

5.2.1 *K2* Photometry

J1351 (a.k.a. EPIC 212759353, $K_p = 18.9$ mag) was observed in *K2* Campaign 6, which lasted roughly 80 days from 14 July 2015 to 30 September 2015. J1351 was targeted with long-cadence (29.4-min) exposures as a high-probability white dwarf candidate, based on its blue colours and high proper motion (Gentile Fusillo et al., 2015).

Several different pipeline extractions of the *K2* photometry were examined, and a final light curve produced from the Pre-search Data Conditioning pipeline from the *Kepler* Guest Observer office (Van Cleve et al., 2016), which uses a 4-pixel fixed aperture.

5.2.2 McDonald/ProEM Photometry

The appearance of periodic variability in the *K2* data of J1351 motivated the collection of follow-up data to characterise the system. Time-series photometry was obtained on J1351 on 2017 March 2, 3, 5, and 6 with a frame-transfer Princeton Instruments ProEM camera on the 2.1m Otto Struve Telescope at McDonald Observatory. Observations were made through a broad (3300 – 6000 Å) BG40 filter to reduce sky noise, and each exposure was 30 s long. Each frame was dark- and flat-corrected with standard IRAF tasks, using calibration data from the start of each observing night. Circular aperture photometry was measured for the target and two bright comparison stars in the field using the IRAF script CCD_HSP (Kanaan et al., 2002). Transparency variations were corrected for and final relative light curves were obtained by dividing the target flux by the weighted mean of the comparison star fluxes.

5.2.3 SOAR/Goodman Spectroscopy

Optical spectra were obtained in April and May 2017 using the 4.1-m Southern Astrophysical Research (SOAR) telescope at Cerro Pachón in Chile. The high-throughput Goodman spectrograph (Clemens et al., 2004) was used with a 930 line mm^{-1} grating, with two different grating/camera angle setups that cover roughly 3600 – 5200 Å and 5200 – 6700 Å. In April 2017 both a 3.21" and 1.19" slits were used, covering a wavelength range of 3600 – 5200 Å. In May 2017, the 1.19" slit was used in both wavelength ranges. The 1.19" spectra have a resolution of 2.4 Å.

All spectra were reduced using the software packages PAMELA and MOLLY. The data obtained with the 1.19" slit were wavelength calibrated using iron arc lamp spectra that were recorded before and after the spectra, as well as being interspersed

every 30 minutes on 30 May. The 3.21" spectra were wavelength-calibrated using a master arc that was recorded prior to observing, but these spectra suffer from large wavelength drifts and are not reliable for precision velocities.

The April spectra were flux-calibrated using the standard star LTT 3218, observed through a 3.21" slit. No flux standard was observed in May due to poor weather conditions. Instead, these spectra were flux-calibrated by comparison with the April data. A third-order spline was fitted to averaged spectra from each night, and uncalibrated spectra were multiplied throughout by the ratio of those splines.

5.2.4 SOAR/Goodman Photometry

J1351 was also followed up with time-series photometry using SOAR/Goodman over two consecutive nights in April 2017. These observations were obtained through a blue, broad-bandpass, red-cutoff S8612 filter. All exposures were 20 s, with roughly 2.1 s dead time for readouts. Each frame was bias- and flat-corrected with standard IRAF tasks, and circular aperture photometry was performed.

5.2.5 NTT/ULTRACAM Photometry

Further follow-up photometry was obtained using ULTRACAM. J1351 was observed in May 2017 using Sloan $u'g'r'$ filters for two nights and the custom 'Super-SDSS' filters $u_s g_s r_s$ for the third night. The latter set of filters are designed to cover the same wavelength range as $u'g'r'$ filters with a higher throughput.

The ULTRACAM data were reduced using the standard ULTRACAM pipeline. Images were bias- and dark-subtracted and were divided throughout by a flat field taken during the run. Due to poor weather conditions no flat field was available using the $u_s g_s r_s$ filters, so $u'g'r'$ flats were used instead. The target was flux-calibrated using a nearby, non-variable SDSS comparison star (SDSS J135203.48-064405.1, $m_{u'} = 17.26$, $m_{g'} = 15.64$, $m_{r'} = 15.11$, all error bars 0.01 mag or less). As no flux standards exist for the $u_s g_s r_s$ filters the absolute flux calibration of data from those filters may not be reliable, but these data are only used for timing purposes. Transparency changes due to clouds or atmospheric thickness were removed using the same comparison star.




Figure 5.1: SOAR spectrum of J1351, showing strong, double-peaked HeII emission at 4686 \AA and almost no emission or absorption of other elements. This spectrum is the average of 38 spectra in the 3000-5200 \AA range (all of those with a slit width of 1.3"), and 4 spectra in the 5200-6700 \AA range. The spectrum in the 5200-6700 \AA range has been rebinned by a factor of 2 for visualisation purposes. The spectrum has been normalised by dividing by the continuum, which increases the apparent strength of lines at longer wavelengths where the continuum is weaker.

5.3 Analysis

5.3.1 Spectroscopy

The SOAR spectra of J1351 (Figure 5.1) show a clear double-peaked emission line at 4686 Å, consistent with He II emission. He II emission lines are also present at 4859 Å, 5412 Å and 6560 Å, corresponding to the Pickering series, though the latter of these occurs at a similar wavelength to H α . Weak He I emission lines are seen at 3889 Å and 5876 Å. A weak feature is also seen at 4102 Å which may correspond to hydrogen, but is difficult to confirm at this S/N.

The lack of significant H β or H γ suggests that the feature at 6560 Å can be attributed to He II rather than H α . The strength of the 6560 Å line is comparable to the strengths of the other lines in the Pickering series. In order to provide an independent test of this identification, I converted that region of the spectrum to velocity space twice. In the first conversion I used a central wavelength corresponding to that of the He II line, and in the second to that of H α . These conversions can be compared to the velocity of the He II 4686 Å to identify the closest match (Figure 5.2). By fitting the line profiles with a double Gaussian shape (first converged on the He II 4686 Å line to constrain its parameters), I measured velocity shifts of -20 ± 70 km/s for the He II case and -140 ± 70 km/s for the H α case. This gives a marginal preference for He II, in accordance with my identification.

I measured the radial velocity (RV) shift of the 4686 Å line in each spectrum using a double-Gaussian fit. The RVs vary on a short-period sinusoidal pattern, as shown in a phase-fold of the measured RVs (Figure 5.3). This frequency was measured by fitting a sine wave to the data. I found that the best results came from fitting to the 30 May data alone, this being the longest stretch of continuous data. The resulting frequency is 92.0 ± 0.7 day $^{-1}$. A Lomb-Scargle periodogram (Lomb, 1976b; Scargle, 1982b; VanderPlas, 2018, as implemented in the Python package `astropy`) combining consecutive nights of RV data splits this peak into multiple aliases (Figure 5.4).

The 4686 Å line has an equivalent width of -16.6 ± 0.2 Å. There is some periodicity in the equivalent widths measured (Figure 5.4). However, its period is closer to the observed photometric period (next section) than the spectroscopic RV period, and I suggest that this results from variations in the continuum rather than in the spectral line.




Figure 5.2: Emission lines of J1351 at 4686 \AA (above) and 6560 \AA (below, offset vertically by -0.3), in both cases converted to velocity space. In the left panel, the 6560 \AA line was converted assuming it is a He II line with a central wavelength of 6560.10 \AA . In the right panel, it is assumed to be the $H\alpha$ line with a central wavelength of 6562.72 \AA . The red line shows a double Gaussian fit to the 6560 \AA line, described in Section 5.3.1. In both cases this line appears to be slightly blue-shifted relative to the 4686 \AA line, but the discrepancy is more significant when the line is treated as $H\alpha$.




Figure 5.3: RV data for J1351, phase-folded on a frequency of 91.7 day^{-1} . This figure includes RV measurements from all 38 spectra that were observed with a $1.31''$ slit and which include the 4686 \AA line.




Figure 5.4: Lomb-Scargle periodograms of various datasets overplotted for comparison. Red and yellow dashed lines show the proposed superhump and orbital periods, with circle and diamond markers in the second panel indicating the predicted amplitudes of those signals (calculated from the *K2* signals according to Equation 5.2, and assuming the amplitude of variation is constant for all wavelengths of light). The purple dashed line shows the harmonic of the *K2* Nyquist frequency. The strong *K2* peaks in the top panel are Nyquist bounces of the $2.1660(2) \text{ day}^{-1}$ peak, which clearly have no corresponding peak in the high-cadence data.

5.3.2 *K2* Photometry

The Nyquist frequency of the *K2* data is 24.47 day^{-1} . For any signal with a frequency higher than the Nyquist frequency of the data, there are several systematic effects which must be taken into account (Bell et al., 2017b). Firstly, any super-Nyquist frequency will be under-sampled, and therefore there will be a sub-Nyquist frequency from which it is indistinguishable. This results in the effect known as ‘Nyquist bounces’, in which a periodogram will show each signal several times, reflected off each harmonic of the Nyquist frequency. Secondly, as the exposure time of the observation is significant compared to the period of any super-Nyquist variability, the signal will be smeared. This causes a reduction in the measured amplitude. For a sinusoidal signal, this reduction in amplitude can be described by

$$A_{\text{measured}}/A_{\text{intrinsic}} = \text{sinc}(\pi t_{\text{exp}}/P) \quad (5.2)$$

where P is the period of the signal, t_{exp} is the exposure time, and $\text{sinc}(x) = \sin(x)/x$.

The *K2* data show a signal at a frequency of $2.1660(2) \text{ day}^{-1}$. A Lomb-Scargle periodogram of these data shows that this signal and its Nyquist bounces have by far the highest measured amplitudes of all signals present (Figure 5.5). A phase fold of the *K2* photometry on the $2.1660(2) \text{ day}^{-1}$ frequency shows that the signal is well-approximated by a sinusoid (Figure 5.6). Given the smearing effect described in Equation 5.2, the intrinsic amplitude of this variability would have to be large ($> 25\%$) if the intrinsic frequency were super-Nyquist. To investigate the constancy of this signal, I separated the *K2* data into 20 non-overlapping sections and produced a periodogram of each section (Figure 5.7). When compared with a constant frequency the frequencies measured from these subdivisions have $\chi^2 = 24.1$, and the phase offsets have $\chi^2 = 16.2$, both for 19 degrees of freedom. I therefore do not find any evidence that the frequency of this signal is variable.

Two lower-amplitude signals are also visible in the periodogram (Figure 5.5). I summarise their sub-Nyquist frequencies in Table 5.2. If these signals are super-Nyquist, their intrinsic amplitudes can be predicted by Equation 5.2 (see Figure 5.5).

These three signals obey Equation 5.1 to within 1σ . The same is true for any Nyquist bounce of the $8.506(2) \text{ day}^{-1}$ and $6.339(2) \text{ day}^{-1}$ signals, provided that both are subject to the same number of Nyquist bounces and that the $2.1660(2) \text{ day}^{-1}$ signal is sub-Nyquist. In order to determine the number of Nyquist bounces which these signals are subject to, it is necessary to examine higher-cadence photometry of the system.




Figure 5.5: Lomb-Scargle periodogram of *K2* data, showing a clear peak at $2.1660(2) \text{ day}^{-1}$ along with its higher-frequency Nyquist bounces. The Nyquist frequency of 24.47 day^{-1} and its harmonics are shown by dashed purple lines. Red circles and yellow diamonds show the predicted frequencies and intrinsic amplitudes of super-Nyquist signals that may correspond to the $8.506(2) \text{ day}^{-1}$ and $6.339(2) \text{ day}^{-1}$ *K2* signals, with frequencies reflected about the Nyquist frequency and its harmonics, and amplitudes corrected according to Equation 5.2.

Table 5.2: The sub-Nyquist peaks in a periodogram of *K2* data. Frequencies and amplitudes were found by fitting sine waves to the *K2* data.

Frequency (day^{-1})	Measured Amplitude (%)
2.1660(2)	1.20(4)
8.506(2)	0.21(5)
6.339(2)	0.26(7)




Figure 5.6: Phase-fold of the *K2* data on a frequency of 2.1660 day^{-1} . Individual data are shown as black circles, with red circles representing 20 bins of the data. Although the individual data show considerable scatter, the binned data is well-described by a sinusoid (red dashed line).




Figure 5.7: A running periodogram of the *K2* data, created by separating the *K2* data into 20 non-overlapping subsections and calculating periodograms of each. There is no significant variation in the frequency of the $2.1660(2) \text{ day}^{-1}$ signal.

5.3.3 High-cadence photometry

The McDonald, SOAR, and ULTRACAM photometric data are shown in Figure 5.8. The mean AB magnitudes and scatters across both ULTRACAM nights with $u'g'r'$ filters are $m_{u'} = 18.46 \pm 0.08$, $m_{g'} = 18.60 \pm 0.07$, and $m_{r'} = 18.95 \pm 0.07$, where the quoted error bars are one standard deviation of the data so as to include the intrinsic variability of the system. The data show variability with a period of approximately 15 minutes ($\approx 90 \text{ day}^{-1}$) and an amplitude on the order of 4–6%, with changes to the lightcurve shape between one cycle and the next. Figure 5.4 shows Lomb-Scargle periodograms produced from the ULTRACAM and McDonald data, each using data from two consecutive nights. Several nights also show long-term trends in brightness (on timescales longer than the observing window of these data), which may correspond to the $2.1660(2) \text{ day}^{-1}$ signal, but are more likely to be due to changes in airmass as has previously been observed with the same McDonald setup (Bell et al., 2017a).

These data can be used to select between the Nyquist bounces of the signals in the *K2* data discussed in Section 5.3.2. All Nyquist reflections of the $2.1660(2) \text{ day}^{-1}$ peak in the *K2* data can be easily ruled out. Given the smearing effect described in Equation 5.2, the intrinsic amplitude of this signal would have to be $> 25\%$ if it were super-Nyquist, and such a signal is clearly not present in the short cadence photometry. It is therefore most likely that the true frequency detected by *K2* is $2.1660(2) \text{ day}^{-1}$, corresponding to a period of $664.82 \pm 0.06 \text{ min}$. This period is not measurable in the ground-based data due to the short observing windows of those data.

The $\approx 90 \text{ day}^{-1}$ signal in the high-cadence data lies between the third Nyquist reflections of the measured frequencies for both the $6.339(2) \text{ day}^{-1}$ and the $8.506(2) \text{ day}^{-1}$ signals, which would give these signals intrinsic frequencies of $89.374(2) \text{ day}^{-1}$ and $91.541(2) \text{ day}^{-1}$. I therefore interpret the variability seen in the high-cadence data as a combination of both intrinsic signals. As shown in Figure 5.4, nightly aliasing makes it nearly impossible to disentangle the two signals using single-site, ground-based data. The amplitude measured in the high-cadence data agrees well with the prediction made from the measured *K2* amplitude by Equation 5.2. The strength of this agreement and the lack of other signals in the short-cadence photometry leads us to interpret $89.374(2) \text{ day}^{-1}$ and $91.541(2) \text{ day}^{-1}$ as the true frequencies of the signals found in the *K2* data. Phase-folding the ULTRACAM and McDonald data on these two frequencies gives very similar, sawtooth-shaped lightcurves (Figure 5.9).

The amplitudes of the two signals are approximately equal at their funda-




Figure 5.8: All short-cadence photometry of J1351 obtained for this project. For the ULTRACAM data, the r'/r_s -band and u'/u_s -band data have been offset by +0.4 and -0.4 respectively. Most observing runs clearly show sawtooth-shaped variability on a frequency of around 90 day^{-1} . On some nights (see April 22) this signal changes significantly in shape and amplitude between cycles. Some runs (see March 06 and May 03) also show evidence for longer-term variability in brightness. The quality of data from March 02, March 03 and May 14 data was affected by poor conditions.




Figure 5.9: Phase folds of all g' ULTRACAM photometry and BG40 McDonald photometry of J1351 on frequencies of $91.541(2) \text{ day}^{-1}$ (top panel) and $89.374(2) \text{ day}^{-1}$ (bottom panel). The red lines show mean values calculated from a series of 20 phase bins. Both signals have a sawtooth shape, with a steep rise and a gentle decline, which may be the result of beating between the two signals.

mental frequencies. At their third harmonic (I use “third harmonic” to refer to $3f_0$, where f_0 is the fundamental harmonic) the higher frequency signal of the two has a higher amplitude than the lower frequency signal.

5.4 Discussion

5.4.1 Classification as an AM CVn star

The presence of strong, double-peaked helium lines in the spectrum of J1351, the absence of spectral features from hydrogen, and the 92.0 day^{-1} RV modulations, are all characteristic of AM CVn-type binaries. The double-peaked He II emission originates from an accretion disc around the central white dwarf, and the spectroscopic RV period corresponds to the orbital period of the system. An orbital period this short is only possible in a binary in which both stars are degenerate or semi-degenerate. The orbital period of $15.65 \pm 0.12 \text{ min}$ is between those of ES Cet (10.3 minutes) and AM CVn itself (17.1 minutes), making J1351 the second-shortest-period AM CVn-type binary to accrete via an accretion disc (see Figure 1.7).

He II emission lines are also seen in HMCnc and ESCet, which are both at shorter orbital periods than J1351. In ESCet the 4686 \AA line is particularly strong with an equivalent width of -80 \AA (Espaillat et al., 2005, cf. -16.6 \AA for the same line in J1351). High-state AM CVns at longer periods than J1351 all show absorption lines rather than emission. It may therefore be the case that J1351 lies close to a transition point between emission-line systems (where some unknown physical mechanism produces emission lines in the shortest-period disc accretors) and absorption-line systems which make up the bulk of high-state systems. The other shorter-period system, V407 Vul, is contaminated by a G star and difficult to study spectroscopically, though Steeghs et al. (2006) searched for evidence of emission lines and did not find any.

It is also worth noting that emission lines are sometimes seen in hydrogen-accreting CVs with high-state discs if the inclination of the system is high. The *Gaia* absolute magnitude of J1351 shows that the system is fainter than other disc-dominated systems, which could also be interpreted as evidence that the disc is viewed edge-on (Ramsay et al., 2018). However, the same explanation is less compelling for the emission lines of ESCet, as that system is relatively bright for a disc-dominated system.

5.4.2 The Nature of the Photometric Periods

In comparison with other AM CVns that have been observed over a long baseline (eg. Skillman et al., 1999; Kupfer et al., 2015), J1351 is unusually well-behaved photometrically. Only three photometric periods have been identified in J1351, all of which appear to be stable over the baseline of observations. These three signals are in good agreement with Equation 5.1. In this section I will use the data presented thus far to establish the physical origin of these signals.

The high frequency signals correspond to periods of 15.7306(3) minutes and 16.1121(4) minutes. These periods are comparable to the spectroscopic orbital period of 15.65 ± 0.12 minutes. The likely interpretation is therefore that one signal is at the orbital period and the other the superhump period, an interpretation backed up by the agreement of these periods with Equation 5.1. The 15.7 minute signal is clearly in closer agreement with the spectroscopic orbital period. The 16.11 minute signal disagrees with the spectroscopic period at the 3.9σ level, and is most likely to be the superhump period. This interpretation fits with the pattern that orbital periods are generally shorter than superhump periods. I therefore consider this the most probable interpretation of the signals, suggesting that the disc is undergoing apsidal precession. However, I note that the possibility that they may be the other way around has not been conclusively ruled out. The χ^2 values of sinusoidal fits to the RV data with these frequencies are 56.2 and 86.0 with 35 degrees of freedom.

This interpretation means that it is the orbital period which has a strong third harmonic ($3f_0$). This may be related to the 3:1 resonance between the orbital period of the binary and the orbital period of a region of the disc, which is crucial to the mechanism by which the disc is driven to be eccentric.

The agreement with Equation 5.1 implies that the low frequency signal of $2.1660(2) \text{ day}^{-1}$, corresponding to a period of 664.82 ± 0.06 min, originates from the precession period of an eccentric disc. The period is of the correct order for this interpretation; variability attributed to disc precession has been detected at a similar period (13.38 hours) in AM CVn itself (Patterson et al., 1993; Skillman et al., 1999). The apparent stability of this signal throughout the *K2* observation period is somewhat surprising given the variable nature of accretion discs. It suggests that the radius of the accretion disc remains approximately constant throughout the period of observation.




Figure 5.10: Measured donor masses or mass ratios for a sample of AM CVns, compared to predicted M – R tracks for donors in AM CVns descended from three proposed formation channels. Error bars are diagonal because of the strong constraint on mean density which comes from orbital period (Faulkner et al., 1972). The thick black line shows the M – R relation for a zero-entropy white dwarf. For systems where M_1 is not known, a value of $0.7 \pm 0.1 M_\odot$ has been assumed.

Table 5.3: Summary of properties of J1351, as derived from the *K2* data. Periods are given in minutes.

Property	Value
Orbital period	15.7306 ± 0.0003
Superhump period	16.1121 ± 0.0004
Disc precession period	664.82 ± 0.06
Superhump excess ϵ	0.02425 ± 0.00003
Estimated mass ratio q	0.111 ± 0.005

5.4.3 Mass Ratio

Taking the measurements of the orbital and superhump frequencies determined in Section 5.4.2, one can estimate the mass ratio, $q = M_2/M_1$, of the binary from the empirical relation with superhump excess (Knigge, 2006)

$$q(\epsilon) = (0.114 \pm 0.005) + (3.97 \pm 0.41) \times (\epsilon - 0.025) \quad (5.3)$$

where $\epsilon = (P_{\text{sh}} - P_{\text{orb}})/P_{\text{orb}}$ is the superhump excess. This relation gives similar results to that of Patterson et al. (2005), but with the inclusion of uncertainties on the fit parameters. When applied to J1351 these uncertainties dominate due to the small uncertainty on ϵ . With caveats discussed in Section 5.1, I find $\epsilon = 0.02425(3)$ which gives an estimate of $q = 0.111 \pm 0.005$.

In Figure 5.10 I put this measurement in context with the other AM CVn systems. I show the donor mass and radius implied by this mass ratio, together with comparable donor properties for other known AM CVns, as listed in Table 5.4. Where the q was derived from superhump excess, I redid this calculation using the Knigge (2006) relation in order to ensure consistency. For systems in which only a mass ratio is known (including J1351), I assume a primary mass $M_1 = 0.7 \pm 0.1$. Error bars are diagonal due to the tight constraint on the mean density of the donor which comes from the orbital period of the system (Faulkner et al., 1972). I also show model M_2 - R_2 tracks for the three evolutionary channels which may contribute to the AM CVn population.

From Figure 5.10, the population as a whole appears to include only donors with a significant amount of thermal support. Given the tracks shown here, this seems to favour the helium star donor channel. This is not conclusive, however: it may also be the case that the effect of irradiation of the donor star has been underestimated. It is also worth noting that mass ratios derived by the superhump relation may yet have an unknown bias in helium-dominated systems.

J1351 is in a region of parameter space that can be explained by either

the white dwarf donor channel or the helium star donor channel. It is reasonably consistent with the other high-state systems HP Lib and CXOGBS J1751. AM CVn itself appears to be something of an outlier, as it falls in a region of parameter space that is difficult to explain by any of these formation channels, unless it is a pre-bounce system forming by the white dwarf donor channel.

Table 5.4: A summary of the AM CVn mass ratios used for Figure 5.10. Where q was derived by the superhump method, we recalculate it using Equation 5.3 for the sake of consistency.

Designation	ϵ	q	$M_2 (M_\odot)$	Method	Reference
SDSSJ1351–0643	0.024 ± 0.000	0.111 ± 0.005	–	Superhumps	1
AM CVn ^a	0.022 ± 0.000	0.180 ± 0.010	0.125 ± 0.012	Spectroscopy	2
HP Lib	0.015 ± 0.000	0.074 ± 0.007	–	Superhumps	3
CXOGBS J1751–2940	0.014 ± 0.001	0.070 ± 0.007	–	Superhumps	4
CR Boo	0.011 ± 0.000	0.058 ± 0.008	–	Superhumps	3
KL Dra	0.019 ± 0.000	0.092 ± 0.006	–	Superhumps	5,6
V803 Cen	0.011 ± 0.003	0.058 ± 0.014	–	Superhumps	3
YZ LMi ^b	0.009 ± 0.000	0.041 ± 0.002	0.035 ± 0.003	Eclipses	7
CP Eri	0.009 ± 0.001	0.051 ± 0.008	–	Superhumps	8
SDSSJ1240–0159	–	0.039 ± 0.010	–	Spectroscopy	9
SDSSJ0129+3842	0.009 ± 0.005	0.051 ± 0.023	–	Superhumps	10
GP Com	–	0.020 ± 0.003	–	Spectroscopy	11
SDSSJ0902+3819	0.005 ± 0.002	0.033 ± 0.012	–	Superhumps	12
Gaia14aae	–	0.029 ± 0.002	0.025 ± 0.001	Eclipses	13
V396 Hya	–	0.014 ± 0.004	–	Spectroscopy	14

References: [1] This work; [2] Roelofs et al. (2006b); [3] Roelofs et al. (2007b); [4] Wevers et al. (2016); [5] Wood et al. (2002); [6] Ramsay et al. (2010); [7] Copperwheat et al. (2011a); [8] Armstrong et al. (2012); [9] Roelofs et al. (2005); [10] Kupfer et al. (2013); [11] Marsh (1999); [12] Kato et al. (2014); [13] This work; [14] Kupfer et al. (2016).

^a The tabulated q for AM CVn itself is derived from spectroscopy; its superhump excess gives $q = 0.101 \pm 0.005$ by Equation 5.3

^b The tabulated q for YZ LMi is derived from eclipse photometry; its superhump excess gives $q = 0.049 \pm 0.008$ by Equation 5.3

5.4.4 Space Density

After SDSS 1908+3940 (Fontaine et al., 2011), J1351 is the second AM CVn to be discovered in the footprints of *Kepler* and *K2*. Given the rarity of AM CVns, this is worthy of note. A survey of AM CVns in 11663 square degrees of SDSS DR7 that was complete to a magnitude $g' < 19$ included only 4 AM CVns within that magnitude limit (Carter et al., 2013, though note the total number of AM CVns discovered by the survey was larger). Based on this, Carter et al. (2013) estimated an AM CVn space density of $(5 \pm 3) \times 10^{-7} \text{pc}^{-3}$. Both high-state systems found by *Kepler* and *K2* are also within the $g' < 19$ limit. Scaling the population found by Carter et al. (2013) by the ratio of that area to the total *Kepler*+*K2* area up to and including Campaign 6 (≈ 750 square degrees) one would expect 0.25 AM CVns with $g' < 19$ in the *Kepler*+*K2* footprint. *Kepler* and *K2* have therefore found more AM CVns than would be expected, but given the small numbers involved and the large uncertainty on the space density, the discrepancy is unlikely to be significant.

Also surprising is the coincidence that both these systems are high-state binaries. Including J1351, the six known high-state systems comprise only ~ 12 per cent of the known AM CVn population. However, the sample of all known systems includes a selection bias toward outbursting systems due to transient surveys (eg. Levitan et al., 2013). If one defines a sample including only systems with magnitude < 19 (the magnitude limit used by Gentile Fusillo et al., 2015, who selected this object as a candidate white dwarf), using quiescent magnitudes for outbursting systems, the bias towards outbursting systems is reduced. High-state systems become ~ 35 – 40 per cent of the sample, and finding two high-state systems then becomes somewhat more probable. Note that high-state systems are expected to make up $\lesssim 2$ per cent of the AM CVn population (Roelofs et al., 2007a), but this is countered in a magnitude-limited sample by their brighter absolute magnitudes.

The *Gaia* parallax of J1351 implies a distance of 1300 ± 500 pc (Ramsay et al., 2018), comparable to other high-state systems ES Cet and SDSS J1908+3940.

5.5 Conclusions

I have presented the discovery of J1351, a system with a spectroscopic period of 15.65 ± 0.12 minutes that was discovered using *K2* data. The spectrum, orbital period, and lightcurve of this object are consistent with a classification as an AM CVn-type binary. This makes J1351 one of only a small number of known disc-accreting high-state AM CVn-type systems, and the second discovered using *Kepler* or *K2* photometry.

J1351 has several visible photometric periods, including a disc precession period at 664.82 ± 0.06 minutes, a signal at 15.7306 ± 0.0003 minutes which is in agreement with its orbital period, and a signal at 16.1121 ± 0.0004 minutes which I identify as the superhump period. Using the empirical relation of Knigge (2006), one can estimate the mass ratio of the binary as $q = M_2/M_1 = 0.111 \pm 0.005$. This mass ratio is presented with the caveat that the relationship between superhump excess and mass ratio may not be reliable for helium-dominated binaries.

As a short-period AM CVn, J1351 is likely to be a bright emitter of low-frequency gravitational waves. Further study may provide the mass estimates required to quantify its emission. The presence of a photometric signature of the orbital period provides an exciting opportunity to track the period evolution of the system over the next few years, providing a constraint on the system component masses. However, the alignment of this period with a nightly alias of the superhump period means such efforts will likely require multi-site observations if performed from the ground. The system has been re-observed by *K2* in Campaign 17 in short-cadence (58.8 s) mode, allowing an opportunity to revisit this analysis and providing a longer baseline with which to constrain the period evolution.

This work highlights the fact that AM CVn binaries have photometric variability on both short and long timescales. Sustained, high-speed photometry can yield a great deal of information on the nature of the system.

The mass ratios of all AM CVn binaries for which these measurements are known are shown in Fig. 5.10. There is some evidence that the donors in these systems across the board are more inflated than would be predicted by white dwarf donor models.

Chapter 6

VLT Spectroscopy of Five Accreting Binaries Below the Period Minimum

But that sunset! I've never seen anything like it in my wildest dreams ... the two suns! It was like mountains of fire boiling into space ... We only ever had the one Sun at home.

— Douglas Adams, *Hitchhikers Guide to the Galaxy* (1979)

6.1 Introduction

In Section 5.4.3, I compared the mass ratios of all AMCVn binaries with appropriate measurements to predictions based on evolutionary models. The results (in Figure 5.10) are suggestive that all AMCVn donors have some amount of thermal support, and few, if any, are fully degenerate. In fact, most are sufficiently inflated as to agree much better with helium star donor models than with white dwarf models, and there are no systems which can uniquely be explained using the white dwarf donor model tracks presented here (given the uncertainties on these measurements). Two possibilities present themselves: it may be that none of these binaries have white dwarf donors, as suggested by Shen (2015); or it may be that some source of heating inflates the donor more than predicted by current models.

A question also remains around the collapse of the donor towards full degeneracy at orbital periods of around 40 minutes, predicted for both the white dwarf donor channel (Deloye et al., 2007) and the helium star donor channel (Yungelson, 2008). Only a few systems at orbital periods longer than 40 minutes have mass ratio measurements: besides Gaia14aae, the others are GP Com (Kupfer et al., 2016), SDSS J0902+3819 (Kato et al., 2014), and V396 Hya (Kupfer et al., 2016). As was seen for Gaia14aae in Chapter 3, none of these systems appear to have undergone the predicted collapse. It may be that this collapse occurs at longer orbital periods than predicted, or it may be that some or all of these systems are descended from the evolved CV channel (the evidence in Chapter 4 suggests Gaia14aae in particular may have formed this way).

For any system descended from the evolved CV channel, its prior state would have been as a sub-period minimum CV with an evolved donor star.¹ Approximately ten of these systems are known (Breedt, 2015). They have orbital periods below the standard CV period minimum, but are distinguished from AM CVn binaries in that the accreted material, while helium-dominated, also contains some amount of hydrogen. Given the similarities and evolutionary links between these systems and AM CVn binaries, it would be interesting to better characterise these systems.

It should also be borne in mind when considering Figure 5.10 that the calibration of the superhump method for deriving mass ratios is poorly tested for AM CVn binaries. It would be useful to find more systems for which this method can be tested.

For all of the above questions, a secure answer can only be found by increasing the number of well-characterised AM CVn systems and sub-period minimum CVs. Given the low space density of both populations, increasing the sample of well-characterised systems means exploring faint systems that require large telescopes to observe. In this chapter I present phase-resolved spectroscopy from the Very Large Telescope (VLT) of five accreting, short-period binary systems. Two of these systems are AM CVn binaries, and three are sub-period minimum CVs. These systems were observed for two hours each, which in many cases is not sufficient to constrain the orbital period well or to fully characterise the mass ratio. The aim of this chapter is to ascertain which systems are amenable to more in-depth follow-up.

The observations undertaken for these data are described in Section 6.2. In Section 6.3, I describe the previously published data on each system, and present and analyse the VLT data. In Section 6.4, I discuss the implications of these preliminary

¹Throughout this chapter I will use ‘sub-period minimum CV’, ‘helium-rich CV’ and ‘He CV’ to refer to systems containing hydrogen, not including AM CVn binaries.

Table 6.1: Targets observed for this chapter. Magnitudes are taken from *Gaia* and are given in the G band, except where otherwise stated. Parallaxes are from *Gaia*. I have not included the *Gaia* magnitude of ASASSN-14ei as it is significantly discrepant from the CRTS V magnitude (at $G=16.4$).

Target	Coordinates	Mag	Outburst	Parallax (mas)
SDSS J1505+0659	15:05:51.58 +06:59:48.7	19.1	–	6.3(5)
ASASSN-14ei	02:55:33.39 –47:50:42.0	$V=18.0$	$V=11.9$	3.92(5)
CRTS J2333–1557	23:33:13.0 –15:57:44	19.8	$V=17.4$	0.04(70)
CRTS J1028–0819	10:28:42.9 –08:19:27	19.4	$V=14.8$	1.4(3)
V418 Ser	15:14:53.64 +02:09:34.54	20.2	$V=15.8$	0.5(6)

Table 6.2: Observations presented in this chapter.

Target	Instrument	Date	Exposure (s)
SDSS J1505+0659	XSHOOTER	2015 April 9	16×300
ASASSN-14ei	XSHOOTER	2015 August 11	22×240
CRTS J2333–1557	FORS2	2015 August 13	11×240
		2015 September 8	11×240
CRTS J1028–0819	XSHOOTER	2015 April 9	18×300
V418 Ser	FORS2	2015 April 9	21×240

results, and the possibility of follow-up on these systems.

6.2 Observations

The targets for these observations were selected from known but poorly characterised accreting binaries believed to have orbital periods shorter than the period minimum. While all have been included in previously published lists of short-period binaries (eg. Ramsay et al., 2018), several have no peer-reviewed data available and none have been the subject of detailed follow-up before now.

All targets were observed using the Very Large Telescope (VLT) at Paranal Observatory, Chile, under observing proposal 095.D-0888. Each target was observed by either XSHOOTER, an Echelle spectrograph mounted on Unit Telescope Two (UT2) of the VLT (Vernet et al., 2011), or by FORS2 (the FOcal Reducer and low dispersion Spectrograph), a spectrograph with lower dispersion which is mounted on UT1 of the VLT (Appenzeller et al., 1998). Fainter targets ($G > 19.5$) were observed with FORS2, while the higher resolution available from XSHOOTER was possible only for brighter targets. In each case, the targets were observed for two sets of hour-long observing blocks; for CRTS J2333–1557 these blocks were separated by one month, while for other targets the observations were consecutive. The dates,

Table 6.3: Classifications and previously known information on the target systems. P_{prev} refers to previous orbital period estimates, which for all systems but SDSS J1505+0659 are photometric in origin.

Target	Class	P_{prev}	P_{sh}
SDSS J1505+0659	AM CVn	50.6	
ASASSN-14ei	AM CVn	40.0	41.63(3)
CRTS J2333–1557	He CV	61.70429(1)	63.6(1)
CRTS J1028–0819	He CV	52.1(6)	54.9(2)
V418 Ser	He CV	–	64.3

instruments and exposure times of each observation are detailed in Table 6.2.

Observations with FORS2 were carried out using the 600B grism and a slit width of 1", giving a resolution ($\lambda/\Delta\lambda$) of 780. Spectra were reduced using PAMELA and MOLLY as discussed in Chapter 2. Wavelength calibrations were performed using a HgCd arc lamp observed the preceding day. Data were flux calibrated against spectrophotometric standard stars observed the same night: LTT 7379 (to calibrate V418 Ser), LTT 1020 (to calibrate CRTS J2333–1557 on 2015 August 13), and LTT 1788 (to calibrate CRTS J2333–1557 on 2015 September 8).

XSHOOTER observations were carried out using a slit width of 1.2" in the VIS arm (5500 to 10200 Å) and 1" in the UVB arm (3200 to 5500 Å), producing resolutions ($\lambda/\Delta\lambda$) of 6500 and 5400 respectively. Data from the NIR arm of XSHOOTER was not used due to the blue nature of these objects. These data were reduced using the standard XSHOOTER pipeline produced by ESO (REFLEX).

6.3 Results

Average spectra of each object are shown in Figures 6.1 (objects observed with XSHOOTER) and 6.2 (objects observed with FORS2). The profile of the 5876 Å He I line for each system is shown in Figure 6.3 to demonstrate the variety of line profiles between systems. Classifications of the targets, as well as previously known periods, are summarised in Table 6.3.

6.3.1 SDSS J1505+0659

SDSS J1505+0659 is a known AM CVn binary reported by Carter et al. (2014b). The system was discovered as part of a spectroscopic survey for AM CVn systems, selecting targets according to their Sloan Digital Sky Survey (SDSS) colours. That paper included four spectra of the system retrieved from the SDSS database, from which they suggested an orbital period of 50.6 min. The uncertainties on their




Figure 6.1: Average spectra of systems observed with XSHOOTER. Identified spectral features have been labelled. Grey shaded regions are dominated by telluric absorption, which has not been corrected for.




Figure 6.2: Average spectra of systems observed with FORS2. Identified spectral features have been labelled.




Figure 6.3: Profiles of He I 5876 \AA lines for all five systems. Also shown is the $H\beta$ emission profile for CRST J2333–1557, for comparison with its He I profile.




Figure 6.4: Trailed spectra of the He I 5875.6 Å line in SDSS J1505+0659, showing the velocity excursions of the central spike. Note that the spike is slightly redshifted from its rest wavelength, as is commonly seen.

measured radial velocities were large compared to the amplitude of the variation, due to the low resolution of SDSS spectra. Given this and the small number of data points available at the time, the orbital period was considered preliminary and subject to significant uncertainty.

The average spectrum of SDSS J1505+0659 is shown in Figure 6.1. The spectrum consists of a series of relatively narrow helium emission lines imposed on a blue continuum. The profile of these helium emission lines is somewhat unusual, being dominated by a strong central spike emission feature with relatively weak emission wings (Figure 6.3). Emission features from sodium and nitrogen are also seen.

The central spike feature undergoes periodic, low-amplitude velocity excursions, as shown in the trailed spectrum (Figure 6.4). These excursions were measured by cross-correlating in each spectrum the eight strongest emission lines with a Gaussian function. The radial velocities (RVs) measured in this way follow a sinusoidal modulation with a semi-amplitude of approximately 5 km/s, around a baseline velocity (γ) which varies between different emission lines. I performed an initial sinusoidal fit to the measured RVs for each line individually, in order to measure γ for each line. These γ values were then subtracted from the RVs in order to perform a combined sinusoidal fit to all emission lines. The γ -subtracted data and their best-fit sinusoid are shown in Figure 6.5, along with a Lomb-Scargle plot.

The orbital period produced by this process is 71.8 ± 2.1 min, where the uncertainty comes from the sinusoidal fit. Using the formal RV uncertainties this fit has $\chi^2 = 105$ for 13 degrees of freedom, suggesting the uncertainties may be significantly underestimated. In Figure 6.5 the uncertainties have therefore been scaled to give a reduced χ^2 of one. This period is significantly different from the 50.6 min period suggested by Carter et al. (2014b). Given the uncertainties on their individual RV measurements this discrepancy may not be surprising. In Figure 6.6 I show the RV measurements of SDSS J1505+0659 folded on our measured orbital period and the Carter period. The preference for our period is clear. This is the longest orbital period of any known AM CVn binary by a significant amount – the longest published orbital period is that of V396 Hya at 65.1 min, which has held this record since 2001. The recently discovered system SDSS J1319+5915 appears to have a slightly longer orbital period than V396 Hya, at 65.6 min (Kupfer, priv. comm.), but still significantly shorter than SDSS J1505+0659. The uncertainty on this 71.8 min period is notable, and some change might be expected if new data is obtained. Fits to the RVs of individual emission lines consistently produced periods in the range of 68–75 min, suggesting that the orbital period is at least of the correct




Figure 6.5: *Above:* Weighted mean RV measurements of SDSS J1505+0659. A sinusoidal fit with a period of 71.8 min is shown in red. Uncertainties are scaled to give a reduced χ^2 of one. *Below:* Lomb-Scargle plot produced from the measured RVs. The strongest peak corresponds to the 71.8 min period fitted above.




Figure 6.6: Phase-folded RVs of SDSS J1505+0659 using the 71.8 min period measured in this chapter (*above*) and the previously published 50.6 min period (*below*).

order.

SDSS J1505+0659 has similarities with other long-period AM CVn binaries. As was noted by Carter et al. (2014b), the line profiles of the system are similar to SDSS J1208+3550 (53.0 min; Kupfer et al., 2013), SDSS J1137+4054 (59.6 min; Carter et al., 2014b), and V396 Hya (65.1 min; Ruiz et al., 2001; Kupfer et al., 2016). The presence of nitrogen emission is reminiscent of GP Com (46.5 min; Nather et al., 1981; Marsh et al., 1990; Kupfer et al., 2016) and V396 Hya (Kupfer et al., 2016). The system has had no detected outbursts in the five years since its discovery, suggesting a relatively low mass transfer rate in common with other long orbital period systems.

When compared to the empirical correlation between orbital period and emission line strength noted by Carter et al. (2013), SDSS J1505+0659 is a clear outlier. For the He I 5876 Å line, I measure an equivalent width (EW) of -37.9 ± 3.2 Å. For a system with an orbital period of 72 min, Carter’s relation would predict an EW on the order of -100 Å. However, there are several other systems with smaller EWs than would be predicted by the correlation, including SDSS J1208+3550 and SDSS J1137+4054. The discrepancy between the EW of SDSS J1505+0659 and this correlation may suggest that the disc is unusually weak, or it may suggest that the central white dwarf is unusually hot for its orbital period (thereby increasing continuum flux relative to the flux of the emission lines). There is also likely to be an inclination dependence in the EW which may cause systems to be outliers.

The mass ratio (q) of a binary can be calculated from the orbital period (P_{orb}) and the velocity amplitude (K_1), if one assumes prior distributions on the orbital inclination (i) and the mass of the accretor (M_1). The calculation uses Kepler’s law and the necessity that orbital velocity $v_1 = K_1/\sin i = 2\pi r/P_{\text{orb}}$, where r is the separation of the accreting star from the centre of mass of the system. From the sinusoidal fit shown in Figure 6.5, I took $P_{\text{orb}} = 71.8 \pm 2.0$ min and $K_1 = 5.3 \pm 1.0$ km/s, assuming normal distributions. A distribution of $M_1 = 0.7 \pm 0.1 M_{\odot}$ was assumed, as is fairly typical for accreting white dwarfs. As no information on i is known, the distribution of $\cos i$ was assumed to be uniform. Combining these gives a median $q = 0.012^{+0.007}_{-0.003}$, where the uncertainties refer to the 16th and 84th percentiles of the distribution. The uncertainties are dominated by the unknown inclination and the uncertainty on K_1 . The distribution of q has a long tail towards higher mass ratios, as shown in Figure 6.7. This mass ratio is notably similar to that of V396 Hya, $q = 0.014 \pm 0.004$ (Kupfer et al., 2016).

A Doppler map of the system was produced using the 71.8 min orbital period. The He I 5015 Å line was chosen because it has the most prominent disc emission.




Figure 6.7: Probability distribution of mass ratios for SDSS J1505+0659, derived from the velocity amplitude of the central white dwarf and several assumptions detailed in the text. The median value of q is marked by the dashed vertical line.




Figure 6.8: Doppler map of SDSS J1505+0659 produced from the 5015 \AA He I line. Emission is dominated by the bright central spike, but an accretion disc and a weak bright spot can also be seen. The white lines show the Roche lobes of the two stars (the accretor star centred approximately on zero, and the donor star above this, centred on approximately $V_y = 450 \text{ km/s}$), as well as the velocities of a ballistic stream of infalling matter, and the Keplerian velocities at each point along the stream. This map has been rotated to the standard orientation based on the RV fit in Figure 6.5.




Figure 6.9: Trailed spectra of the He II 4686 Å and He I 4713 Å lines in ASASSN-14ei. Each spectrum has been mean-subtracted. At some phases, weak S-waves can be seen in both lines (visible as signals moving in parallel) resulting from the bright spot of the system.

The Doppler map is shown in Figure 6.8. The quality of this map is limited by the small number of spectra, giving the map spoke-like artefacts. Despite this, several features can be distinguished. As well as the central spike and disc emission, a weak bright spot can be seen in the top left quadrant. Figure 6.8 also shows predicted velocity tracks for the infalling matter stream, based on $i = 60^\circ$, $q = 0.012$, and $M_1 = 0.7M_\odot$, which approximately agree with the weak bright spot.

6.3.2 ASASSN-14ei

ASASSN-14ei is a known AM CVn, first detected in outburst by the All-Sky Automated Survey for SuperNovae (ASAS-SN, Shappee et al., 2014). The ASAS-SN detection was reported in ATel #6340 (Prieto et al., 2014a), who associated the outburst with the star USNO-B1.00421–0026203. They also note possible detections by ROSAT in the X-ray (1RXS J025533.1–475020) and by WISE in the infrared. Follow-up observations reported in ATel #6475 (Prieto et al., 2014b) include a spectrum showing He I and He II lines, as well as optical, UV and X-ray imaging with SWIFT, and a series of rebrightenings following the original outburst. In vsnet-alert #17575² was reported a photometric period during outburst (likely a superhump period) of 41.63(3) min. Based on archival data from the Catalina Real-time Tran-

²<http://ooruri.kusastro.kyoto-u.ac.jp/mailarchive/vsnet-alert/17575>




Figure 6.10: Doppler maps of ASASSN-14ei produced simultaneously from the 4686 Å He II line (*left*) and the 4713 Å He I line (*right*). Regularly spaced artefacts around the accretion disc result from aliasing between the observation times of the spectra and the orbital period. The He II line features a strong central spike, an accretion disc, and possibly a bright spot on the left hand side. The He I line shows accretion disc emission, but no central spike. There may be also be a bright spot on the left hand side.

sient Survey (CRTS), they suggest a photometric period during quiescence (possibly the orbital period) of ≈ 40.0 min.

The mean XSHOOTER spectrum of ASASSN-14ei is shown in Figure 6.1. A series of helium emission lines are seen, including both He I and He II. Several metal lines are also seen, including calcium and silicon in emission and magnesium and sodium in absorption. The presence of magnesium absorption is interesting. A trailed spectrum shows no notable velocity excursions from the magnesium triplet, suggesting that it originates from the central white dwarf. This may imply a low temperature for the central white dwarf, given the low ionisation energy of magnesium.

The He I line profiles of ASASSN-14ei are unusual, as shown in Figure 6.3. It is not easy to identify any feature as the central spike – the spike-like feature in Figure 6.3 is blue-shifted rather than red-shifted. Trailed spectra of some helium lines show marginally-visible bright spot S-waves. In Figure 6.9, each component spectrum of the trailed spectrum has been continuum- and mean-subtracted, in order to remove stationary components and highlight the velocity-varying bright spot. The period of these S-waves is roughly consistent with the orbital period of ~ 40.0 min proposed previously. Unfortunately the strength of the bright spot is not enough to be detected by the cross-correlation method used for SDSS J1505+0659.

In order to explore the orbital period, Doppler maps were produced for a range of periods between 38 and 43 min. I attempted to select by eye those maps in which the bright spot was most clearly visible. It should be noted that this process was made more difficult by aliasing between the cadence of observations and the orbital period, producing artefacts in the maps. Bearing this difficulty in mind, I found that I preferred maps with orbital periods in the range of 41 to 42.5 min. In Figure 6.10 are shown Doppler maps produced using periods of 41.6 min.

The orbital periods preferred by this Doppler tomography method are slightly longer than the CRTS photometric orbital period of 40.0 min, and agree better with the proposed superhump period of 41.63(3) min. It may be that the modulation reported in vsnet-alert #17575 was orbital in origin. Variation at the orbital period can be seen in some systems during some stages of the outburst (eg. Isogai et al., 2019). However, given the uncertainties involved in this Doppler method (in particular the aliasing of cadence and orbital period) one cannot conclusively rule out an orbital period of 40.0 min using these data.




Figure 6.11: Trailed, continuum-subtracted spectra of CRTS J2333–1557, showing the $H\beta$ emission line at 4861 Å and the magnesium triplet in absorption at 5170 Å. Weak emission from He I 4921 Å and 5015 Å can also be seen. The magnesium triplet undergoes velocity excursions that are clearly greater than any excursion from the hydrogen or helium emission, suggesting that it originates from the secondary star.




Figure 6.12: Doppler maps of CRTS 2333–1557, produced from the 5876 Å He I line (*top left*), the 4686 Å He II line (*bottom left*), the H α line (*top right*), and the H β line (*bottom right*). All maps are affected by artefacts, but there appears to be little consistency between the maps. Possible bright spots are seen in He I (at the left edge), and He II (at the bottom edge).




Figure 6.13: Doppler map combining absorption from all three components of the 5167/73/84 Å magnesium triplet. The dark absorption feature is believed to originate at the donor star. This map is rotated according to the ephemeris in Equation 6.1. One would expect the donor star to be offset in the positive y -direction; the fact that it is not suggests some drift from the ephemeris, likely due to a slight inaccuracy in the period. The offset of this absorption feature from the origin is approximately 500 km/s.

6.3.3 CRTS J2333–1557

CRTS J2333–1557 is an eclipsing, helium-rich CV below the period minimum. It was discovered in outburst by CRTS in 2011 (CSS111019) and reported in ATel #3705 (Woudt and Warner, 2011), who quoted an orbital period of 61.70(2) min. The same ATel also reported a superhump period of 63.6(1) min. A small amount of follow-up photometry with ULTRACAM, covering six eclipses from 2011 and 2014, was presented by Hardy et al. (2017). From the mid-times of these eclipses the following ephemeris was derived:

$$\text{BMJD(TDB)} = 56866.13556(2) + 0.04285020(1)E \quad (6.1)$$

where E represents the number of cycles. The refined orbital period from this paper is 61.70429(1) min. Hardy et al. (2017) also noted an additional outburst in 2014.

The eclipsing nature of the system makes it a good candidate for lightcurve modelling by the same method as was performed on Gaia14aae in Chapter 3. As the required photometry is not currently available, in this chapter I will concentrate on spectroscopy.

The mean FORS2 spectrum of CRTS J2333–1557 is shown in Figure 6.2. It shows a series of emission features from both hydrogen and helium. The profiles of hydrogen and helium lines are notably different, as can be seen in Figure 6.3. Absorption lines are also seen from magnesium. A trailed spectrum of CRTS J2333–1557 shows large velocity excursions of these magnesium absorption lines, suggesting that they originate from the donor star rather than the central white dwarf.

The visibility of absorption lines from the donor star in the spectrum of CRTS J2333–1557 likely suggests a high-mass, large-radius donor. This is consistent with the relatively high outburst recurrence rate, which suggests a high mass transfer rate. This may imply that CRTS J2333–1557 is a pre-bounce system, still evolving towards shorter orbital periods.

Doppler maps of hydrogen and helium lines created using the ephemeris of Hardy et al. (2017) are shown in Figure 6.12. Several possible bright spots are visible (especially in the He I and He II lines), but none appear to be consistent between emission lines. Figure 6.13 shows a Doppler map of the magnesium absorption lines, again created using the published ephemeris. The offset of the donor star feature in Figure 6.13 from the origin is approximately 500 km/s. The donor is rotated by 120° from the expected position of the donor star (the positive y -direction). Given that the orbital period is short and that the last eclipses observed by Hardy et al. (2017)




Figure 6.14: Trailed spectra of the $H\alpha$ emission feature of CRTS J1028–0819. Each spectrum has been continuum-subtracted.

were 11 months prior to the earliest of these spectra, a relatively small inaccuracy in the ephemeris can result in these maps being significantly rotated. Future eclipse photometry could be used to refine the ephemeris.

6.3.4 CRTS J1028–0819

CRTS J1028–0819 was discovered in outburst by CRTS in 2009 (CSS090331). A superhump period was measured during outburst of 54.9(2) min (Kato et al., 2009). Woudt et al. (2012) measured a photometric period during quiescence of 52.1(6) min, which they suggest to be the orbital period.

The mean spectrum of CRTS J1028–0819 is shown in Figure 6.1. It shows a series of hydrogen and helium emission lines, as well as emission from calcium, sodium, magnesium, and nitrogen. Alone of the systems studied here, CRTS J1028–0819 does not appear to show He II emission at 4686 Å.

Trailed spectra of the $H\alpha$ line in CRTS J1028–0819 are shown in Figure 6.14. The system does not appear to show any bright spot S-wave, and no central spike is visible. In order to search for RV variations, the $H\alpha$ line in each spectrum was cross-correlated with a double-peaked Gaussian function, fixed so that the separation of the two Gaussians was 1500 km/s and each Gaussian had a FWHM of 200 km/s. The results are shown in Figure 6.15. A sinusoidal fit to these RV measurements produces a period of 51.4 ± 3.0 min, consistent with the proposed photometric orbital




Figure 6.15: *Above:* RV measurements of the $H\alpha$ line in CRTS J1028–0819, measured by cross-correlating a double-peaked Gaussian function with the emission lines. A sinusoidal fit is shown, with a period of 51.4 ± 3.0 min. *Below:* A Lomb-Scargle plot produced from these RV measurements.




Figure 6.16: Doppler maps of CRTS J1028–0819 produced from the $H\alpha$ line (*left*) and the 5876 \AA He I line (*right*).

period but not precise enough to exclude the proposed superhump period.

Doppler maps of the $H\alpha$ and He I 5876 \AA lines were produced using the photometric orbital period of 52.1 min. Both lines show disc emission. The width of the disc is less than the other systems discussed here, perhaps implying a relatively face-on inclination. A feature at the upper end of the map may be a possible bright spot.

6.3.5 V418 Ser

V418 Ser was classed as a dwarf nova due following its outburst in 2004 (Rykoff et al., 2004), but initially few follow-up observations were performed. Following a CRTS-detected outburst in 2014, photometry showed the system to have a superhump period of ≈ 64.3 min (Kato et al., 2015). A spectrum reported in ATel #6287 (Garnavich et al., 2014) detected the presence of hydrogen.

The mean FORS2 spectrum is shown in Figure 6.2. The system shows a series of emission lines from hydrogen and helium. Emission is also seen from magnesium. The $H\beta$ line shows a notable absorption wing on its blue-ward side (the red-ward wing is hidden by nearby He I emission). The breadth of this absorption suggests it to be photospheric in origin. No absorption wings are visible around the helium lines.

Trailed spectra of V418 Ser show no visible S-wave from any bright spot or central spike, and no RV excursions were measurable for any emission line. It has therefore not been possible to measure the orbital period of V418 Ser.

The orbital period is likely to be long, given the reported superhump period. The lack of any visible velocity excursions in V418 Ser suggests a low-mass donor star. This would imply a low mass-transfer rate, perhaps explaining the lack of any detectable bright spot. The presence of photospheric hydrogen absorption may imply that the disc is unusually faint relative to the central white dwarf, allowing this absorption to be seen, which would also be consistent with a low accretion rate into the disc. Given that the system is particularly hydrogen-depleted even compared to other helium-rich CVs, V418 Ser may have already passed through a period bounce at a shorter orbital period and be evolving towards longer orbital periods.

6.4 Discussion

6.4.1 Evolutionary Natures

The systems discussed in this chapter cover a range of stages of AM CVn binary and sub-period minimum CV evolution. The sub-period minimum CVs appear to span the evolutionary phases for this track. CRTS J2333–1557 is likely to be a pre-bounce CV, featuring a large donor star and a high mass transfer rate, evolving towards shorter orbital periods. CRTS J1028–0819 may similarly have a high mass donor star, given its measurable RV shifts, though the donor is not directly visible. Its orbital period is shorter than CRTS J2333–1557, suggesting it has evolved further along the evolved CV track. V418 Ser shows evidence for a low mass transfer rate and has no measurable RV shifts, perhaps implying that it is a post-bounce CV with a low-mass companion.

ASASSN-14ei is the shortest period system of the sample, and the only system with a period less than 50 min. Its bright spot appears to be weaker than that of Gaia14aae, despite its period being 7-10 min shorter. This may imply a lower accretion rate and lower-mass donor than Gaia14aae. As such, it may favour a white dwarf or helium star donor channel over an evolved CV channel (though this is far from conclusive).

As the longest-period AM CVn system known, SDSS J1505+0659 represents an interesting phase of AM CVn binary evolution. If it is a white dwarf or helium star donor system, it must be the oldest known AM CVn binary descended from these channels. If it is descended from an evolved CV, it must have evolved through




Figure 6.17: The positions of known AM CVn binaries in donor mass-radius space, reproduced from previous chapters. The position of SDSS J1505+0659 implied by the q measured in this chapter is highlighted in red and labelled.

the period bounce some time ago.

The mass ratio of SDSS J1505+0659 measured in the previous section is shown in the donor star mass-radius plot in Figure 6.17. SDSS J1505+0659 appears very similar to V396 Hya, and perhaps forms a sequence with GP Com. All three systems lie in a space between the white dwarf/helium star tracks (which, according to Deloye et al., 2007, should have fully degenerate donor stars at this orbital period) and the evolved CV tracks. Given their uncertainties, both SDSS J1505+0659 and V396 Hya can easily be moved upwards into the evolved CV region, and GP Com can perhaps be explained by white dwarf or helium star tracks. Regardless, further modelling to explore whether any of these tracks can explore this region of parameter space would be informative.

The similarity between the strong central spike feature in the line profiles of

SDSS J1505+0659 and those of other long period systems suggests that the mechanism causing the central spike must be related to orbital period, perhaps via the mass transfer rate. The significant difference in EW between SDSS J1505+0659 and V396 Hya may imply a difference in the luminosity of the disc or of the central white dwarf, perhaps due to different evolutionary histories or to different white dwarf masses. It is interesting to note that none of the sub-period minimum CVs in this sample show central spikes, despite being of similar orbital periods to SDSS J1505+0659 and V396 Hya. However, note that central spikes have previously been seen in sub-period minimum CVs, even in the H α line (Breedt et al., 2012).

6.4.2 Prospects for Future Study

As has been discussed throughout this thesis, some AM CVn binaries (and, by extension, sub-period minimum CVs) are more amenable to study than others. One objective of the observations in this chapter is to determine which of these systems is suitable for characterisation by more extensive future observations. For systems in which the orbital period is measurable, the measurement precision could be improved using a different observing strategy, in which observations are split over consecutive nights in order to extend the baseline of coverage.

SDSS J1505+0659 is a prime candidate for future study. The strong central spike is ideal for measuring RV shifts of the central white dwarf. The 71.8 ± 2.0 min period measured here makes SDSS J1505+0659 an interesting target, as the longest-known AM CVn. A more precise orbital period could be obtained with further observations. More importantly, further observations may help to better constrain the amplitude of the RV curve of the central white dwarf, and may perhaps allow for the bright spot to be better constrained in Doppler space; both of these will help to constrain the mass ratio of the binary.

ASASSN-14ei also merits further study. The Doppler maps in Figure 6.10 show a bright spot which can be used to constrain the orbital period. With the baseline of these observations it is not possible to choose between the photometric signals seen from the system. More extensive observations may allow for a precise orbital period to be determined, leading to a prediction of q via the superhump relation. If RV variations of the central spike become visible, this would provide another system for which the superhump and spectroscopic methods for finding q can be compared.

CRTS J2333–1557 is an eclipsing system, and as such will be best characterised photometrically. The unusual appearance of the Doppler maps in Figure 6.12 may limit the usefulness of spectroscopic observations of this system. However, the

visibility of spectral features from the donor star in CRTS J2333–1557 is promising, and can perhaps be used to constrain photometric results (if, for instance, red noise proves to be problematic in this high-accretion-rate system).

CRTS J1028–0819 shows RV excursions of its accretion disc, and perhaps a weak bright spot in its Doppler maps (Figure 6.16). As for ASASSN-14ei, future observations may allow one to choose between the photometric signals seen from the system, and thereby determine q using the superhump method.

V418Ser is an interesting system, as perhaps the sub-period minimum CV which most closely resembles an AM CVn binary. If it were possible to characterise the system, valuable insight into this evolutionary channel might be gained. However, the absence of any visible bright spot or signatures of either star makes spectroscopic study of this binary difficult, and applying for further spectroscopy would be something of a gamble. It is notable that a photometric signal (likely a superhump signal) was detected during outburst. If a photometric signature of the orbital period could be detected during quiescence, this star may be more amenable to photometric characterisation.

6.5 Conclusions

I have presented XSHOOTER and FORS2 data of five binaries with orbital periods between 40 and 75 min, including two AM CVn binaries and three sub-period minimum CVs. These systems include SDSS J1505+0659, for which the orbital period measured here (71.8 ± 2.1 min) makes it the longest-period AM CVn system known. Many of the systems presented here merit further follow-up observations in order to better constrain their orbital periods and, in some cases, provide measurements of their stellar mass ratios.

Chapter 7

Conclusions and Outlook

Then Tiffany made a hat out of the sky and sat on the old pot-bellied stove, listening to the wind around the horizons while the sun went down.

As the shadows lengthened, many small shapes crept out of the nearby mound and joined her in the sacred place, to watch.

The sun set, which is everyday magic, and warm night came.

The hat filled up with stars ...

— Terry Pratchett, *A Hat Full of Sky* (2004)

Many uncertainties remain surrounding the evolution of AM CVn binary stars. Numerous factors – their low space density, the faintness of their donor stars, the intrinsic, accretion-driven red noise, and the questions surrounding superhumps in helium accretion discs – limit our ability to study these systems from an evolutionary perspective. However, recent results, including those presented in this thesis, have made some progress towards understanding these uncertainties.

The most important result in this thesis comes from the eclipsing AM CVn binary, Gaia14aae. As the first known fully-eclipsing AM CVn binary, Gaia14aae provides an unrivalled opportunity to study the component stars of an AM CVn system. By fitting a geometric model to Gaia14aae’s eclipse lightcurves, I measured a mass ratio $q = 0.0287 \pm 0.0020$ and a donor mass $M_2 = 0.0250 \pm 0.0013 M_\odot$. Direct mass measurements such as this are possible in only a few AM CVn binaries, and the precision possible for Gaia14aae is unprecedented. The measured donor mass is surprising when compared with existing models. The donor is not degenerate,

ruling out (for Gaia14aae) the white dwarf donor channel which has historically been favoured.

The mass and radius measured fit best with the evolved CV channel. Spectroscopy of Gaia14aae presented in Chapter 4 provided marginal detections of nitrogen and oxygen compared to a non-detection of carbon. If believable, a high ratio of nitrogen to carbon also favours the evolved CV evolutionary channel over the helium star channel. The problem of the hydrogen fraction in the accreted material still favours channels other than the evolved CV channel.

It is thought that the evolved CV channel can only contribute a minority of AM CVn binaries to the total population. However, there is evidence that the evolutionary origin of Gaia14aae may be somewhat unusual when compared with other AM CVn binaries. Ramsay et al. (2018) attempted to measure the mass transfer rate of a number of AM CVn binaries, including Gaia14aae, based on a spectral energy density (SED) fit and its *Gaia* parallax. Their results suggest the mass transfer rate of Gaia14aae may be somewhat elevated compared to other AM CVn binaries of the same orbital period. This is consistent with results from the same paper that Gaia14aae is unusually bright for its orbital period, and with the system having experienced an outburst despite its long orbital period (Gaia14aae remains the second-longest orbital period outbursting AM CVn system). It is also notable that in Figure 5.10, the donor of Gaia14aae has the largest radius of any donor measured.

Phase-resolved spectroscopy of Gaia14aae also made it possible to confirm the origin of the ‘central spike’ emission, which has long been thought to originate from the central white dwarf but has never been confirmed to do so until now. This test validates an assumption that has been used in numerous spectroscopic mass ratio measurements. I also attempted to use these spectra to explore another feature of AM CVn spectroscopy, the ‘second bright spot’, though the results were somewhat inconclusive.

Chapter 5 presented the discovery of an AM CVn binary with a 16 minute orbital period, one of the shortest period binaries known. *K2* photometry of the system showed three clear periodic signals, which follow-up observations confirmed to be the orbital period, the superhump period, and the disc precession period. This system will be a strong emitter of gravitational waves, and it may be possible to track the period evolution of the system photometrically, although the presence of a nightly alias with the superhump period may necessitate multi-site observations for such a project.

In the same chapter, I collected together AM CVn mass ratio measurements

from across the literature and compared these to theoretical models. It appears that almost all AM CVn donor stars are more inflated than would be expected in the canonical degenerate donor model. A similar result was found by Ramsay et al. (2018) when considering the absolute magnitudes and mass transfer rates of AM CVn binaries.

Finally, I discussed the necessity of characterising a greater number of AM CVn binaries, as well as the evolutionarily-related sub-period minimum CVs. I presented phase-resolved spectroscopy of accreting binaries with orbital periods below the period gap, including two AM CVn binaries and three sub-period minimum CVs. For some systems the data presented here was sufficient to obtain an approximate orbital period, though the limited coverage only allows for a precision of a few minutes, from which I was unable to confirm the proposed photometric orbital periods. I suggest that four of these systems will benefit from further study, and may help to shed further light on the evolutionary channels of AM CVn binaries, especially the evolved CV channel.

7.1 Implications for AM CVn Formation

There is a growing weight of evidence that the donors in most AM CVn binaries are not fully degenerate. Unequivocally, the donor in Gaia14aae is not a white dwarf. Considering the systems shown in Figure 6.17, some donors may be consistent with hot and thermally inflated white dwarfs. However, the majority of systems lie above the white dwarf tracks. Given the size of the uncertainties on these donors, the scatter of systems is plausibly consistent with a population in which almost no white dwarf donors exist.

As the white dwarf donor channel has generally been considered the favoured channel by population synthesis models, this is certainly of interest. The space density of AM CVn binaries has long been short of theoretical predictions (Carter et al., 2013). This discrepancy could perhaps be explained if the white dwarf donor channel were not possible. This might relate to the potential stability issues for double-degenerate accreting binaries discussed by various works (Marsh et al., 2004; Shen, 2015).

However, at this stage the evidence is only suggestive. The model tracks that I am comparing these data to are, in most cases, over a decade old. Newer models may be able to reconcile the white dwarf donor channel with observations, for instance if some form of heating, such as irradiation, is able to increase the inflation of the donors. Other theoretical questions include the absence of hydrogen

in Gaia14aae, and whether this can be reconciled with the evolved CV track.

From an observational perspective, there are a number of future observations which may help to provide further constraints on AM CVn binary evolution. I discuss these in the next section.

7.2 Future Perspective

There are a number of goals which future observations of AM CVn binaries may build towards in order to shed light on the evolutionary pathways discussed in this thesis. Increasing the number of systems with precise mass measurements is a clear way to explore the nature of the donor. Measurements of mass transfer rate or the composition of the accreted material will also place constraints on this. Spectroscopic or eclipse-derived mass ratios for systems which also show superhumps would allow a more rigorous test of the superhump relation for short-period, helium-dominated binaries, an invaluable step forward given that the majority of mass ratio measurements to date have used this method. Lastly, a more precise measurement of the space density of AM CVn binaries will provide another constraint on evolutionary models.

Many of the known AM CVn binaries have more to offer with further follow-up. Obvious candidates for this are the systems discussed in Chapter 6. Beside these, a number of systems have known superhump periods without known orbital periods or published spectroscopic data – bright examples are NSV 1440 and ASASSN-14mv. Phase-resolved spectroscopy may allow the orbital period to be measured, providing further mass ratio measurements by the superhump method. If spectroscopy-derived mass ratios are also possible for these systems, this would provide further calibration points for the superhump relation.

It was shown in Chapter 4 that the ‘central spike’ emission shown by many AM CVn systems originates on or near the surface of the white dwarf, and that the redshift between the He II central spike and the wings of the He I emission is gravitational in origin. By measuring this redshift for a number of systems, it may be possible to estimate the primary star mass. UV spectroscopy of known systems may also allow for measurements of the primary star masses and of the mass transfer rate (Sion et al., 2011, eg.), as well as an exploration of the CNO element abundances.

Several AM CVn systems show photometric signatures of their orbital periods, including the eclipsing and partially-eclipsing systems Gaia14aae, YZ LMi, and PTF J1919+4815, and the short-period systems HM Cnc, V407 Vul, ES Cet, and SDSS J1351-0643. Long-term period tracking of these systems will allow for the

orbital period change, \dot{P} , to be measured. This will provide a constraint upon the nature of the donor via its response to mass loss, ξ . For systems which already have measured orbital period changes (HMCnc, V407 Vul, and ESCet), further observations may allow for the second derivate of period, \ddot{P} , to be measured. This is particularly of interest for HMCnc and V407 Vul, both of which have negative \dot{P} and both of which are believed to have formed by the double white dwarf channel. A measurement of \ddot{P} could help constrain whether these systems will continue to inspiral until they merge, or whether they will reach a state of stable accretion. Eclipse photometry of the eclipsing, helium-rich CV CSS J2333-1557 will allow for mass measurements using the method applied in Chapter 3, as well as period evolution tracking over a longer timescale.

Perhaps the most important goal for those wishing to find further observational constraints on AM CVn evolution is to increase the number of known AM CVn systems. The advent of large-scale astronomical surveys such as *Gaia*, ZTF, GOTO, J-PAS, SDSS V, DESI, 4MOST, WEAVE, LSST and LISA offers many opportunities for this. *Gaia* provides absolute magnitude information which can be used to remove many of the contaminants that plagued previous AM CVn surveys, and also provides variability information that can select AM CVn systems over isolated, non-variable white dwarfs (though the intrinsic faintness of AM CVn binaries limits the number of systems with precise *Gaia* parallaxes). ZTF is the successor to PTF, and can be expected to find many more AM CVn systems, especially outbursting systems. LSST will be a photometric survey going deeper than 20 magnitudes. Like ZTF, it can be expected to find many outbursting AM CVn binaries. It may also be possible to find low- and high-state systems through their variability; I have shown in Chapter 5 that this is possible in *K2*. As a photometric survey, it may be possible to preferentially target eclipsing systems (Korol et al., 2017). The photometric survey GOTO may provide an opportunity to test this method, and will also be able to detect outbursting systems. Another photometric survey, J-PAS, will provide narrow-band observations at multiple wavelengths that may be useful to select for systems with helium emission lines.

Multi-object spectroscopy surveys such as DESI, 4MOST, WEAVE, and SDSS (the latter of which is now starting its fifth stage with an extension of its coverage to the southern hemisphere) will provide spectroscopy of a huge number of targets. Many of these surveys will deliberately target white dwarfs for calibration purposes, and will incidentally observe similar objects such as CVs and AM CVn binaries.

In the more distant future, LISA is planned to launch in 2034. Long-period

AM CVn binaries will make a contribution to the unresolved background of Galactic sources, and ideally one would like the AM CVn population to be well characterised by this point. On the other hand, short-period AM CVn binaries are expected to be among the brightest resolved sources LISA will detect (Kupfer et al., 2018), and undoubtedly a large number of new systems will be discovered.

Recent results have made significant progress towards understanding the elusive nature of AM CVn binaries. Future observations, and particularly the goldmine of data provided by current and upcoming surveys, will offer numerous opportunities to shed further light on these systems. We are entering a period of great excitement for the study of AM CVn systems.

Bibliography

- Abbott B.P. et al., 2016. *Physical Review Letters*, 116(6):061102. ISSN 0031-9007.
- Abbott T.M.C. et al., 1992. *The Astrophysical Journal*, 399:680. ISSN 0004-637X.
- Aldis T.A., 1870. *Phil. Mag.*, 39:363.
- Althaus L.G. et al., 2010. *The Astronomy and Astrophysics Review*, 18:471–566. ISSN 0935-4956.
- Anderson S.F. et al., 2005. *The Astronomical Journal*, 130:2230–2236.
- Anderson S.F. et al., 2008. *The Astronomical Journal*, 135(6):2108–2113. ISSN 0004-6256.
- Appenzeller I. et al., 1998. *The Messenger*, 94:1–6.
- Armitage P.J. and Livio M., 1998. *The Astrophysical Journal*, 493:898–908.
- Armstrong E., Patterson J., and Kemp J., 2012. *Monthly Notices of the Royal Astronomical Society*, 421(3):2310–2315.
- Augusteijn T., van Kerkwijk M.H., and van Paradijs J., 1993. *Astronomy and Astrophysics*, 267(2):L55–L58.
- Augusteijn T. et al., 1996. *Astronomy and Astrophysics*, 311:889.
- Baron F. et al., 2012. *The Astrophysical Journal*, 752:20. ISSN 0004-637X.
- Bell K.J. et al., 2017a. *The Astrophysical Journal*, 851(1):24. ISSN 1538-4357.
- Bell K.J. et al., 2017b. *The Astrophysical Journal*, 835(2):180. ISSN 0004-637X.
- Belloni D. et al., 2018. *Monthly Notices of the Royal Astronomical Society*, 478(4):5626–5637. ISSN 0035-8711.
- Bergeron P. et al., 2011. *The Astrophysical Journal*, 737(1):28. ISSN 0004-637X.

- Bildsten L. et al., 2006. *The Astrophysical Journal*, 640:466–473. ISSN 0004-637X.
- Breedt E., 2015. *Proceedings of The Golden Age of Cataclysmic Variables and Related Objects - III (Golden2015)*, 25.
- Breedt E. et al., 2012. *Monthly Notices of the Royal Astronomical Society*, 425(4):2548–2556. ISSN 00358711.
- Breedt E. et al., 2014. *Monthly Notices of the Royal Astronomical Society*, 443:3174. ISSN 13652966.
- Breivik K. et al., 2018. *The Astrophysical Journal Letters*, 854:1. ISSN 0004-637X.
- Campbell H.C. et al., 2015. *Monthly Notices of the Royal Astronomical Society*, 452(1):1060–1067.
- Cannizzo J.K. and Nelemans G., 2015. *The Astrophysical Journal*, 803:19. ISSN 0004-637X.
- Carroll B. and Ostlie D., 2013. *An Introduction to Modern Astrophysics*. Pearson, London, 2nd edition.
- Carter P.J. et al., 2013. *Monthly Notices of the Royal Astronomical Society*, 429:2143–2160. ISSN 00358711.
- Carter P.J. et al., 2014a. *Monthly Notices of the Royal Astronomical Society*, 437:2894.
- Carter P.J. et al., 2014b. *Monthly Notices of the Royal Astronomical Society*, 439:2848. ISSN 13652966.
- Clemens J.C., Crain J.A., and Anderson R., 2004. In A. Moorwood and I. Masanori, editors, *Ground-based Instrumentation for Astronomy*, volume 5492, 331.
- Cook M.C. and Warner B., 1984. *Monthly Notices of the Royal Astronomical Society*, 207(4):705–729. ISSN 0035-8711.
- Copperwheat C.M. et al., 2010. *Monthly Notices of the Royal Astronomical Society*, 402:1824–1840.
- Copperwheat C.M. et al., 2011a. *Monthly Notices of the Royal Astronomical Society*, 410(2):1113–1129. ISSN 00358711.
- Copperwheat C.M. et al., 2011b. *Monthly Notices of the Royal Astronomical Society*, 413(4):3068–3074.

- Cowles M.K., 2013. *Applied Bayesian Statistics*. Springer, New York.
- Cropper M. et al., 1998. *Monthly Notices of the Royal Astronomical Society*, 293(2):L57–L60. ISSN 0035-8711.
- D’Antona F. et al., 2006. *The Astrophysical Journal*, 653:1429–1434. ISSN 0004-637X.
- de Miguel E. et al., 2018. *The Astrophysical Journal*, 852:19.
- Deloye C.J. and Taam R.E., 2006. *The Astrophysical Journal*, 649:L99–L102. ISSN 0004-637X.
- Deloye C.J. et al., 2007. *Monthly Notices of the Royal Astronomical Society*, 381(2):525–542.
- Dhillon V.S. et al., 2007. *Monthly Notices of the Royal Astronomical Society*, 378(3):825–840. ISSN 0035-8711.
- Dhillon V.S. et al., 2014. *Monthly Notices of the Royal Astronomical Society*, 444(4):4009–4021. ISSN 0035-8711.
- Eggleton P.P., 1983. *The Astrophysical Journal*, 268:368. ISSN 0004-637X.
- Einstein A., 1916. *Sitzungsberichte der Königlich Preußischen Akademie der Wissenschaften (Berlin)*, 688.
- Eldridge J.J. et al., 2017. *Publications of the Astronomical Society of Australia*, 34:e058. ISSN 1323-3580.
- Espaillet C. et al., 2005. *Publications of the Astronomical Society of the Pacific*, 117(828):189–198. ISSN 0004-6280.
- Faulkner J., 1971. *The Astrophysical Journal*, 170:L99. ISSN 0004-637X.
- Faulkner J., Flannery B.P., and Warner B., 1972. *The Astrophysical Journal*, 175:L79. ISSN 0004-637X.
- Feline W.J. et al., 2004. *Monthly Notices of the Royal Astronomical Society*, 355(1):1–10.
- Fontaine G. et al., 2011. *The Astrophysical Journal*, 726(2):92. ISSN 0004-637X.
- Foreman-Mackey D. et al., 2013. *Publications of the Astronomical Society of Pacific*, 125:306–312. ISSN 0004-6280.

- Frank J., King A., and Raine D., 2002. *Accretion Power in Astrophysics*. Cambridge University Press, Cambridge, 3 edition.
- Garnavich P. et al., 2014. *The Astronomer's Telegram*, 6287.
- Gentile Fusillo N.P., Gänsicke B.T., and Greiss S., 2015. *Monthly Notices of the Royal Astronomical Society*, 448(3):2260–2274.
- Gianninas A. et al., 2013. *The Astrophysical Journal*, 766:3.
- Godon P. et al., 2004. *The Astrophysical Journal*, 602:336–341.
- Goliashch J. and Nelson L., 2015. *The Astrophysical Journal*, 809:80.
- Goodman J. and Weare J., 2010. *Communications in Applied Mathematics and Computational Science*, 5(1):65–80. ISSN 2157-5452.
- Goodricke J., 1783. *RSPTA*, 73:474.
- Green M.J. et al., 2018a. *Monthly Notices of the Royal Astronomical Society*, 477:5646–5656. ISSN 0035-8711.
- Green M.J. et al., 2018b. *Monthly Notices of the Royal Astronomical Society*, 476:1663.
- Green M.J. et al., 2019. *Monthly Notices of the Royal Astronomical Society*, 485(2):1947–1960. ISSN 0035-8711.
- Green P.J., 2013. *The Astrophysical Journal*, 765:1.
- Greenstein J.L. and Matthews M.S., 1957. *The Astrophysical Journal*, 126:14. ISSN 0004-637X.
- Hakala P., Ramsay G., and Byckling K., 2004. *Monthly Notices of the Royal Astronomical Society*, 353(2):453–456. ISSN 00358711.
- Hakala P. et al., 2003. *Monthly Notice of the Royal Astronomical Society*, 343:L10–L14. ISSN 0035-8711.
- Harding L.K. et al., 2016. *Monthly Notices of the Royal Astronomical Society*, Volume 457, Issue 3, p.3036-3049, 457:3036–3049. ISSN 0035-8711.
- Hardy L.K. et al., 2017. *Monthly Notices of the Royal Astronomical Society*, 465:4968–4984. ISSN 0035-8711.

- Heber U., 2016. *Publications of the Astronomical Society of the Pacific*, 128(8):082001. ISSN 0004-6280.
- Hellier C., 2001. *Cataclysmic Variable Stars*. Springer, Berlin.
- Hill C.A. et al., 2017. *Monthly Notices of the Royal Astronomical Society*, 472:2937–2944. ISSN 0035-8711.
- Holberg J.B. and Bergeron P., 2006. *The Astronomical Journal*, 132(3):1221–1233. ISSN 0004-6256.
- Horne K., 1986. *Publications of the Astronomical Society of the Pacific*, 98:609. ISSN 0004-6280.
- Horne K., 1994. In *Astronomical Society of the Pacific Conference Series*, volume 69, 23. Astronomical Society of the Pacific (ASP).
- Howell S.B., 2006. *Handbook of CCD Astronomy*. Cambridge University Press, Cambridge, 2 edition.
- Hulse R.A. and Taylor J.H., 1975. *The Astrophysical Journal*, 195:L51. ISSN 0004-637X.
- Iben I.J. and Tutukov A.V., 1987. *The Astrophysical Journal*, 313:727.
- Isogai K. et al., 2019. *eprint arXiv:1901.11425*.
- Israel G. et al., 2004. *ASP Conference Proceedings*, 315:338.
- Israel G.L. et al., 1999. *Astronomy and Astrophysics*, 349:L1–L4.
- Israel G.L. et al., 2002. *Astronomy and Astrophysics*, 386:L13–L17. ISSN 0004-6361.
- Ivanova N. et al., 2013. *Astronomy and Astrophysics Review*, 21(1):59.
- Jetsu L. et al., 2013. *The Astrophysical Journal*, 773:1. ISSN 0004-637X.
- Jha S. et al., 1998. *International Astronomical Union Circular*, 6983:1.
- Joergens V. et al., 2000. *Astronomy and Astrophysics*, 354:579–588. ISSN 0004-6361.
- Kanaan A., Kepler S.O., and Winget D.E., 2002. *Astronomy and Astrophysics*, 389(3):896–903. ISSN 0004-6361.
- Kaplan D.L., Bildsten L., and Steinfadt J.D.R., 2012. *The Astrophysical Journal*, 758:64.

- Kato T., Hamsch F.J., and Monard B., 2015. *Publications of the Astronomical Society of Japan*, 67:L2. ISSN 0004-6264.
- Kato T. et al., 2009. *Publications of the Astronomical Society of Japan*, 61:S395–S616. ISSN 0004-6264.
- Kato T. et al., 2014. *Publications of the Astronomical Society of Japan*, 66(5):L7.
- Knigge C., 2006. *Monthly Notices of the Royal Astronomical Society*, 373:484–502. ISSN 0035-8711.
- Knigge C., Baraffe I., and Patterson J., 2011. *The Astrophysical Journal Supplement Series*, 194(2):28. ISSN 0067-0049.
- Knigge C., Leigh N., and Sills A., 2009. *Nature*, 457:288–290.
- Koester D. and Kepler S.O., 2015. *Astronomy and Astrophysics*, 583:A86. ISSN 0004-6361.
- Kong X. et al., 2018. *Publications of the Astronomical Society of the Pacific*, 130(8):084203.
- Korol V. et al., 2017. *Monthly Notices of the Royal Astronomical Society*, 470:1894–1910.
- Kotko I. et al., 2012. *Astronomy and Astrophysics*, 544:A13. ISSN 0004-6361.
- Kowalski P.M. and Saumon D., 2006. *The Astrophysical Journal*, 651:L137–L140. ISSN 0004-637X.
- Kraft R.P., Mathews J., and Greenstein J.L., 1962. *The Astrophysical Journal*, 136:312. ISSN 0004-637X.
- Kruszewski A., 1966. *Advances in Astronomy and Astrophysics*, 4:233–299. ISSN 0065-2180.
- Kupfer T. et al., 2013. *Monthly Notices of the Royal Astronomical Society*, 432:2048–2060. ISSN 0035-8711.
- Kupfer T. et al., 2015. *Monthly Notices of the Royal Astronomical Society*, 453(1):483–496. ISSN 0035-8711.
- Kupfer T. et al., 2016. *Monthly Notices of the Royal Astronomical Society*, 457(2):1828–1841. ISSN 0035-8711.

- Kupfer T. et al., 2018. *Monthly Notices of the Royal Astronomical Society*, 480:302–309. ISSN 0035-8711.
- Landau L.D. and Lifshitz E.M., 1971. *The classical theory of fields*. Pergamon Press, Oxford, 3rd edition.
- Leonard P.J.T., 1989. *The Astronomical Journal*, 98:217. ISSN 00046256.
- Levitan D. et al., 2011. *The Astrophysical Journal*, 739(10pp):68.
- Levitan D. et al., 2013. *Monthly Notices of the Royal Astronomical Society*, 430:996–1007. ISSN 00358711.
- Levitan D. et al., 2014. *The Astrophysical Journal*, 785(2):114. ISSN 0004-637X.
- Levitan D. et al., 2015. *Monthly Notices of the Royal Astronomical Society*, 446(1):391–410. ISSN 0035-8711.
- Littlefair S.P. et al., 2014. *Monthly Notices of the Royal Astronomical Society*, 445:2106L.
- Littlefield C. et al., 2013. *The Astronomical Journal*, 145:145. ISSN 0004-6256.
- Lomb N.R., 1976a. *Astrophysics and Space Science*, 39(2):447–462. ISSN 0004-640X.
- Lomb N.R., 1976b. *Astrophysics and Space Science*, 39(2):447–462. ISSN 0004-640X.
- Lubow S.H., 1989. *The Astrophysical Journal*, 340:1064. ISSN 0004-637X.
- Lubow S.H. and Shu F.H., 1975. *The Astrophysical Journal*, 198:383. ISSN 0004-637X.
- Manser C.J. et al., 2016. *Monthly Notices of the Royal Astronomical Society*, 455:4467–4478.
- Maoz D., Mannucci F., and Nelemans G., 2014. *Annual Review of Astronomy and Astrophysics*, 52:107–70.
- Margon B. et al., 2018. *The Astrophysical Journal Letters*, 856:1.
- Marsh T., Parsons S., and Dhillon V., 2017. *The Astronomer’s Telegram*, 10354.
- Marsh T.R., 1989. *Publications of the Astronomical Society of the Pacific*, 101:1032. ISSN 0004-6280.

- Marsh T.R., 1999. *Monthly Notices of the Royal Astronomical Society*, 304:443–450. ISSN 0035-8711.
- Marsh T.R. and Horne K., 1988. *Monthly Notices of the Royal Astronomical Society*, 235(1):269–286.
- Marsh T.R., Horne K., and Rosen S., 1991. *The Astrophysical Journal*, 366:535–543.
- Marsh T.R., Nelemans G., and Steeghs D., 2004. *Monthly Notices of the Royal Astronomical Society*, 350(1):113–128. ISSN 00358711.
- Marsh T.R. and Steeghs D., 2002. *Monthly Notices of the Royal Astronomical Society*, 331:L7–L11. ISSN 0035-8711.
- Marsh T.R. et al., 1990. *The Astrophysical Journal*, 364:637. ISSN 0004-637X.
- McAllister M.J. et al., 2015. *Monthly Notices of the Royal Astronomical Society*, 451:114–125. ISSN 0035-8711.
- McAllister M.J. et al., 2017. *Monthly Notices of the Royal Astronomical Society*, 464(2):1353–1364.
- Moffat A.F.J., 1969. *Astronomy and Astrophysics*, 3:455. ISSN 0004-6361.
- Morales-Rueda L. et al., 2003. *Astronomy and Astrophysics*, 405:249–261. ISSN 0004-6361.
- Motch C. et al., 1996. *Astronomy and Astrophysics*, 307:459–469. ISSN 0004-6361.
- Nagel T., Rauch T., and Werner K., 2009. *Astronomy and Astrophysics*, 499(3):773–781. ISSN 0004-6361.
- Nather R.E., 1985. *Interacting Binaries*. D. Reidel, Dordrecht.
- Nather R.E., Robinson E.L., and Stover R.J., 1981. *The Astrophysical Journal*, 244:269. ISSN 0004-637X.
- Naylor T., 1998. *Monthly Notices of the Royal Astronomical Society*, 296(2):339–346. ISSN 0035-8711.
- Nelemans G., 2003. *Classical and Quantum Gravity*, 20(10):S81–S87. ISSN 02649381.
- Nelemans G., Steeghs D., and Groot P.J., 2001a. *Monthly Notices of the Royal Astronomical Society*, 326:621–627. ISSN 0035-8711.

- Nelemans G., Yungelson L.R., and Portegies Zwart S.F., 2004. *Monthly Notices of the Royal Astronomical Society*, 349:181–192.
- Nelemans G. et al., 2001b. *Astronomy and Astrophysics*, 368:939–949.
- Nelemans G. et al., 2010. *Monthly Notices of the Royal Astronomical Society*, 401:1347. ISSN 00358711.
- Neustroev V.V., Zharikov S.V., and Borisov N.V., 2016. *Astronomy and Astrophysics*, 586:A10.
- Nissanke S. et al., 2012. *The Astrophysical Journal*, 758:131.
- O’Donoghue D., Menzies J.W., and Hill P.W., 1987. *Monthly Notices of the Royal Astronomical Society*, 227(2):347–357. ISSN 0035-8711.
- O’Donoghue D. et al., 1994. *Monthly Notices of the Royal Astronomical Society*, 271(4):910–918. ISSN 0035-8711.
- Paczynski B., 1967. *Acta Astronomica*, 17:287.
- Paczynski B., 1971. *Annual Review of Astronomy and Astrophysics*, 9(1):183–208. ISSN 0066-4146.
- Paczynski B., 1976. In P. Eggleton, S. Mitton, and J. Whelan, editors, *IAU Symp. 73, Structure and Evolution of Close Binary Systems*, 75. D. Reidel, Dordrecht.
- Paczynski , 1977. *The Astrophysical Journal*, 216:822. ISSN 0004-637X.
- Patterson J., Halpern J., and Shambrook A., 1993. *The Astrophysical Journal*, 419:803. ISSN 0004-637X.
- Patterson J. et al., 1997. *Publications of the Astronomical Society of the Pacific*, 109:1100. ISSN 0004-6280.
- Patterson J. et al., 2005. *The Publications of the Astronomical Society of the Pacific*, 117:1204–1222.
- Pearson K.J., 2007. *Monthly Notices of the Royal Astronomical Society*, 379:183–189. ISSN 0035-8711.
- Pickering E.C., 1881. *Proc. American Academy*, 16:1.
- Podsiadlowski P., Han Z., and Rappaport S., 2003. *Monthly Notices of the Royal Astronomical Society*, 340(4):1214–1228.

- Prieto J.L. et al., 2014a. *The Astronomer's Telegram*, 6475.
- Prieto J.L. et al., 2014b. *The Astronomer's Telegram*, 6475.
- Provencal J.L., 1994. *Applications of spectroscopic abundance analysis to late-type stars*. Ph.D. thesis.
- Provencal J.L. et al., 1991. *White Dwarfs, Proceedings of the 7th European Workshop, NATO Advanced Science Institutes (ASI) Series C*, 336:449.
- Ramsay G., Hakala P., and Cropper M., 2002. *Monthly Notices of the Royal Astronomical Society*, 332:L7–L10. ISSN 0035-8711.
- Ramsay G. et al., 2010. *Monthly Notices of the Royal Astronomical Society*, 407:1819–1825. ISSN 0035-8711.
- Ramsay G. et al., 2018. *eprint arXiv:1810.06548*.
- Rappaport S., Joss P.C., and Webbink R.F., 1982. *The Astrophysical Journal*, 254:616. ISSN 0004-637X.
- Rau A. et al., 2010. *The Astrophysical Journal*, 708(1):456–461. ISSN 0004-637X.
- Rebassa-Mansergas A. et al., 2016. *Monthly Notices of the Royal Astronomical Society*, 458:3808–3819. ISSN 0035-8711.
- Ritter H. and Kolb U., 2003. *Astronomy and Astrophysics*, 404:301–303. ISSN 0004-6361.
- Roelofs G. et al., 2006a. *Monthly Notices of the Royal Astronomical Society*, 365:1109–1113. ISSN 0035-8711.
- Roelofs G.H.A., 2007. *The AM Canum Venaticorum stars*. Ph.D. thesis, University of Nijmegen.
- Roelofs G.H.A., Nelemans G., and Groot P.J., 2007a. *Monthly Notices of the Royal Astronomical Society*, 382:685–692. ISSN 00358711.
- Roelofs G.H.A. et al., 2005. *Monthly Notices of the Royal Astronomical Society*, 361(2):487–494. ISSN 0035-8711.
- Roelofs G.H.A. et al., 2006b. *Monthly Notices of the Royal Astronomical Society*, 371:1231–1242. ISSN 0035-8711.
- Roelofs G.H.A. et al., 2007b. *The Astrophysical Journal*, 666(2):1174–1188.

- Roelofs G.H.A. et al., 2007c. *Monthly Notices of the Royal Astronomical Society*, 379:176. ISSN 00358711.
- Roelofs G.H.A. et al., 2007d. *Monthly Notices of the Royal Astronomical Society*, 382:1643–1647. ISSN 0035-8711.
- Roelofs G.H.A. et al., 2009. *Monthly Notices of the Royal Astronomical Society*, 394:367–374. ISSN 0035-8711.
- Roelofs G.H.A. et al., 2010. *The Astrophysical Journal*, 711:L138.
- Ruiz M.T. et al., 2001. *The Astrophysical Journal*, 552:679–684. ISSN 0004-637X.
- Rutkowski A. et al., 2016. *Monthly Notices of the Royal Astronomical Society*, 463:3290–3304. ISSN 0035-8711.
- Rykoff E.S. et al., 2004. *Information Bulletin on Variable Stars*, 5559. ISSN 0374-0676.
- Saio H. and Jeffery C.S., 2002. *Monthly Notices of the Royal Astronomical Society*, 333(1):121–132. ISSN 00358711.
- Sánchez B. et al., 2012. In I.S. McLean, S.K. Ramsay, and H. Takami, editors, *Proceedings of SPIE - The International Society for Optical Engineering*, volume 8446, 84464T. International Society for Optics and Photonics. ISBN 9780819491473. ISSN 0277786X.
- Savonije G.J., de Kool M., and van den Heuvel E.P.J., 1986. *Astronomy and Astrophysics*, 155:51–57.
- Savoury C.D.J. et al., 2011. *Monthly Notices of the Royal Astronomical Society*, 415(3):2025–2041.
- Scargle J.D., 1982a. *The Astrophysical Journal*, 263:835. ISSN 0004-637X.
- Scargle J.D., 1982b. *The Astrophysical Journal*, 263:835. ISSN 0004-637X.
- Schlafly E.F. and Finkbeiner D.P., 2011. *The Astrophysical Journal*, 737(2):103. ISSN 0004-637X.
- Schlegel D.J., Finkbeiner D.P., and Davis M., 1998. *The Astrophysical Journal*, 500:525–553. ISSN 0004-637X.
- Schreiber M.R., Zorotovic M., and Wijnen T.P.G., 2016. *Monthly Notices of the Royal Astronomical Society*, 455:L16–L20. ISSN 0035-8711.

- Shakura N.I. and Sunyaev R.A., 1973. *Astronomy and Astrophysics*, 24:337–355.
- Shappee B.J. et al., 2014. *The Astrophysical Journal*, 788:48.
- Shen K.J., 2015. *The Astrophysical Journal Letters*, 805:L6.
- Sion E.M. et al., 2011. *The Astrophysical Journal*, 741(1):63.
- Skillman D. et al., 1999. *Publications of the Astronomical Society of the Pacific*, 111(764):1281–1291. ISSN 0004-6280.
- Smak J., 1967. *Information Bulletin on Variable Stars*, 182:1. ISSN 0374-0676.
- Solheim J.E., 2010. *Publications of the Astronomical Society of the Pacific*, 122(896):1133–1163.
- Solheim J.E. and Yungelson L.R., 2005. *Astronomical Society of the Pacific Conference Series*, 334:357.
- Spark M.K. and O’Donoghue D., 2015. *Monthly Notices of the Royal Astronomical Society*, 449(1):175–190. ISSN 0035-8711.
- Stanway E.R. and Eldridge J.J., 2018. *Monthly Notices of the Royal Astronomical Society*, 479:75–93. ISSN 0035-8711.
- Steehls D. et al., 2006. *The Astrophysical Journal*, 649:382–388. ISSN 0004-637X.
- Strohmayer T.E., 2002. *The Astrophysical Journal*, 581:577–584. ISSN 0004-637X.
- Strohmayer T.E., 2004. *The Astrophysical Journal*, 610:416–421.
- Strohmayer T.E., 2005. *The Astrophysical Journal*, 627:920–925.
- Szkody P. et al., 2012. *The Astrophysical Journal*, 753:158.
- Taylor J.H. and Weisberg J.M., 1982. *The Astrophysical Journal*, 253:908. ISSN 0004-637X.
- Tremblay P.E., Bergeron P., and Gianninas A., 2011. *The Astrophysical Journal*, 730(2):128. ISSN 0004-637X.
- Van Cleve J.E. et al., 2016. *Publications of the Astronomical Society of the Pacific*, 128(7):075002.
- VanderPlas J.T., 2018. *The Astrophysical Journal Supplement Series*, 236:16.

- Verbunt F. and Rappaport S., 1988. *The Astrophysical Journal*, 332:193. ISSN 0004-637X.
- Verbunt F. and Zwaan C., 1981. *Astronomy and astrophysics.*, volume 100. EDP Sciences [etc.].
- Vernet J. et al., 2011. *Astronomy and Astrophysics*, 536:A105.
- Wagner R.M. et al., 2014. *The Astronomer's Telegram*, 6669.
- Wang L. et al., 2017. *Monthly Notices of the Royal Astronomical Society*, 466(2):2261–2271. ISSN 0035-8711.
- Wang L. et al., 2018. *Monthly Notices of the Royal Astronomical Society*, 478:5174–5183. ISSN 0035-8711.
- Warner B., 1972. *Monthly Notices of the Royal Astronomical Society*, 159(3):315–319. ISSN 0035-8711.
- Warner B., 1995. *Cataclysmic Variable Stars*. Cambridge University Press, Cambridge.
- Warner B. and Woudt P., 2002. *Publications of the Astronomical Society of the Pacific*, 114(792):129–131. ISSN 0004-6280.
- Webbink R.F., 1984. *The Astrophysical Journal*, 277:355. ISSN 0004-637X.
- Wevers T. et al., 2016. *Monthly Notices of the Royal Astronomical Society: Letters*, Volume 462, Issue 1, p.L106-L110, 462:L106–L110. ISSN 0035-8711.
- Whitehurst R., 1988. *Monthly Notices of the Royal Astronomical Society*, 232(1):35–51. ISSN 0035-8711.
- Wood J. et al., 1986. *Monthly Notices of the Royal Astronomical Society*, 219:629–655.
- Wood J.H. et al., 1989. *Monthly Notices of the Royal Astronomical Society*, 239(3):809–824.
- Wood M.A. et al., 2002. *Monthly Notices of the Royal Astronomical Society*, 334(1):87–93. ISSN 0035-8711.
- Woudt P.A. and Warner B., 2003. *White Dwarfs: Galactic and Cosmologic Probes*, 25th meeting of the IAU, 5.

- Woudt P.A. and Warner B., 2011. *The Astronomer's Telegram*, 3705.
- Woudt P.A., Warner B., and Motsoaledi M., 2013. *The Astronomer's Telegram*, 4726.
- Woudt P.A., Warner B., and Rykoff E., 2005. *IAU Circ.*, 8531:3. ISSN 0081-0304.
- Woudt P.A. et al., 2009. *The Astrophysical Journal*, 706:738–746. ISSN 0004-637X.
- Woudt P.A. et al., 2012. *Monthly Notices of the Royal Astronomical Society*, 421:2414–2429. ISSN 0035-8711.
- Yungelson L.R., 2008. *Astronomy Letters*, 34:620–634.
- Zeilik M. and Gregory S., 1988. *Introductory Astronomy and Astrophysics*. Thomson Learning, London, 4th edition.