University of Sheffield

Photometric Follow-up of Eclipsing Post-Common-Envelope Binaries

Alexander James Brown

Department of Physics & Astronomy

A dissertation submitted in candidature for the degree of Doctor of Philosophy at the University of Sheffield

May 2023

“

From slime to ape
we learn to fly.

”

— Low Low Low - James

Contents

List of Figures

List of Tables

Declaration

I declare that, unless otherwise stated, the work presented in this thesis is my own. No part of this thesis has been accepted or is currently being submitted for any other qualification at the University of Sheffield or elsewhere. The following chapters are based on papers submitted and published during the course of my PhD:

Chapter 3 is based on: Brown, A. J., Parsons, S. G., Littlefair, S. P., Wild, J. F., Ashley, R. P., Breedt, E., Dhillon, V. S., Dyer, M. J., Green, M. J., Kerry, P., Marsh, T. R., Pelisoli, I., & Sahman, D. I. "Characterizing eclipsing white dwarf M dwarf binaries from multiband eclipse photometry" MNRAS, 513, 3050 (2022)

Chapter 4 is based on: Brown, A. J., Parsons, S. G., van Roestel, J., Rebassa-Mansergas, A., Breedt, E., Dhillon, V. S., Dyer, M. J., Green, M. J., Kerry, P., Littlefair, S. P., Marsh, T. R., Munday, J., Pelisoli, I., Sahman, D. I., & Wild, J. F. "Photometric follow-up of 43 new eclipsing white dwarf plus main-sequence binaries from the ZTF survey" MNRAS, 521, 1880 (2023)

Acknowledgments

Obviously the first thanks have to go out to Steven Parsons. You've been a class supervisor and I've learnt a lot from you over the years (particularly to never go to a CV conference). This wouldn't have been possible without the many the hours you've dedicated to guiding me along. I don't think you've done half bad for a first go!

Cheers to Vik Dhillon, for bringing me along to mount HiPERCAM, for trusting me to use your precious cameras, and for revealing your secret cep spot. I'm especially grateful for your not-so-subtle persuasion to go to La Palma for the year. I really don't know if I'd have ended up doing a PhD if it wasn't for getting bitten by the research bug while I was out there. And for that, a massive thanks go out to Dave Jones. As far as cunts go, you're a sound one. Cheers for the first paper, showing me a lot of the best food spots on La Palma, and getting me into Oceansize and Vennart.

Thanks to Stu Littlefair, whose seemingly endless stories of observing cock-ups make me feel better about any I might make. You're also both to thank and to blame for the quality of any code I write – not my fault. And thanks to the rest of the ULTRACAM team, Dave, Ingrid, James, Elmé, Martin, Matt, and Pablo for kindly observing so many of my endless ZTF eclipsers and for being on hand for any middle-of-the-night troubleshooting. I'd also like to mention Tom Marsh who sadly passed away last year. Tom always seemed happy to drop everything he was doing to try and help with any problem and was always a good laugh to be around.

Cheers to the rest of the astro PhD students and postdocs at Sheffield, past and present, who've made the office a fun place to be, particularly Rebecca, Martin, Mark, James W., Emma, Umar, Summer, and George. And thanks to the rest of the astrogroup who've made me feel so welcome here.

Finally, I need to mention my family, without whom I'd probably not have got this far. Thanks to my mum for always supporting me and making sure that I do what makes me happy in life, thanks to my dad for passing on his curiosity and fascination for the Universe and many of the things in it (I think you showing me Hale-Bopp when I was 2 is my first memory), and thanks to my sister for not letting me get too full of myself. Most importantly, thanks to our dog Billy who always gives me a good welcome when I come home.

Summary

Following the rapidly increasing interest in a wide variety of astronomical transients, the desire for large-area time-domain photometric surveys has exploded. From dedicated supernova searches to multi-messenger gravitational wave follow-up facilities to transient surveys simply designed to detect whatever they can find, the total area of the sky covered by these as well as the temporal sampling is only increasing. One of the outcomes of this is that periodically varying sources are becoming easy to find and this includes eclipsing systems – particularly those with orbital periods of order a few days or less such as the post-common-envelope white dwarf-M dwarf binaries. As such, the numbers of these eclipsing post-common-envelope binaries are increasing dramatically. Although this rapidly increasing population encompasses a wealth of information and potential insight into white dwarf and close binary astrophysics, the sheer numbers being discovered make the detailed follow-up that is required to capitalise on this sample difficult and time consuming.

In this thesis I present a new technique of characterising these eclipsing white dwarf binaries from their eclipse light curves and demonstrate it to be significantly more robust and precise than previous methods while requiring less telescope time overall. I then go on to use this new technique to measure the parameters for 33 new eclipsing systems, finding a variety of interesting examples including some with pulsating WDs, magnetic WDs, or brown dwarf companions. Finally, I summarise some of the future research that naturally follows on from the work presented here.

Chapter 1

Introduction

1.1 Preface

Whenever we go out and look up at the stars in the night sky, many of us automatically assume them to be single stars sitting in their own, otherwise empty, region of space. Perhaps surprisingly though, many of the points of light that we see in the sky and refer to as stars are binaries, triples or higher-order systems. In fact, of the ten brightest 'stars' in the night sky, three are known to be binaries (Sirius, Procyon, and Achernar), one is known to be a triple (Alpha Centauri), and two are known to be quadruples (Capella and Rigel). For solar-type stars, around 50 per cent are thought to have at least one companion (Raghavan et al., 2010) – with binaries making up the majority of these. The study of how these binary systems evolve therefore has wide-ranging implications, but in order to attempt to understand binary evolution, first we must take a look at how a star would evolve unaccompanied.

1.2 Single star evolution

The majority of a star's life is spent on the main sequence, stably fusing the hydrogen in its core into helium. For a star like the Sun, this phase may last around 10 Gyr. However, the hydrogen fuel in a star's core is not in infinite supply, and eventually, this hydrogen fuel starts to become depleted. At this point, the star begins to evolve off the main sequence, passing through a number of phases before its subsequent demise. The nature of these phases is heavily dependent on the initial mass of the star, with low-mass stars evolving slightly differently to intermediate-mass stars and those differently to high-mass stars. For the white dwarfs (WDs) that are the focus of this thesis, it is the low- and intermediate-mass stars which we are most interested in.

1.2.1 Low-mass stars

As low-mass stars ($M \lesssim 2$ M$_\odot$) begin to exhaust the hydrogen supply in their cores, hydrogen burning continues in a shell around the core which is now composed of the helium ash from core hydrogen burning. As the core is no longer undergoing fusion it begins to contract, releasing gravitational energy and becoming close to isothermal. The release of gravitational energy heats up the surrounding hydrogen shell, increasing the rate of hydrogen burning and therefore its luminosity and causing the envelope to expand slightly. The continued piling up of helium ash in the centre causes the core to exceed the Schönberg-Chandrasekhar limit – the maximum fraction of a star's mass that can exist as an isothemal core while still supporting the layers above (Carroll & Ostlie, 2017). With the core now unable to support the layers above, it contracts rapidly until it becomes degenerate, releasing more gravitational energy and, as before, causing the star to expand.

With the expansion and cooling of the envelope the opacity increases, initiating a convection zone in the outer layers which deepen as the star continues to evolve. The resulting increase in the efficiency of energy transport through the star causes it to expand further, now evolving up the red giant branch (RGB). As the star continues up the RGB, the core temperature increases until it is high enough to ignite the helium in the core. Due to the degeneracy of the core, the increase in luminosity and temperature from the helium burning is not slowed by a corresponding expansion and so the core experiences a runaway ignition known as the helium flash (Boffin & Jones, 2019). After a short time, the immense energy released by the ignition of the core is able to lift the degeneracy and bring the helium burning back under control. The, now stable, helium burning continues in the core – which becomes convective due to the strong temperature dependence of the triple-α process. The increase in effective temperature also reduces the opacity of the envelope and pushes the convection zone of the envelope back out

to the surface layers (Carroll & Ostlie, 2017). The star is now on the horizontal giant branch (HGB), moving blue-ward on the HR-diagram with nearly constant luminosity.

As helium-burning continues in the core, with hydrogen-burning occurring in a concentric shell, carbon and oxygen ash is deposited in the centre, gradually moving the helium-burning into a shell around the carbon and oxygen core. Mirroring the sub-giant and red-giant phases, the core begins to contract, heating the helium-burning shell and causing the star to expand again, now evolving up the asymptotic giant branch (AGB). As the star progresses towards the tip of the AGB, thermal pulses begin to occur. These are caused by a slowing of the hydrogen shell burning which reduces the fuel available to the helium shell below. As the helium shell cools the hydrogen shell can then contract again, reigniting and increasing the helium available for the shell below which consequently reignites. This forces the hydrogen layer to expand and cool again. As this cycle continues, the star is also releasing a strong stellar wind – with typical mass loss rates of $10^{-5} - 10^{-4}$ $M_\odot$ yr^{-1} (Boffin & Jones, 2019). This superwind eventually exposes the carbon-oxygen core which heats up and contracts, traversing the HR-diagram to join the WD cooling sequence.

1.2.2 Intermediate-mass stars

Intermediate mass stars ($2 \lesssim M \lesssim 8$ $M_\odot$) evolve slightly differently to low-mass stars. As hydrogen becomes depleted in the core, the nuclear burning stops rather than immediately continuing in a surrounding shell. With the suspension of energy production the star begins to contract, increasing its effective temperature as gravitational potential energy is liberated. Once the temperature and density in the hydrogen shell is sufficient, it ignites, and the star continues on the sub-giant branch in a similar way to low-mass stars – with the core contracting as its mass is built up, resulting in the expansion and cooling of the envelope (Carroll & Ostlie, 2017). Again, the expanding and cooling enve-

lope causes an increase in the opacity of the envelope, initiating a deep convection zone. The star ascends the RGB until the temperature and density in the core is sufficient for helium fusion. In intermediate stars, however, the core is not degenerate and so the ignition can be controlled by a corresponding expansion and cooling. This prevents a helium flash from occurring and instead the star experiences a blue loop – heating up then cooling again before exhausting the helium supply in the core. As in low-mass stars, the helium burning continues in a concentric shell around the core – now composed of carbon and oxygen ash and becoming increasingly degenerate. The star ascends the AGB, again experiencing a series of thermal pulses before losing its envelope to the dense AGB superwind and progressing bluewards across the HR-diagram to join the WD cooling sequence.

1.3 Binary star evolution

During these later stages of stellar evolution, stars expand to many times their main-sequence radii. A 1 $M_\odot$ star is predicted to inflate from its main-sequence radius of 1 $R_\odot$ to 180 $R_\odot$ on the RGB, and up to around 350 $R_\odot$ at the peak of the thermal pulses on the AGB. For a 2 $M_\odot$ star these values are 60 $R_\odot$ and 500 $R_\odot$ respectively and for a 5 $M_\odot$ star, 100 $R_\odot$ and 650 $R_\odot$ (Figure 1.1, Dotter 2016; Choi et al. 2016). For binary systems or higher order multiples, it seems reasonable to believe that these extreme expansions could impact the evolution of the system as a whole.

1.3.1 Roche lobes

In order to assess how these large post-main-sequence increases in radii would affect a given binary, it is useful to know when and how any mass transfer episodes would occur. Roche lobes are the key to understanding this. A Roche lobe is the maximum contour of equipotential (gravitational and centrifugal) that totally encloses a star (red

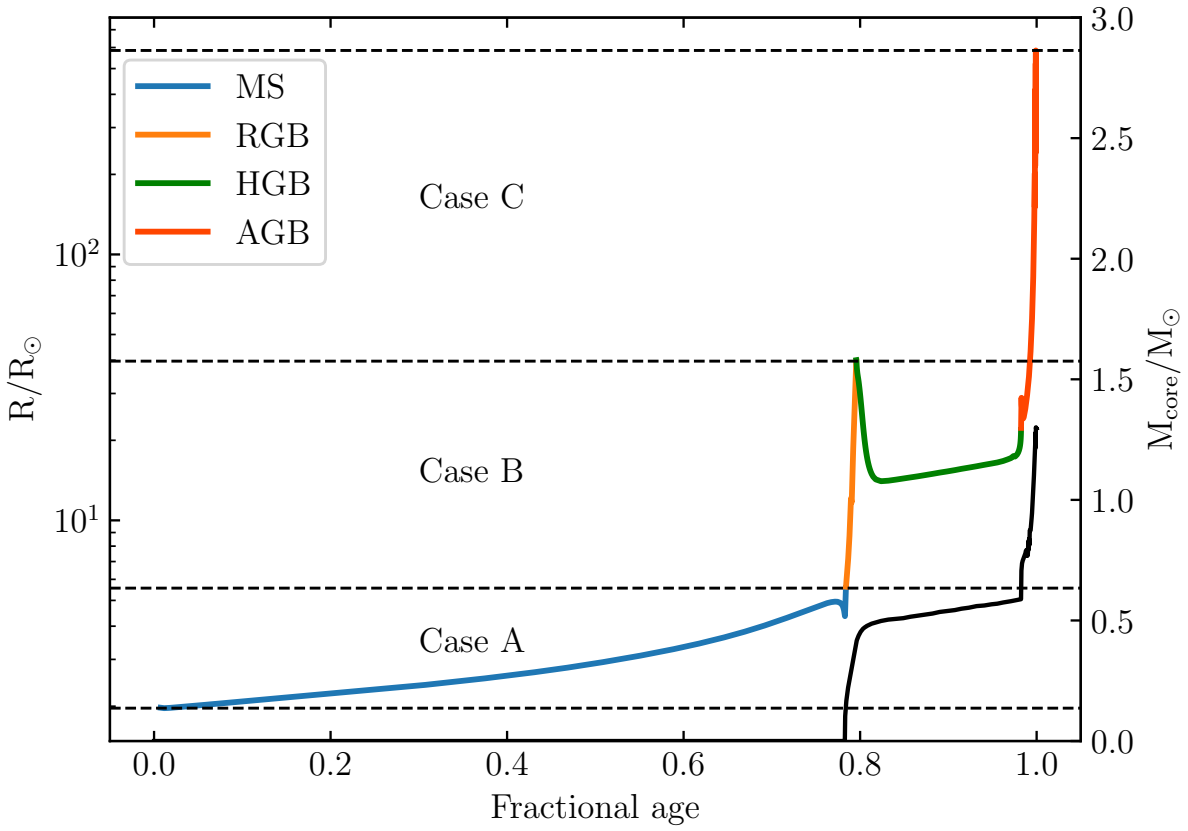

Figure 1.1: Radius evolution of a 3 $M_\odot$ star as a function of its fractional age (its age divided by its total lifetime) (Dotter, 2016; Choi et al., 2016). The phases of stellar evolution from the main sequence to the AGB are plotted with the corresponding mass-transfer cases. The total mass of the core is shown by the black line.

contours in Figure 1.2), with the two Roche lobes of a binary system forming a figure-of-eight that crosses at the first Lagrange point. Matter located within a star's Roche lobe can usually be considered as bound to that star. Outside the Roche lobe, however, it can either be bound to the binary as a whole or unbound completely. The radius of the sphere containing an equivalent volume as a star's Roche lobe, $R_{L,A}$, scaled by the orbital separation, a, depends on the mass ratio of the binary, $q = \frac{M_A}{M_B}$. An approximation for this, given by Eggleton (1983) and accurate to 1 per cent over the full range of q, is

$$\frac{R_{L,A}}{a} \approx \frac{0.49q^{2/3}}{0.6q^{2/3} + \ln(1 + q^{1/3})}. \tag{1.1}$$

As the surface of a star will typically match an equipotential surface, stars with small radii relative to this Roche lobe radius will be close to spherical. Binaries such as these where neither star fills their Roche lobe are known as detached. If a star's radius were to increase then its surface would become progressively more distorted by the gravitational influence of its companion, approaching a teardrop shape as it fills its Roche lobe. Systems in which only one component fills its Roche lobe are considered semi-detached. Finally, if both stars fill their respective Roche lobes, the system is known as a contact binary with both stars touching at the first Lagrange point.

1.3.2 Mass transfer

In the systems that are semi-detached, material from the Roche-filling star will be transferred onto the companion via the first Lagrange point. This is known as Roche lobe overflow. How mass-transfer will proceed once initiated often depends on the evolutionary phase of the Roche-filling star, being heavily influenced by its properties and response to mass-loss. Based on the radius evolution of a star over its lifetime there are three main phases in which a star is likely to overflow its Roche lobe. These are on the main-sequence (case A), ascending the RGB (case B), and ascending the AGB (case C)

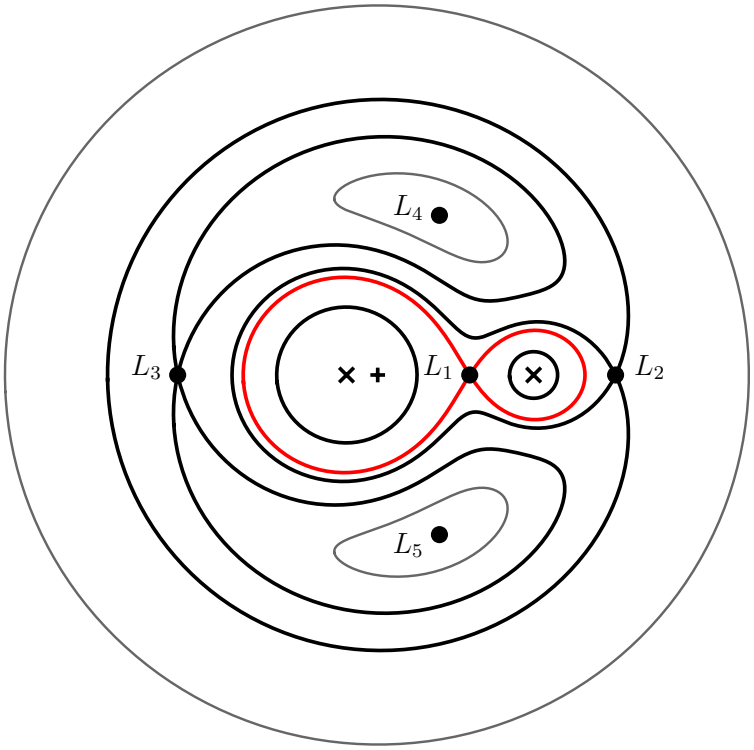

Figure 1.2: Equipotentials around a binary with a mass ratio of 0.2 (where the more massive star is on the left). Red lines denote the Roche lobes of the two stars, the crosses mark the centre of masses of the two stars, and the plus marks the centre of mass of the binary. The five Lagrangian points are also shown.

(Figure 1.1).

The ensuing mass-transfer can further be described as stable or unstable. By comparing the response of the donor's radius to mass-loss with the change in the size of its Roche lobe as a result of the mass-transfer, the stability for a particular binary can be determined. These responses can be compared using the mass-radius exponent, ζ, with the radius of the donor, R_*, being related to its mass, M_*, by $R_* \propto M_*^{\zeta_*}$, and the radius of the donor's Roche lobe, R_{RL}, following the similar relation, $R_{RL} \propto M_*^{\zeta_{RL}}$. If $\zeta_* > \zeta_{RL}$ then the mass-transfer will be self-regulating with the radius of the donor shrinking relative to its Roche lobe. This can be due to the mass ratio at the onset of the mass-transfer being such that the binary separation will increase as material moves from one star to the other or because the donor possesses a significant radiative envelope which contracts in response to mass-loss.

If $\zeta_* < \zeta_{RL}$, however, then the mass-loss from the donor will only increase the mass-transfer rate leading to a runaway unstable mass-transfer. This is often the case with donors that are on the RGB or AGB – possessing deep convective envelopes that can be approximated with a mass-radius exponent of $\zeta_* = -1/3$ (Boffin & Jones, 2019) and meaning that they expand in response to adiabatic mass-loss. Unstable mass-transfer can also be a result of the donor star being more massive than its companion at the onset of mass-transfer, with the angular momentum of any transferred material increasing as it moves outwards from the binary centre of mass and therefore removing orbital angular momentum from the orbit and causing the Roche lobe to shrink further.

1.3.3 Common envelope evolution

For many of the binary systems that are formed in the Universe, the initial orbital separation is large enough that the Roche lobes can accommodate the radius expansion of each star as they evolve. The stars in these systems will evolve as if they were

single, ending their lives as wide double WD binaries. For up to 25 per cent of binaries, however, their initial orbital separations are small enough that at some point during their evolution one of the stars will overflow its Roche lobe and begin transferring mass to its companion (Willems & Kolb, 2004). As mentioned above, this mass-transfer can be stable or unstable depending on the properties of the donor star and the binary as a whole at the onset of the mass-transfer.

For the majority of systems, the first occurrence of Roche lobe overflow will be as the more massive star – which will evolve off the main sequence first – ascends up the RGB or AGB, usually initiating an unstable mass-transfer due to both the deep convective envelope of the donor and the mass ratio of the binary. The runaway mass-transfer will occur on a dynamical timescale, much shorter than the thermal timescale of its companion, causing it to be thrown out of thermal equilibrium by this rapid pileup of material. As a result the companion will swell until both stars are filling their Roche lobes. Any further expansion of the primary's envelope will now overflow both Roche lobes, engulfing the binary in its entirety and earning the name "common envelope" (Izzard et al., 2012). Initially, the orbit begins to decay due to loss of corotation (Ivanova et al., 2013). This can either be due to the tidal forces from the binary not being sufficient to keep the expanding envelope rotating along with it, or due to dynamical instabilities such as the Darwin instability (Darwin, 1879). With the binary and envelope no longer corotating, the stars (the denser core of the primary and the still main sequence secondary) start to experience drag forces in the form of gravitational torques as they pass through the envelope. This leads to a loss of orbital energy and angular momentum from the binary and causes them to rapidly spiral in to shorter orbital periods. The orbital energy and angular momentum is transferred to the envelope during this process and results in its expansion and possible ejection. If the binary has not merged during the rapid spiral in, then the significant expansion of the envelope reduces the drag forces that act on the binary leading to a slowing of the inspiral. At this point the frictional energy

can be transported to the outer surface of the envelope and radiated away (Ivanova et al., 2013). This final phase can last hundreds of years and ends with the ejection of the envelope.

The immediate product of a non-merger common-envelope phase is most-often a short period binary composed of the hot remnant core of the primary – which will contract and cool as a WD – and its main sequence companion (Willems & Kolb 2004, I will refer to these as WD + MS binaries). In fact it was the existence of these short period binaries containing compact objects – where the giant progenitor of the compact object would not fit within the present-day orbit – that first prompted the development of the common envelope idea (Paczynski, 1976). Because the onset of the common-envelope phase will curtail the evolution of the giant star, the chemical composition of the remnant core, and therefore the resulting WD, is dictated by its evolutionary phase at the onset of the common-envelope. Mass-transfer initiated on the AGB (case C mass-transfer), after core helium-burning has begun, will mostly result in WDs composed of a mixture of carbon and oxygen (CO) with a thin hydrogen atmosphere (or even oxygen and neon (ONe) if the AGB is massive enough). Stars that filled their Roche lobes on the RGB, however, (case B mass-transfer) – before the ignition of helium in the core – will result in low-mass WDs primarily made up of helium (He) (Figure 1.1), again, with a thin atmosphere of hydrogen. In-fact, due to the long main-sequence lifetimes of stars that are not massive enough to ignite helium burning ($M \lesssim 0.6$ M$_\odot$; Dotter 2016; Choi et al. 2016), binary evolution is believed to be the only way that the He-core WDs that we observe in the Universe can form (Marsh et al., 1995).

1.3.4 Energetics of the Common Envelope

Unfortunately, the common-envelope phase is a particularly poorly understood aspect of binary evolution, with the wide range of timescales involved making it especially difficult

to simulate and the relatively short duration of the phase making it observationally rare (although possible common-envelope mergers have been seen Tylenda et al. 2011). Much of the work has focused on attempting to match simulated and observed populations, aiming to quantify the energy required to eject the envelope (i.e., its binding energy), E_{bind}, and comparing this with the energy lost from the binary orbit in a non-merger scenario, ΔE_{orb} (Paczynski, 1976; van den Heuvel, 1976; Webbink, 1984). For this purpose, the common-envelope efficiency parameter, α_{CE}, was introduced (Livio & Soker, 1988; Iben & Livio, 1993) (a similar formalism based on the conservation of angular momentum rather than energy has also been used, Nelemans et al. 2000; Nelemans & Tout 2005),

$$E_{bind} = \alpha_{CE} \Delta E_{orb}. \tag{1.2}$$

Expanding both terms and adding in an additional parameter, λ, to account for the structure of the donor star (which is of order unity) (de Kool et al., 1987) gives

$$\frac{GM_{1,i}(M_{1,i} - M_{1,c})}{R_{1,i}\lambda} = \alpha_{CE} \left(\frac{GM_{1,c}M_2}{2a_f} - \frac{GM_{1,i}M_2}{2a_i} \right), \tag{1.3}$$

where $M_{1,i}$ and $R_{1,i}$ are the mass and radius of the primary at the onset of the common envelope phase, $M_{1,c}$ and M_2 are the core mass of the primary and the mass of the secondary respectively, and a_i and a_f are the orbital separations immediately before and after the common envelope phase. Using stellar evolution models – and a magnetic braking law if necessary – an attempt to reconstruct the common envelope phase can be made, constraining the product, $\alpha_{CE}\lambda$. Zorotovic et al. (2010) used this method to constrain α_{CE} to between 0.2 and 0.3 for a large sample of WD + MS binaries found in SDSS. De Marco et al. (2011) find a slightly higher average value of 0.43 ± 0.08 assuming the efficiency is constant for all binaries observed. Similar values for the common-envelope efficiency are found for WDs with brown dwarf (BD) companions

(Scherbak & Fuller, 2023; Zorotovic & Schreiber, 2022) and population synthesis methods have also shown some preference for low values of α_{CE} (Toonen & Nelemans, 2013; Camacho et al., 2014; Cojocaru et al., 2017). Note that low values of α_{CE} mean that more orbital energy must be extracted from a binary in order to eject the envelope. The implication of this, population-wise, is that many of the systems that emerge from the common-envelope have very short orbital periods, typically ranging from a few hours to a few days. Large and representative samples of post-common-envelope binaries are key to both of these methods, allowing the common-envelope efficiency to be measured over a range of binary and stellar parameters and maybe even revealing any dependence on these such as with orbital period or mass ratio (Zorotovic et al., 2010; De Marco et al., 2011; Davis et al., 2012).

1.3.5 Post common envelope binaries (PCEBs)

It is now believed that many interesting and exotic astrophysical phenomena in the Universe have passed through this turbulent phase of binary evolution. These include but are not limited to cataclysmic variables (CVs), axisymmetric planetary nebulae, X-ray binaries, both channels to type Ia supernovae, and both the binary WD and binary neutron star gravitational wave sources – the latter now confirmed as an originator of short gamma-ray bursts (Abbott et al., 2017). It is therefore a key stage in the past evolution of a wide range of systems and, consequently, our knowledge and understanding of this phase (or lack thereof) impacts many fields of research. As the most numerous outcome of common-envelope evolution, close binaries composed of a WD and a main sequence star offer one of the best insights into this poorly understood phase. Additionally, as the progenitors of many further classes of objects and phenomena (Figure 1.3), they are an important population for in-depth study. The collation of large and representative samples of WD + MS PCEBs is therefore crucial for progress on many fronts –

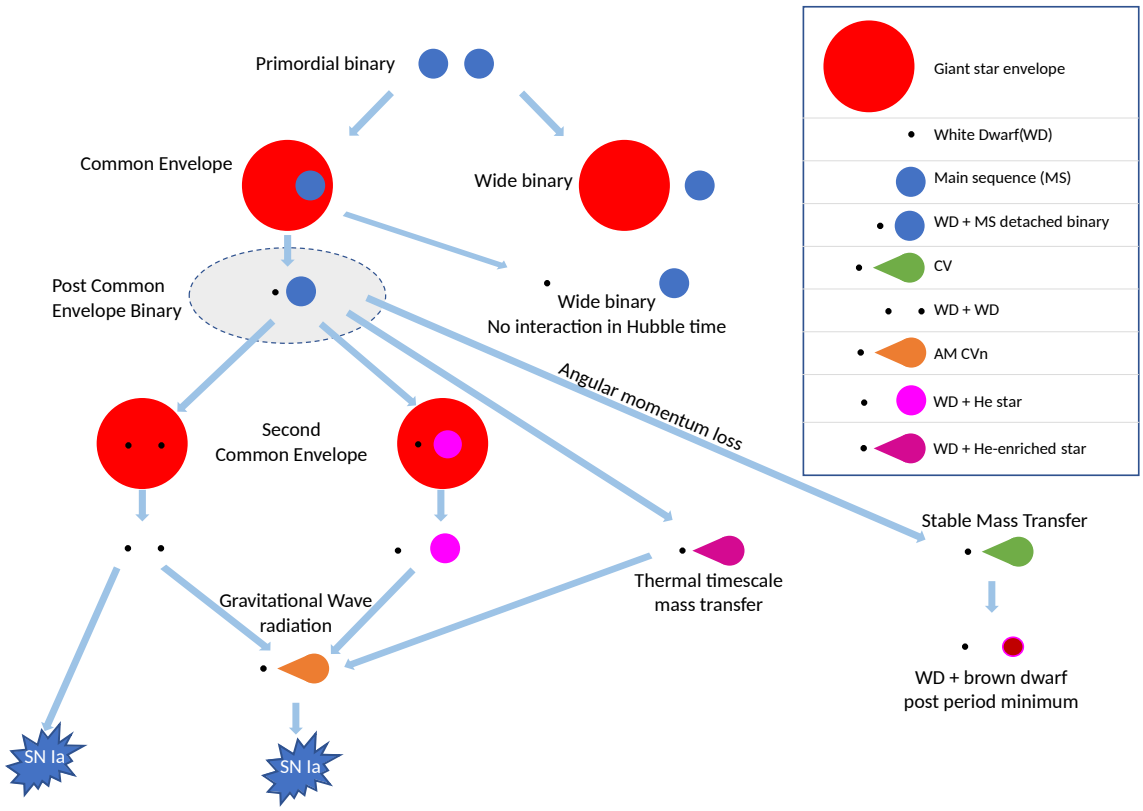

Figure 1.3: Some of the evolutionary channels leading to close WD binaries that involve a common-envelope phase. The close WD + MS PCEBs that are the subject of this thesis are highlighted by the grey ellipse. Figure adapted from Inight et al. (2021).

from investigating the common-envelope phase and linking the various populations that evolve from these binaries, to constraining the evolutionary channels that lead to the cosmologically important type Ia supernovae.

WD + MS binaries

WD + MS binaries are often split into two subgroups, mainly based on their observational properties. The first of these are WDs with late-type companions, mostly M dwarfs but usually including late K-type stars as well. These are often referred to collectively as WD + M systems and, due to the relatively low luminosities of the late-type stars, both components are usually detectable at optical wavelengths. This has allowed them to be

found in optical spectroscopy or from ground-based photometric surveys. Their relative ease of detection has meant that the WD + M systems have made up the majority of the known population of WD + MS binaries, with around 3300 being found from their SDSS spectra (Rebassa-Mansergas et al., 2007, 2010, 2012, 2016) and around 900 found from their LAMOST (Large sky Area Multi-Object fibre Spectroscopic Telescope, Luo et al. 2012) spectra (Ren et al., 2014, 2018). Of these around 25 per cent are expected to be close, and therefore PCEBs (Nebot Gómez-Morán et al., 2011), however only those that show significant radial velocity variation between spectra can be confirmed as such, resulting in a confirmed sample of over 200 post-common-envelope systems from SDSS and 78 from LAMOST (Ren et al., 2014, 2018).

Recently, all-sky space-based photometry from *Gaia* has also been used to produce a volume-limited sample of WD + M binaries, making use of a 100 pc distance limit to define any unresolved systems as close binaries and therefore PCEBs (Rebassa-Mansergas et al., 2021). This volume-limited sample totals 112 systems but, unfortunately, the data cuts required to minimise contamination to this sample from single main-sequence stars – namely, the removal of all targets redder than the blue edge of the main-sequence – also means that this sample only represents around 9 per cent of the total population. Additionally, radial velocity data is still required for these systems in order to confirm whether or not they truly are PCEBs. This would not be the case, however, for systems that are found to be eclipsing or photometrically variable, with the photometric period clearly defining them as close binaries.

The second subgroup are WDs with higher-mass solar-type companions, referred to as the WD + FGK binaries. In these systems the hotter and larger (compared to M dwarfs) solar-type stars significantly outshine the small WD component and make them difficult to discern. Fortunately, the fact that the WDs are usually a lot hotter than their companions means that they can be detected from their UV colours, standing out from single main-sequence stars on the UV colour-magnitude diagram (Parsons et al., 2016).

This method has allowed the sample of WD + FGK binaries to start to catch up with the WD + M systems with over 2000 now known (Parsons et al., 2016; Rebassa-Mansergas et al., 2017) of which around 60 have been found to be close binaries (Rebassa-Mansergas et al., 2017; Ren et al., 2020).

Characterisation of WD + M binaries

Large samples of these systems are important, but they are not particularly useful without some binary and stellar parameters to go along with them. Characterising the binaries in these samples is therefore an important task. Until now, WD + M binaries have mostly been characterised via spectral fitting methods – fitting a composite spectral model made up of a WD spectrum and an M dwarf spectrum (Figure 1.4).

For systems where both components contribute a similar proportion of the flux in the spectral range this method can work reasonably well as an initial parametrisation. However, for systems in which one component dominates spectral fitting starts to struggle, only able to return reasonable parameters for the brighter component and often failing to characterise the fainter star altogether.

Eclipsing systems

One of the significant advantages of PCEBs is that, due to their short orbital periods, and therefore small orbital separations, a relatively large fraction (around 10 per cent Parsons et al. 2013a) are orientated in such a way that the two stars occult each other as they move through their orbits. For the WD + M systems – in which the WD contributes a measurable fraction of the total flux in the optical – this results in detectable eclipses when the M dwarf passes in front of the WD. Due to the small size of the WD, these eclipses are almost always flat-bottomed total eclipses – making them easy to discern as WD + M systems and, since the WDs often contribute a large fraction of the systemic flux in the optical, the eclipses are often deep. These properties have been exploited to find

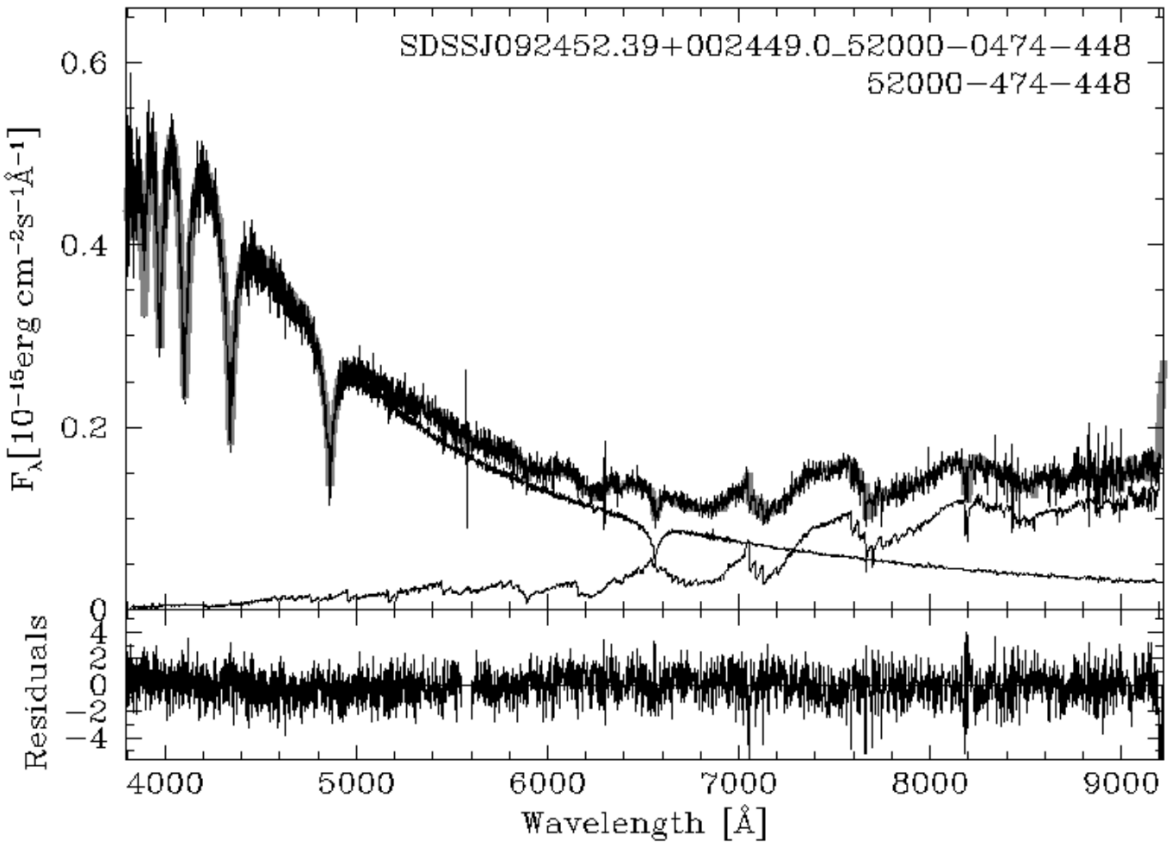

Figure 1.4: SDSS spectrum of a WD + M binary (SDSS J0924+0024) with the best-fit composite spectral model. Both the WD and M dwarf spectral models are shown by the thin black lines and the residuals to the fit are displayed in the panel below. Figure from `www.sdss-wdms.org` (Rebassa-Mansergas et al., 2012)

PCEBs in time-domain photometric surveys such as the Catalina Real-time Transient Survey (CRTS; Drake et al. 2010; Parsons et al. 2013a, 2015) with the published number currently standing at around 70 systems. However, finding these binaries in photometric surveys requires that a system is observed in-eclipse at least twice in order to confirm the presence of periodic eclipses. Fortunately, with only 200 photometric measurements there is a 99 per cent chance of measuring two separate eclipses (Parsons et al., 2013a) and therefore detecting an eclipsing WD + M system.

These eclipsing systems offer more than just an easy way to detect PCEBs, though. The sharp features of the eclipse – namely the ingress and egress of the WD – enable some of the most precise measurements of stellar radii possible in astrophysics, often at or below the 1 per cent level. Combining these radii with independent mass measurements from radial velocities allows a binary to be fully characterised, making these eclipsing systems some of the best laboratories for stellar and binary astrophysics. As such they have been used to test the mass-radius relations of WDs (Parsons et al., 2017a), measure and confirm the over-inflation of M dwarfs relative to theoretical models (Parsons et al., 2018), discern the mass at which the transition from He-core to CO-core composition occurs in WDs (Parsons et al., 2017a), as well as to investigate systems with unusual evolutionary pasts such as those that have likely formed as the result of a merger (O'Brien et al., 2001). They additionally serve as one of the few places where brown dwarfs can be found relatively easily and have their radii measured precisely (Beuermann et al., 2013; Littlefair et al., 2014; Parsons et al., 2017b; Casewell et al., 2020a; van Roestel et al., 2021), with the eclipse of the WD by the brown dwarf companion usually causing a total disappearance of the target at optical wavelengths.

Going forward

As the interest in astronomical transients has exploded recently, so too has the desire for wide-field time-domain photometric surveys with which to find and monitor them – for

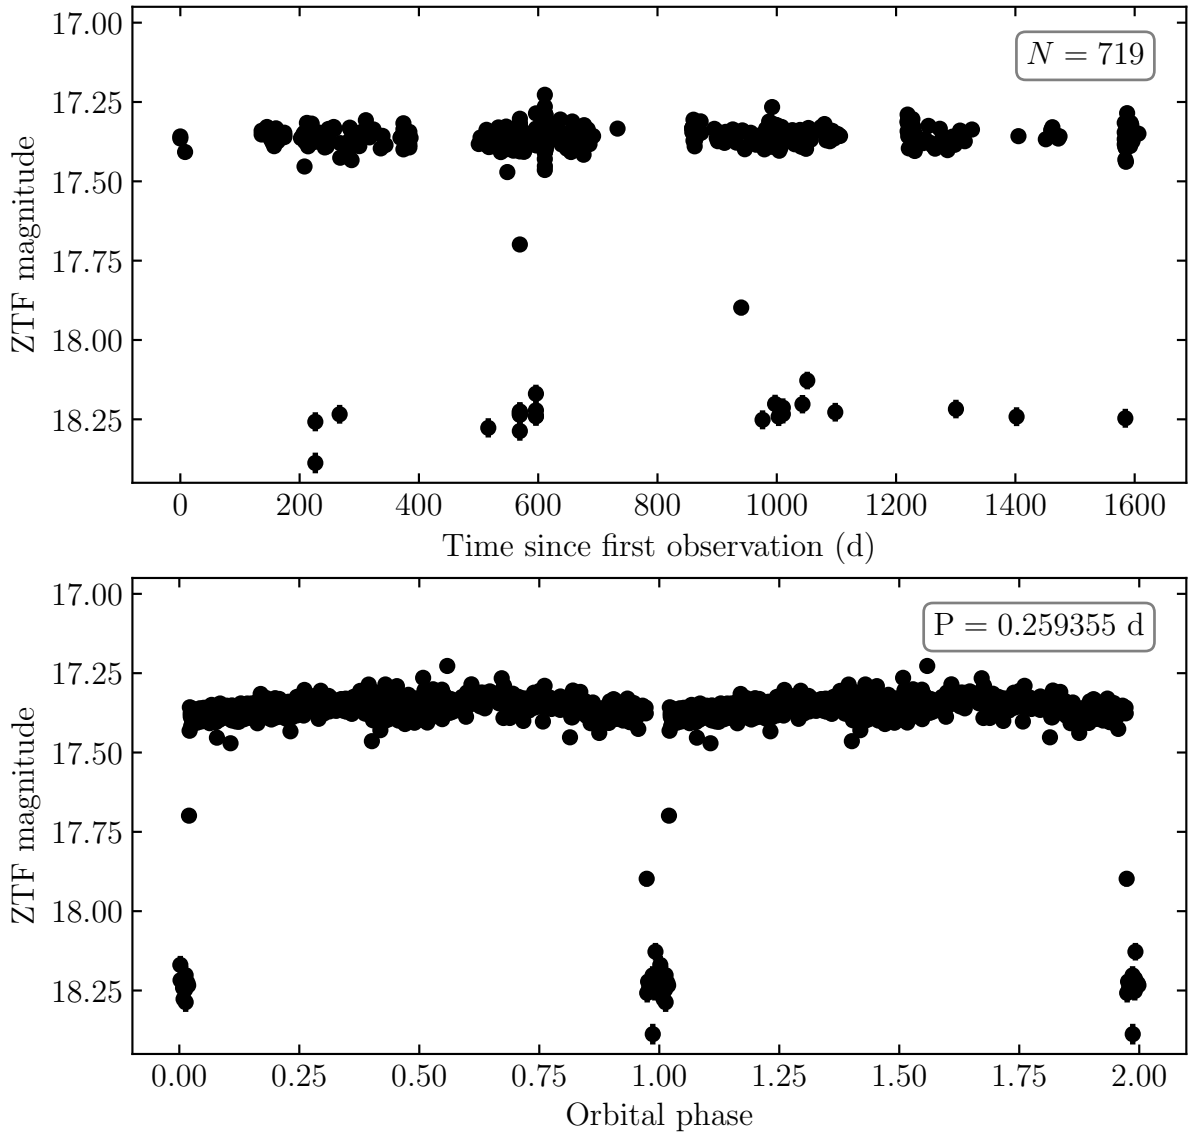

Figure 1.5: ZTF r-band photometry for the WD + M system, ZTF J0639+1919. The top panel shows the raw ZTF light curve with a significant number of points clearly fainter than the typical source magnitude, indicating the presence of eclipses. This r-band light curve consists of 719 individual photometric measurements. The bottom panel shows the same data phase-folded on an orbital period of P = 0.259355 d, with the fainter points now forming a clean eclipse and the out-of-eclipse points demonstrating a slight reflection effect.

example CRTS (Drake et al., 2009), GOTO (the Gravitational wave Optical Transient Observer; Steeghs et al. 2022), BlackGEM (Bloemen et al., 2016), PTF (the Palomar Transient Factory; Law et al. 2009), and ZTF (the Zwicky Transient Facility; Graham et al. 2019; Bellm et al. 2019; Masci et al. 2019). As such, the numbers of WD + M PCEBs found from their periodic eclipses are increasing rapidly. This increase will only accelerate with LSST (the Legacy Survey of Space and Time; Ivezić et al. 2019) and as fainter systems are discovered, the follow-up and characterisation that is required to extract meaningful results from this growing population – a population that encompasses a wealth of information on WD and binary astrophysics – will become far more challenging. Spectroscopic methods, in particular, will struggle as fainter systems are detected, requiring far more time on ever larger telescopes. A method to characterise this ever-growing population in an efficient manner, while still able to deal with fainter systems will therefore be an important tool going forward.

1.4 Summary

In this chapter I have explained the processes of stellar and binary evolution that lead to the WD + M binaries that are the subject of this thesis. I have also mentioned the importance of the common-envelope phase through which they pass and how WD + MS PCEBs exist as some of the best tracers for this relatively poorly understood phase. Finally, I have shown that the eclipsing examples of these systems are some of the best systems available to help improve our understanding of both stellar and binary astrophysics and that their subsequent follow-up and characterisation is an important challenge that must be overcome in order to make the most of the increasing population.

1.5 Thesis overview

In this thesis I will begin by outlining the key methods and techniques I have used and relied on (Chapter 2), before describing the work I have undertaken in the past three-and-a-half years, with Chapter 3 explaining the development and evaluation of a novel method with which to characterise eclipsing WD + M systems from their eclipse photometry alone, and Chapter 4 making use of this method to determine the parameters of 34 new systems found in the ZTF survey, as well as identifying some particularly interesting examples. Finally, in Chapter 5, I will wrap up with a final discussion and describe some of the future research that would naturally follow on from the work presented in this thesis.

Chapter 2

Methods and Techniques

2.1 Introduction

In this chapter I will outline the key methods and tools required for the analyses presented in later chapters. I will begin with an overview of astronomical detectors and their specific data reduction requirements and methods, before moving on to more general tools and techniques that are employed throughout the rest of this thesis.

2.2 Observations and Reductions

2.2.1 Charge Coupled Devices

Charge coupled devices (CCDs) have been the standard detector for optical astronomy since the 1980s, with the first astronomical images obtained by a CCD camera taken in 1976. Their high quantum efficiencies, high dynamic range, and linearity were a vast improvement on the photographic plates used previously.

Astronomical CCDs consist of a wafer of silicon semi-conductor with a thin layer of insulating material on top. This silicon wafer is covered in many electrodes arranged in rows across the top of the insulating layer with three rows of electrodes defining one pixel row. The pixel columns are defined by channel stops arranged perpendicular to the electrodes which separate the charge collecting areas of each pixel column. A schematic diagram of a section of a CCD chip is shown in Figure 2.1. When exposing, the central electrode row of each pixel is held at a higher voltage than the adjacent electrodes, attracting the photo-electrons towards the centre of the pixel. After the exposure has finished the collected electrons from each pixel must be measured individually, referred to as being 'read out'.

The usual analogy here involves buckets, representing pixels, on parallel conveyor belts collecting rain, representing photons (Figure 2.2). Once the rain has been collected, each parallel conveyor transfers the bottom row of buckets into the empty buckets

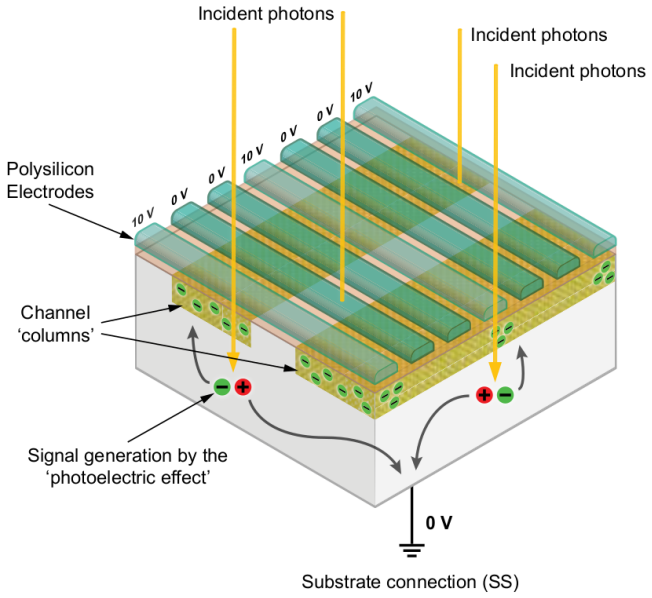

Figure 2.1: A diagram of a section of a front-illuminated CCD chip. Photons hit the chip from the top, producing electron–hole pairs in the silicon substrate which are then attracted to the nearest electrode being held at a higher voltage. Figure from Teledyne Imaging (2022).

arranged on a perpendicular conveyor. This perpendicular conveyor can then transfer the buckets one-by-one to a measurement device, measuring the contents of each bucket until all are empty. The cycle can then repeat and in this way measure the rain collected in each bucket. Leaving the analogy behind, the part of the CCD that works like the perpendicular conveyor is the serial register. The collected electrons from the bottom row of pixels are transferred to the serial register where they are passed to the amplifier pixel-by-pixel and measured. The process of moving the collected electrons is known as clocking and involves a sequence of steps in which the electrode voltages are altered in a specific pattern that passes the collected electrons down one pixel row with every clock cycle. This sequence of steps is shown in Figure 2.3. The measurement of the charge (i.e., the number of electrons) contained in a pixel consists of transferring it to a readout capacitor, inducing a small voltage across the capacitor. This voltage can then be amplified and fed to a (usually 16 bit) analogue-to-digital converter (ADC)

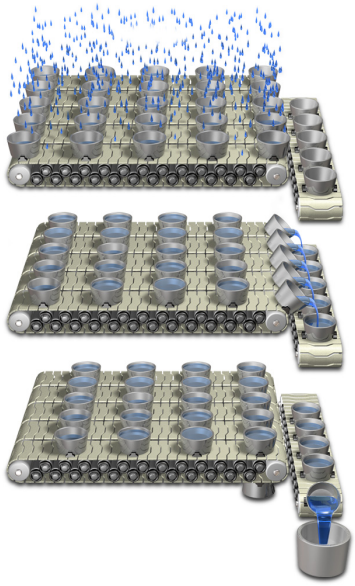

Figure 2.2: Diagram of the 'rain in buckets' analogy for CCD readout. Figure from Spring et al. (2023).

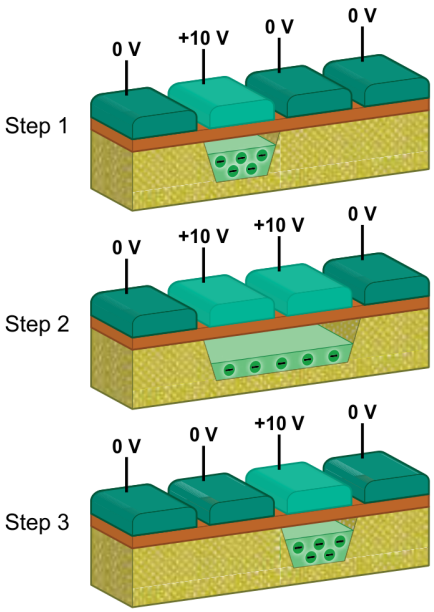

Figure 2.3: The sequence of voltages required to clock the charge to the next electrode. Figure from Teledyne Imaging (2022)

which will translate the amplified voltage to an integer value between 0 and 65535. The measurement of the amplified voltage brings with it a corresponding uncertainty, known as readout noise, which is typically around a few electrons per pixel for an astronomical CCD.

2.2.2 Frame Transfer CCDs

The process of reading out each pixel one-by-one can take an appreciable amount of time – particularly if a more precise and therefore slower readout mode is used (for example the readout time of the VIS arm of X-Shooter, a spectrograph on the VLT, can be as long as 92 s; Vernet et al. 2011). Unfortunately, typical CCD formats must have the instrument shutter closed while the previous image is reading out in order to prevent the smearing of the image as pixels are slowly clocked down the chip. This results in lost time, known as dead time. For high-speed imaging this is not ideal. If the readout

time of your detector is, for example, 10 seconds and an exposure time of 10 seconds is desired then your detector is only collecting photons for half of the time, resulting in a duty cycle of 50 per cent. Obviously this inefficiency worsens as the exposure time is reduced. For the eclipsing binaries which are the subject of this thesis it is essential that the ingress and egress of the eclipses can be resolved, and so consecutive exposures need to be as close together as possible, with typical exposure times of a few seconds. Some method of reducing the dead time due to the CCD readout is therefore needed.

The obvious way to achieve this is to reduce the number of pixels that need to be read out. This can be done by windowing the detector whereby only the pixels within a specified area of the detector are read out, with the rest being discarded. This comes along with the obvious disadvantage that much of the data in the frame is lost. If only the data in small pre-specified regions is required then this can be an acceptable compromise. Another way is to reduce the effective number of pixels by binning, combining the charge from adjacent pixels so that they act as the one pixel. In this way the charge from, for example, 4 pixels in 2x2 binning are read out all in one go. For 2x2 binning it follows that the readout time will be approximately quartered as the effective number of charge measurements is also quartered. The issue here is that the spatial information contained across those 4 pixels is lost. If matched to the atmospheric seeing so that 2 binned pixels approximately equal the seeing then no real information is lost but in good seeing conditions this is often not favourable.

Another method of reducing the readout time, and the one employed by both ULTRACAM and HiPERCAM, is known as frame-transfer. This requires a CCD detector twice the area of the desired image size. Frame transfer CCDs don't reduce the readout time as such but remove the bottleneck of needing to wait for the previous image to finish reading out before the next can start exposing. This is done by clocking the image, a relatively fast operation (7.8 ms in the case of HiPERCAM; Dhillon et al. 2021), into a storage area of the same size as the imaging area. This storage area is coated in an

opaque material – often an aluminium screen – to prevent photo-electrons from being produced in this area. The charge held in this storage area can then begin reading out while the imaging area, which is now clear, can begin exposing again. In this way, as long as the desired exposure time is more than or equal to the readout time for the selected windowing and binning parameters, then the dead time is dictated by the clock speed which can usually be considered negligible. The means that there can be effectively no delay between one exposure ending and the next beginning and the duty cycle can approach 100 per cent. Both ULTRACAM (Dhillon et al., 2007) and HiPERCAM (Dhillon et al., 2021) make use of this frame-transfer technique to enable the high-cadence imaging required by this work. For even higher cadence imaging, 'drift-mode' can be used. This significantly reduces the time needed to clock the image into the storage area and therefore allows frame-rates of ∼500 Hz with a corresponding duty cycle of ∼75 per cent. It is only required for frame-rates above ∼10 Hz, though, and so is not used in this work.

2.2.3 Photometry

Photometry, literally light-measure, is the act of carefully quantifying the amount of light detected from a source, usually in well defined spectral windows. Precise and accurate measurement of the flux of an object is the foundation on which the work in this thesis rests and as such cannot be overlooked. The process of extracting and calibrating an accurate and precise flux measurement relies on a series of steps outlined here.

Bias subtraction

Due to the statistical fluctuations caused by readout noise, the measurement of the number of photo-electrons contained in a pixel could return a negative value. Representing this negative value would require one of the 16 bits that encodes the pixel counts to be reserved for storing the sign of the number, reducing the available bits for the integer part to 15 and therefore halving the dynamic range. To prevent this, the output elec-

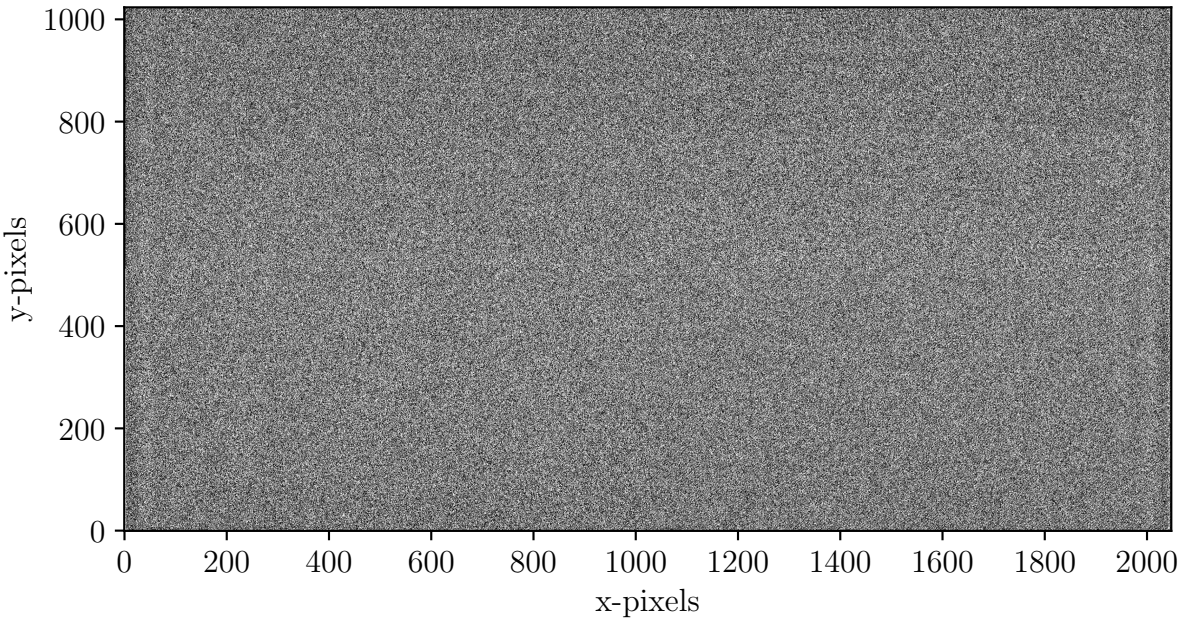

Figure 2.4: A single HiPERCAM bias frame plotted with the black point at the 10th percentile and the white point at the 90th percentile. The scatter of values is a result of noise sources in the readout electronics and is known as readout noise.

tronics of the CCD are biased positive such that a negative value is extremely unlikely. The extra bias voltage translates through the ADC as an additional number of counts on each pixel which must be measured and corrected for. Due to the readout noise, there will be some scatter around the bias level and so a series of bias frames are taken and averaged together (an example HiPERCAM bias frame is shown in Figure 2.4). The averaged bias frame can then be subtracted from all subsequent images. Additionally, the bias frames also allow the readout noise to be measured – an important quantity that defines the noise floor of the instrument.

As the aim of a bias frame is to characterise the output electronics only, the number of electrons collected in the pixels themselves need to be minimised. Bias frames are therefore taken with an exposure time as close to zero as possible (minimum exposures are limited by the time taken to clock the whole frame off the CCD) and with precautions taken to minimise the light impinging on the detector (i.e., dome lights off and the instrument shutter closed).

Dark subtraction

Even without any light illuminating a CCD, electrons can collect in the pixels of the detector. These electrons arise from the thermal energy of the detector supplying enough energy to create electron-hole pairs within the semiconductor material, with the electrons collecting in the nearest pixel. This accumulation of charge from thermal electrons is known as dark current and contributes to the measured counts in a pixel as a function of time. Firstly, this effect can be significantly mitigated by cooling the detectors, reducing the available thermal energy and therefore the chance of a thermal electron being released. Additionally, as the charge accumulated due to dark current is proportional to time, its effects can be reduced by using shorter exposure times.

If the remaining dark current in a detector is still a non-negligible source of counts (relative to the readout noise, sky background, or target) then dark frames can be taken to characterise this effect and subtract it from the science images. Since dark frames are intended to characterise the emission rate of thermal electrons, contamination from photo-electrons needs to be reduced. Therefore, as with bias frames, precautions are taken to prevent light from hitting the CCD. In the case of dark frames, though, the exposure time should be matched to the exposure times of the science images.

To subtract dark current from a science image the average bias frame is first subtracted from the dark frames, which can then be averaged and subtracted from the science images. If the dark frames have a different exposure time to the science images then the averaged dark frame can be scaled according to the ratio of exposure times before subtraction. How well this works depends on the stability of the dark current in the particular instrument though and should therefore be used with caution.

For the short exposures used in the high-speed photometry of eclipsing binaries – typically under 10 seconds – there is not enough time for an appreciable amount of dark current to accumulate and so dark subtraction is not required. Additionally, for optical

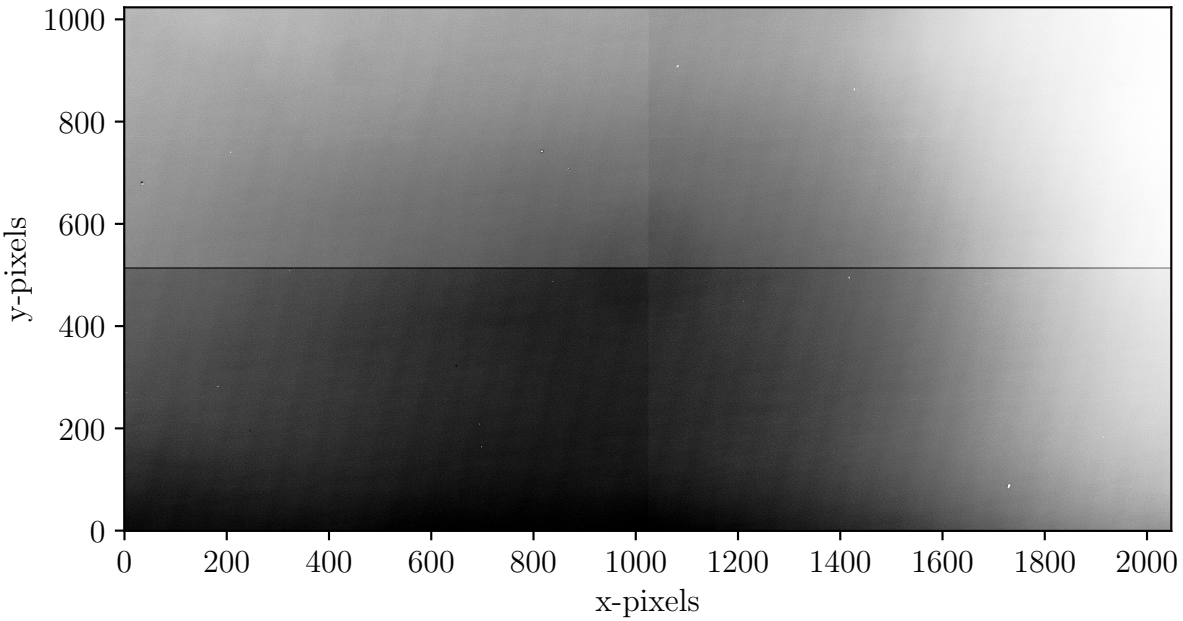

Figure 2.5: A HiPERCAM r_s-band debiased and averaged flat-field frame. The four quadrants are due to each section having its own readout electronics and, consequently, slightly different gain values

instruments, the noise due to dark current is typically insignificant compared to that from the sky background.

Flat-field correction

The effective sensitivity to photons across different pixels in a detector can be influenced by a few factors including things like dust or blemishes somewhere in the optical path and vignetting of the field as a whole, but also pixel-to-pixel variations in quantum efficiency. The consequence of this is that a star near the edge of the image may produce fewer counts that it would have if it were near the centre. In order to remove this effect, images are taken of a source known to have a uniform brightness across the field, usually the twilight sky. These are known as flat-field frames.

When taking flat-field frames, a few things need to be kept in mind. First, the exposure time should be long enough such that effects due to differential exposure times across the frame are made negligible. A mechanical shutter will block the incoming

light from one side of the detector before the other leading to slightly different exposure times. These differences in exposure time will be small fractions of a second and so exposures of at least a few seconds are enough to minimize the impact of this. For frame transfer cameras like ULTRACAM that do not possess a mechanical shutter this effect is removed, with the first row of pixels clocked onto the chip also being the first row clocked into the storage area and therefore cancelling out any differential exposure time. Second, in the case of twilight sky flats, stars will need to be removed from the flat-field images. This is performed by dithering the telescope between images, allowing them to be removed by taking the pixel-by-pixel median of a set of images.

The average bias frame is subtracted from the individual flat-field frames which can then be averaged together. The HiPERCAM pipeline (Dhillon et al., 2021) scales groups of debiased flat-field frames to the same mean value before taking the pixel-by-pixel median of the mean-scaled group. Each group median frame is then scaled such that its mean value is equal to the sum of the means of its input frames. These scaled median frames can then be averaged together and normalised to the mean value. In this way the frames with better signal are more heavily weighted. An example of the resulting master flat-field is shown in Figure 2.5. The debiased science images can then be divided by this master flat-field image, correcting the variable throughput across the detector.

Fringe correction

Many astronomical CCDs are arranged such that the electrodes are on the opposite side of the chip to the incoming light – known as back-illumination. This is an effective method of improving the quantum efficiency of the detector but requires the silicon substrate to be thinned such that shorter wavelength light is not absorbed before it reaches the light sensitive region of the chip. A side effect of this thinning process is that light at the redder end of the spectrum has a wavelength longer than the thickness of the silicon chip. This allows it to reflect inside the chip a number of times before being

absorbed. If the incident light is coherent then this can result in interference fringes. Narrow emission lines from the night sky (primarily from OH radicals in the z_s-band) can provide this coherent light source and result in images like those shown in the top panel of Figure 2.6. In the case of HiPERCAM, this is especially noticeable in the z_s band. Fortunately, as the interference pattern is caused by the characteristics of the CCD chip, it is stable over time and can therefore be corrected for.

Fringing is not removed by flat-field frames as it is produced by narrow emission lines that are not present at twilight. It therefore needs its own calibration step. This involves taking a series of fringe frames, taken with a dither pattern so that stars can be removed. These are then debiased and flat-field corrected before being median averaged together (an example HiPERCAM fringe frame is shown in the middle panel of Figure 2.6). The scale of the fringe frame can then be measured by selecting a series of peak-trough pairs. When defringing a science image, these same peak-trough pairs are used to measure the scale of the fringing in the science image which can then be used to rescale the premade fringe frame to the amplitude present in the science image. This scaled fringe frame can then be subtracted from the science images.

Aperture photometry

With some nicely calibrated science images, the next step is to extract photometry from any sources that we are interested in. The most common method for this is called aperture photometry and typically involves placing a series of three concentric circular 'apertures' around each target (Figure 2.7). The smallest of these apertures is known as the target aperture and the values of all pixels within this aperture are summed together. Pixels that are not totally contained by the aperture can be assigned a weight according to the fraction of their area located within the aperture.

The counts that have been summed from our target aperture contain contribution from the sky background. This needs to be subtracted to retrieve an accurate measure

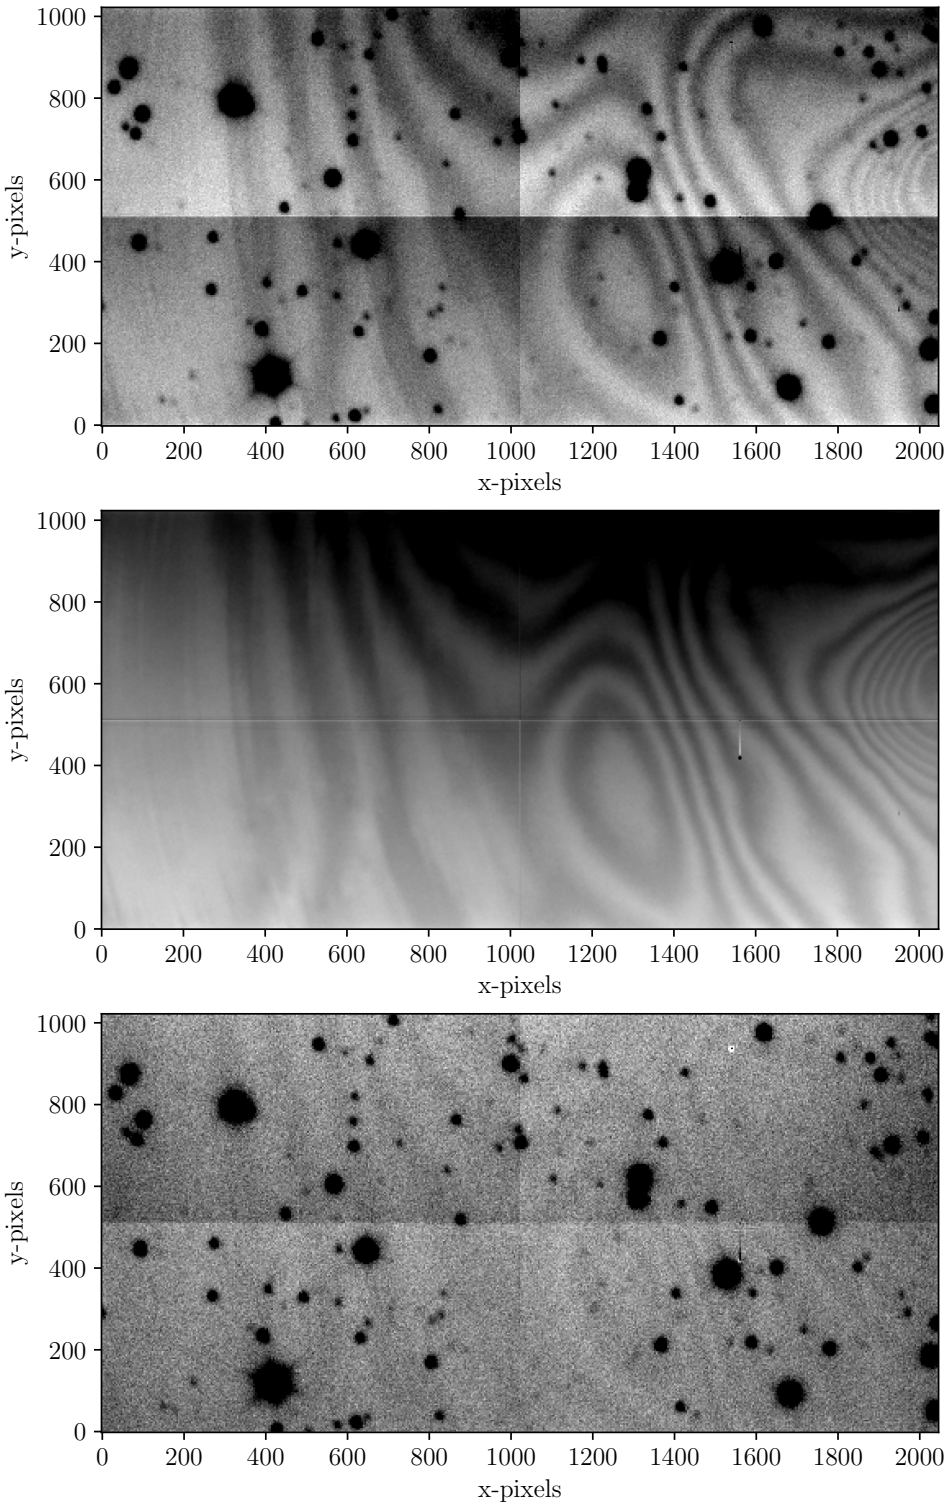

Figure 2.6: *Top*: A debiased and flat-field corrected HiPERCAM z_s-band science image. Note that the images have been scaled to accentuate the effects of fringing. The typical amplitude of the z_s-band fringing is around 1 per cent. *Middle*: A calibrated and averaged HiPERCAM z_s-band fringe frame showing the characteristics of the fringe pattern for this CCD. *Bottom*: The same debiased and flat-field corrected HiPERCAM z_s-band science image as in the top panel, but with the scaled fringe frame subtracted. The image scaling is the same as the top panel.

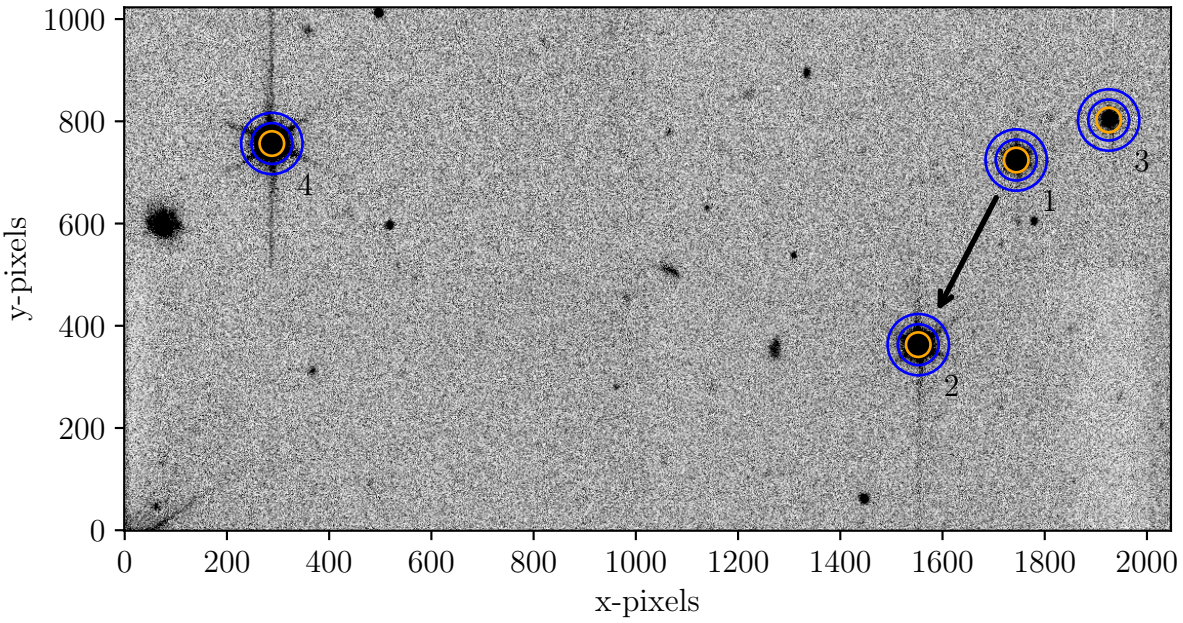

Figure 2.7: A calibrated HiPERCAM g_s-band science image with target apertures (shown in orange) and sky annuluses (shown in blue). The arrow indicates that the position of aperture 1 is linked to the position of aperture 2 so that the aperture remains centred on the target while it is in-eclipse.

of counts from our target. Assuming that the sky background is relatively constant spatially, it can be measured with another aperture close to the target and subtracted. Typically, a sky annulus centred around the target is used for this. The counts in this annulus are averaged using some method that is resistant to outliers, either a median or a sigma-clipped mean. This measure of the sky background is per square pixel and so to subtract the sky background from the target, it must first be multiplied by the number of pixels contained within the target aperture to give the total counts contributed by the sky in this area. This number is then subtracted from the target to give the total counts from the target.

The radius of the target aperture as well as the inner and outer radii of the sky annulus can be chosen to optimise the signal-to-noise ratio of the extracted flux. For the target radius this is a balance between increasing the fraction of the total counts from the target that lie within the target aperture, and minimising the additional noise

added from low-signal pixels. With the sky annulus, a larger area means a larger sample of pixels over which the average sky background is calculated, leading to a more precise measure. However, a larger annulus will also have a higher chance of including background stars, skewing the sky background measurement. It may also begin sampling sky background that is a slightly different level to that around the target, possibly due to an imperfect flat-field correction. The inner radius of the annulus should be chosen such that contribution from the extended point-spread function (PSF) of the target is negligible.

When performing photometry on targets with a low number of counts, it can be useful to link the aperture of the target with the aperture of a nearby star such that the position of the nearby star defines the position of the target by way of a known offset (shown by the arrow in Figure 2.7). This technique is often used in the reduction of the eclipsing systems, studied in later chapters, in order to keep the apertures centred during the eclipse where it is difficult to measure the centroid of the source.

Optimal extraction

Optimal extraction offers a slightly different approach to just summing the counts within the aperture, weighting each pixel's contribution according to a fitted PSF. Given a source position and a PSF, a single pixel value alone provides an estimate of the total source flux. Optimal extraction effectively calculates a weighted mean of these flux estimates using each pixel within an aperture where the weights are given by the relative flux of that pixel defined by the given PSF, and the pixel value's uncertainty. For faint targets this method can extract more of the available signal than aperture photometry, however, the benefit is negligible for brighter targets and so this extraction is only used for systems where it will produce an appreciable improvement in the quality of the light curve.

Differential photometry

For time-series photometric observations, variations in sky transparency and airmass will impact the extracted light curve. The typical method of removing these effects is to extract the flux of the target star divided by the flux of another star in the frame, known as a comparison star. The key assumption being that any transparency or airmass variations that affect the target will affect the comparison equally and therefore be removed.

When performing differential photometry it is important that a suitable comparison star is selected. The primary concern is that the comparison star should be non-variable, at least on the timescale of the observations. This can be confirmed by extracting differential photometry of the potential comparison star against a few other stars in the frame. The chosen comparison star should also be brighter than the target so that the signal-to-noise ratio of the extracted light curve is not degraded by the higher noise of a fainter comparison. A third, but usually negligible, point is that stars with different colours will be affected differently by airmass variations, with bluer sources experiencing more atmospheric extinction than redder sources. This effect is usually small but can be minimised by selecting a comparison star with a similar colour to the target. This is only a concern for very high precision photometry that covers a significant range of airmass.

2.2.4 Flux calibration

For many scientific goals differential photometry is sufficient, with the shape of the light curve providing enough information. In some cases, though, placing the photometry onto a defined flux scale is required – for example, the fitting method developed in Chapter 3 relies on the absolute flux scale to constrain the radii and temperatures of the two stars in the binary. Flux calibration is the process of accounting for the instrumental and

atmospheric effects specific to an observation, so that it can be placed on this defined flux scale.

Atmospheric extinction

As incoming photons encounter the Earth's atmosphere, some fraction are absorbed and/or scattered out of the line of sight of an observer. This phenomenon is known as atmospheric extinction and takes the form of an exponential decay with respect to airmass, X, with the number of photons scattered or absorbed being proportional to the total number of photons traversing the atmosphere. It is therefore linear in magnitudes,

$$m = k_x X + m_0, \tag{2.1}$$

where m is the instrumental or observed magnitude and the subscript, 0, indicates its above-atmosphere value, k_x is the extinction coefficient which corresponds to a specific filter, x. The instrumental magnitude – shown by the subscript i – is defined as

$$m_i = -2.5 \log(N_{\mathrm{ADU}}/\Delta t), \tag{2.2}$$

with N_{ADU} being the number of measured counts from the source and Δt being the time interval over which the counts were collected.

By observing stars over a range of airmass, the extinction coefficient can be found by fitting a straight line to their instrumental magnitudes over this airmass range. As atmospheric extinction is highly dependent on the wavelength of light, this must be performed for each filter being used.

Instrumental Zeropoint

With the atmospheric extinction measured, the photometry can be corrected to above-atmosphere values. However, this above-atmosphere magnitude is still dependent on

the throughput of the observing system and the collecting area of the telescope. These are corrected for by using an observation of a standard star, with a defined magnitude, $m_{0,std,x}$, in a specific photometric system. By comparing the above-atmosphere instrumental magnitude of the standard star, $m_{i,std,x} - k_x X_{std}$, with its catalogue magnitude, $m_{0,std,x}$, the offset – known as the instrumental zeropoint,

$$m_{zp,x} = m_{0,std,x} - m_{i,std,x} + k_x X_{std}, \tag{2.3}$$

can be found. Like the atmospheric extinction, the zeropoint is bandpass specific. This then permits the calibration of observations taken with that telescope, instrument, and filter via,

$$m_{0,targ,x} = m_{i,targ,x} - k_x X_{targ} + m_{zp,x}, \tag{2.4}$$

or in its expanded form,

$$m_{0,targ,x} = m_{i,targ,x} + m_{0,std,x} - m_{i,std,x} + k_x(X_{std} - X_{targ}). \tag{2.5}$$

For precise photometry, it is best to observe a flux standard star at an airmass as close to that of the target when it was observed as possible, minimising the impact of the atmospheric extinction coefficient.

Flux standards

There are a wide range of filter systems in use at telescopes around the world. The Johnson-Cousins UBVRI photometric system was very common until the Sloan Digital Sky Survey (SDSS) introduced their *ugriz* set (shown by the dashed lines in Figure 2.8; Fukugita et al. 1996). The SDSS filters have less overlap between filters and higher transmissions, making them increasingly popular. In order to use a flux standard to

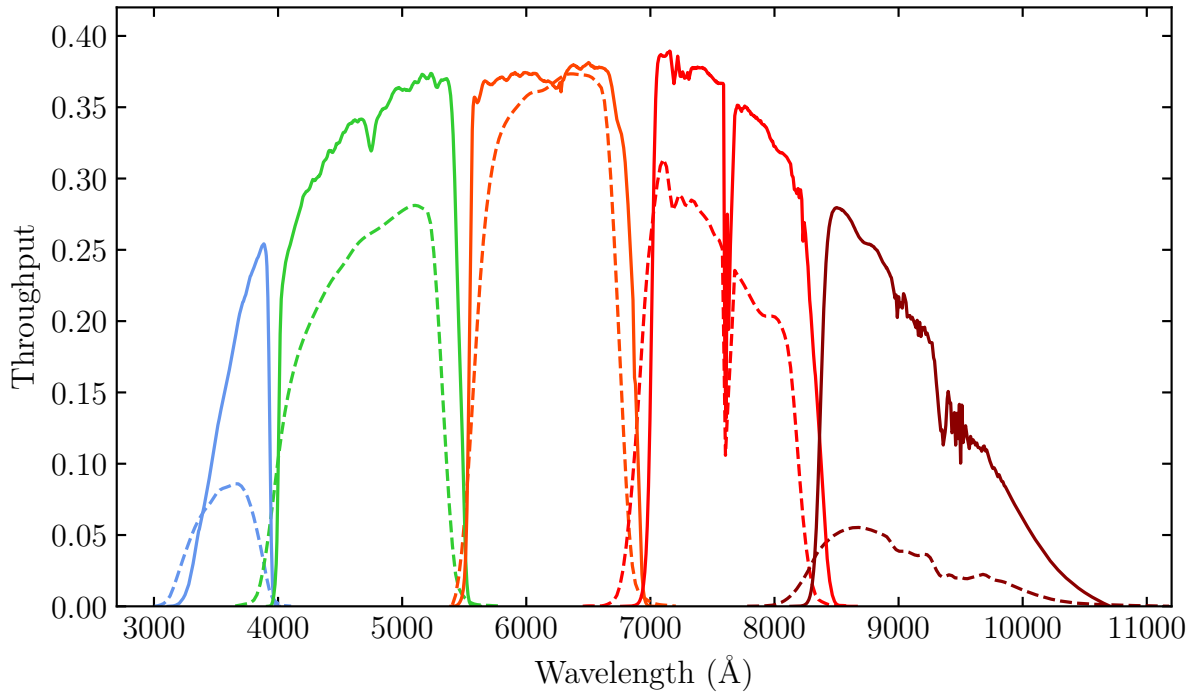

Figure 2.8: Filter profiles plotted for both HiPERCAM Super-SDSS ($u_s g_s r_s i_s z_s$) (Dhillon et al., 2021) and SDSS ($u'g'r'i'z'$) (Fukugita et al., 1996) photometric systems, both including the instrument, telescope, and atmosphere. Solid lines indicate the HiPERCAM system while dashed lines show the SDSS system.

calibrate photometry it must have a catalogued brightness in the same bandpass as the photometry. For precise photometric measurements, even small differences between filters can have a significant effect.

The Super-SDSS photometric system used by ULTRACAM and HiPERCAM (shown by the solid lines in Figure 2.8) benefits from higher throughput than the standard SDSS filters but also have slightly different bandpasses. These differences need to be taken into account when calibrating data using these Super-SDSS filters. Colour terms can often be used to correct for small differences, defining a magnitude offset as a function of colour,

$$m_{std,x} - m_{std,y} = cC_{std} + d, \tag{2.6}$$

where C_{std} is the colour of the standard star (for example $g' - r'$), c is a colour term

indicating the gradient of the magnitude difference as a function of C_{std}, and d is the magnitude offset between the photometric systems when $C_{std} = 0$. This approach can easily be applied to a flux standard to transform its catalogue magnitudes from one system into the photometric system being used. Unfortunately, it can sometimes struggle with certain spectral types of star, or in spectral regions with absorption lines that may be strongly dependent on other parameters – for example the Balmer jump with surface gravity. In these cases, it can sometimes be best to produce a catalogue of standard stars with magnitudes defined in the specific photometric system. This is a relatively easy process if the bandpass, $S_x(\lambda)$, is known and a flux calibrated spectrum covering the wavelength range of all filters, $f_\lambda(\lambda)$ or $f_\nu(\lambda)$, is available. The photon-weighted mean flux is then

$$\langle f_\lambda \rangle_x = \frac{\int f_\lambda(\lambda) S_x(\lambda) \lambda d\lambda}{\int S_x(\lambda) \lambda d\lambda}, \tag{2.7}$$

for a spectrum in per-unit-wavelength units, and

$$\langle f_\nu \rangle_x = \frac{\int f_\nu(\lambda) S_x(\lambda) \lambda^{-1} d\lambda}{\int S_x(\lambda) \lambda^{-1} d\lambda}, \tag{2.8}$$

for a spectrum in per-unit-frequency units. These fluxes can then be transformed into a chosen magnitude system. This process of generating synthetic photometry for a catalogue of standard stars has been performed for the ULTRACAM and HiPERCAM Super-SDSS systems and the resulting standard star catalogue is included in Appendix A.

Magnitude systems

In the past, most astronomical observations made use of the Vega magnitude system. In this system, the star, Vega, is defined to have an apparent magnitude of 0.03 mag in all bands (originally its apparent magnitude was defined to be 0 mag). Tying a magnitude system to an easily observable and apparently non-variable star was a practical idea,

however as telescopes and instruments have become more sensitive, zeroth-magnitude stars have become too bright to observe, negating the main benefit of the Vega magnitude system. Additional drawbacks of this system are that Vega does not have a flat spectral energy distribution and so a colour of 0 mag can represent a significant departure from the relatively constant flux it suggests. This can be misleading, particularly in the blue and ultraviolet where Vega's spectral energy distribution is particularly steep, exacerbating this effect.

Other magnitude systems have since been developed that remove this misleading colour effect, one of these being the AB mag system. In this system, a colour of 0 mag represents a constant flux per unit frequency. In absolute terms it is defined as

$$m_{AB} = -2.5 \log(f_\nu) - 48.60, \tag{2.9}$$

where f_ν is the spectral flux density in units of $\mathrm{erg\,s^{-1}\,cm^{-2}\,Hz^{-1}}$, or

$$m_{AB} = -2.5 \log(f_\nu) + 8.90, \tag{2.10}$$

for f_ν in units of Janskys. With the exception of magnitudes from the *Gaia* survey, all magnitudes given throughout this thesis will be in the AB magnitude system.

2.3 Tools

With the data reduced and calibrated, the resulting light curves can then be modelled and analysed. Some general tools are required for this stage, including light curve modelling routines as well as some statistical techniques and algorithms to enable the fitting of the light curve models to the data. These modelling and fitting tools and techniques are covered here.

2.3.1 LROCHE

LROCHE is a routine, distributed as part of the LCURVE software (Copperwheat et al., 2010), that can model the light curve of a binary system. The major benefit of LROCHE is that, unlike a lot of other light curve modelling codes, it is written to take Roche geometry into account – important for correctly modelling the close PCEBs studied here.

LROCHE models a star as a collection of flat elements on a surface, each with an area, orientation, and blackbody temperature – together defining a flux. Calculating the total observed flux is as simple as summing the flux of the visible elements, each scaled by a given limb-darkening law for its angle to the line of sight. By modelling two stars like this, with all elements lying on a gravitational equipotential, the light curve for a Roche-distorted eclipsing binary can be generated.

The surface on which the points lie is defined by the radius of each star measured from its centre of mass towards the first Lagrangian point, L_1, scaled by the orbital separation between the two stars, a. As LROCHE works in these scaled units, no masses need to be provided and the mass ratio, $q = M_2/M_1$, is enough to define the gravitational equipotentials that in-turn define the stellar surfaces. LROCHE can additionally take gravity darkening and reflection effect into account, both important variability sources for these close binaries. A summary of the parameters required to define an LROCHE model is provided in Table 2.1.

Radius-inclination degeneracy

An important thing to note when modelling binary eclipse light curves is that there is a strong degeneracy between the radii of the two stars and their orbital inclination. If we take a binary initially inclined at 90 degrees and then reduce the inclination, the length of the chord that the WD travels behind the main sequence star will be reduced.

This shortening of the chord length can be compensated for by increasing the radius of the main sequence star. Reducing the binary inclination will also have the effect of increasing the ingress and egress times of the eclipse as the WD is now being eclipsed by an edge which is angled to its effective direction of travel. This can also be compensated for though by reducing the radius of the WD (see Figure 2.9).

Breaking this degeneracy is key to being able to constrain a binary from an eclipse light curve. There are several different methods that can be used to do this including measuring the gravitational redshift of the WD – which is sensitive to its radius – or measuring the rotational broadening of the secondary star which is sensitive to the secondary radius. Unfortunately, both of these methods rely on high quality spectroscopic observations.

If only photometry is available then the depth of the secondary eclipse can be used instead, with the amount of flux obscured by the WD as it transits the secondary being simply the ratio of the areas of the two stars, $(R_{WD}/R_{MS})^2$, and therefore providing the additional information necessary to break the degeneracy. One of the issues with this method, however, is that the small size of WDs relative to main sequence stars means that the area eclipsed by the disk of the WD is usually not much more than 1 per cent, leading to very shallow secondary eclipses and making them difficult to measure.

Finally, the radii can be constrained using the luminosity of the two stars. If the distance to the system is known along with the effective temperatures of the two stars, then the radii can be constrained from their calibrated observed fluxes. This is the method that I will use to constrain the binaries studied in the following chapters.

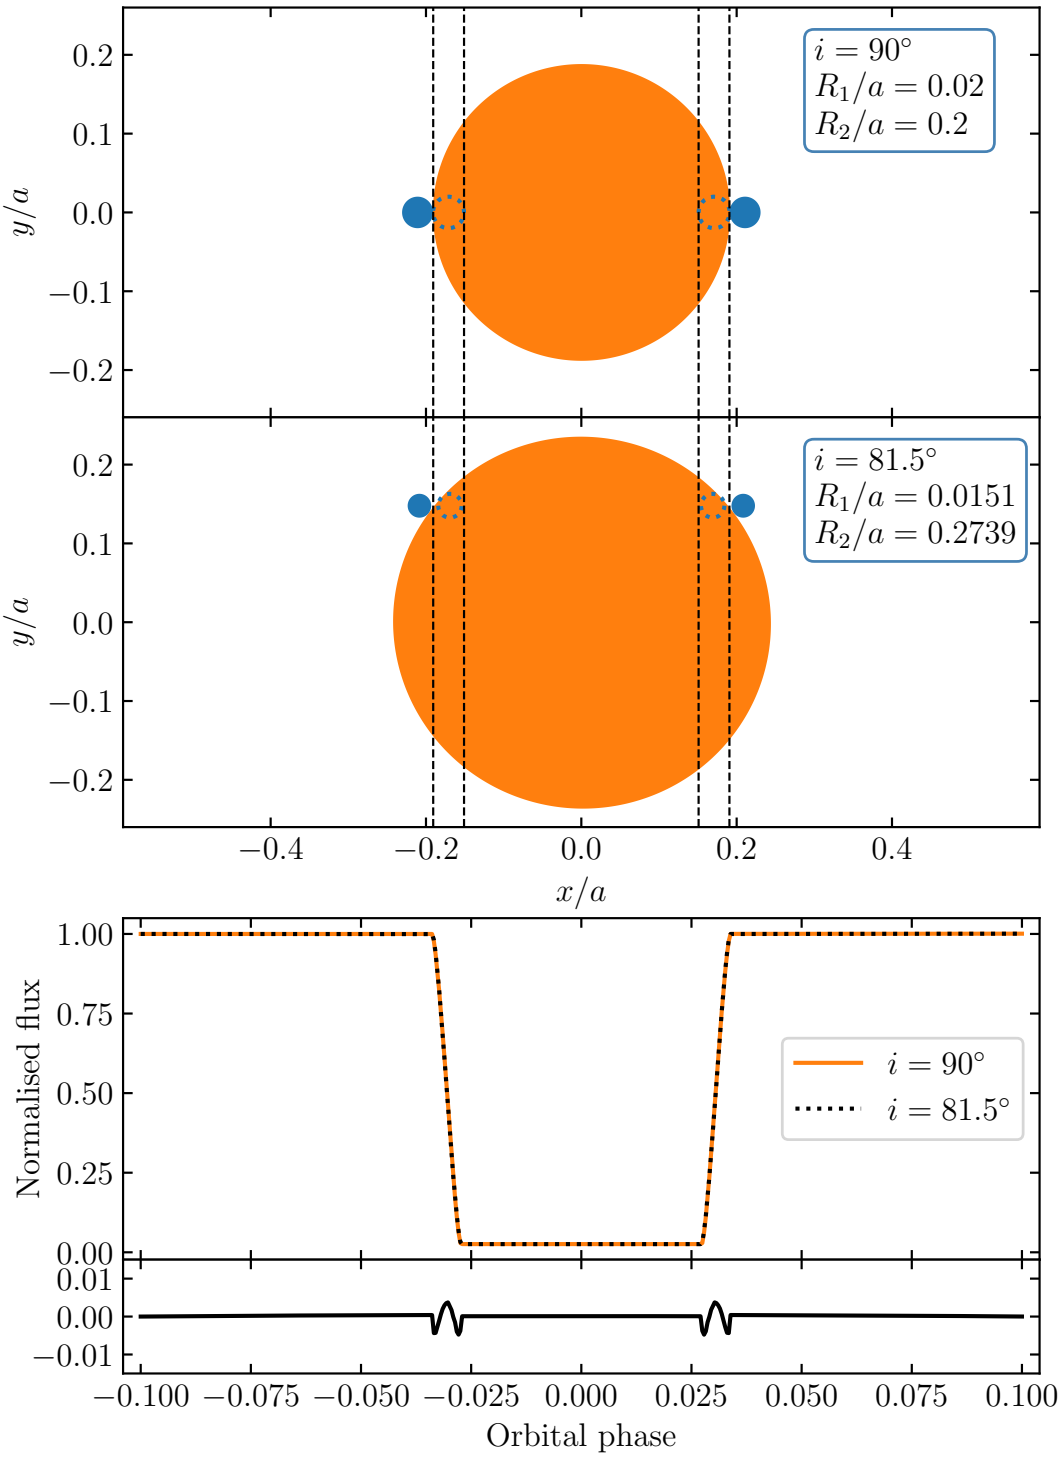

Figure 2.9: *Top*: WD (blue) plus main sequence (orange) binaries projected onto the line of sight at the four eclipse contact phases for inclinations of 90 and 81.5 degrees. The vertical dashed lines mark the inside edge of the WD. The origin is taken as the centre of mass of the main sequence star. *Bottom*: Normalised modelled eclipse light curves for the two binaries pictured above with the difference shown in the lower panel. Note that the largest difference between the two models is around 0.5 per cent – below the level that typical observations can measure.

Parameter	Description
Binary and stars:	
q	Binary mass ratio, M_2/M_1.
i	Binary inclination relative to line of sight. $i = 90$ is edge-on.
R_1/a	Scaled radius (towards L_1) of star 1.
R_2/a	Scaled radius (towards L_1) of star 2.
T_1	Blackbody brightness temperature for star 1.
T_2	Blackbody brightness temperature for star 2.
$a_{1,n}$	Limb-darkening coefficients for star 1, $n = [1-4]$.
$a_{2,n}$	Limb-darkening coefficients for star 2, $n = [1-4]$.
y_1	Gravity darkening coefficient for star 1.
y_2	Gravity darkening coefficient for star 2.
General:	
T_0	Ephemeris zeropoint, midpoint of the eclipse of star 1 by star 2.
P	Orbital period.
$\lambda_{\rm eff}$	Effective wavelength of model. Combined with blackbody temperatures to define flux.
$F_{\rm abs}$	Fraction of flux incident on star 2 from star 1 that is absorbed and re-radiated.
c_1, c_2, c_3	Coefficients of the polynomial, $c_3x^3 + c_2x^2 + c_1x + 1$, which the the fit is multiplied by to account for any additional data trends. Here, x is the time re-scaled to lie between -1 and 1.
Computational:	
$\delta\phi$	Phase accuracy of eclipse calculations.
$n_{\rm lat,c,1}$	Number of latitude strips for the coarse grid of star 1.
$n_{\rm lat,f,1}$	Number of latitude strips for the fine grid of star 1.
$n_{\rm lat,c,2}$	Number of latitude strips for the coarse grid of star 2.
$n_{\rm lat,f,2}$	Number of latitude strips for the fine grid of star 2.
$\phi_{\rm fine,1}$	Distance in phase from $\phi = 0$ where the fine grid of star 1 is used. $\phi_{\rm fine,1} = 0.05$ uses the fine grid between -0.05 and 0.05.
$\phi_{\rm fine,2}$	Similar to above but for star 2 centred around $\phi = 0.5$ and given as the phase where the fine grid should start. $\phi_{\rm fine,2} = 0.45$ uses the fine grid between 0.45 and 0.55.
`npole`	Whether the North pole of the stellar grids are true North poles (1) or substellar points (0).
`roche1`	Take Roche distortion of star 1 into account (1) or not (0).
`roche2`	Take Roche distortion of star 2 into account (1) or not (0).
`eclipse1`	Model the eclipse of star 1 (1) or not (0).
`eclipse2`	Model the eclipse of star 2 (1) or not (0).
`use_radii`	Set the stellar radii directly (1) or from contact points (0).
`gdark_bolom`	Bolometric gravity darkening law (1) or filter specific form (0).
`limb1`	Limb-darkening law for star 1 ("Poly" or "Claret").
`limb2`	Limb-darkening law for star 2 ("Poly" or "Claret").

Table 2.1: Summary of the LROCHE model parameters required to model a detached eclipsing binary.

2.3.2 Bayesian inference

For many modelling problems, it can be desirable to take our prior held beliefs or knowledge about the situation into account. In many cases this is necessary to obtain a reasonable answer. Bayesian inference is the method that allows this prior information to influence the final result – for example, I will use this method in later chapters in order to include the previously measured parallax of a system in its light curve fit.

In the context of fitting a model, θ, to some observed data, D, Bayes' rule (Equation 2.11) specifies that the probability of that model given the data, $P(\theta \mid D)$ – which is known as the posterior probability and is the quantity that we want to measure in order to assess the quality of a fit – is equal to the probability of the observed data given the model, $P(D \mid \theta)$, multiplied by the probability of the model (also known as the prior), $P(\theta)$, all divided by the probability of observing the data, $P(D)$.

$$P(\theta \mid D) = \frac{P(D \mid \theta) \cdot P(\theta)}{P(D)} \tag{2.11}$$

For the purposes of fitting a model to some data, since the data do not change, the $P(D)$ term is a constant and can therefore be ignored. The $P(D \mid \theta)$ term is often referred to as the likelihood of the model given the observed data, denoted as $\mathcal{L}(\theta \mid D)$, reducing Bayes' rule to

$$P(\theta \mid D) \propto \mathcal{L}(\theta \mid D) \cdot P(\theta). \tag{2.12}$$

Transforming this into log-space makes the calculation a simple addition, as well as making it much easier to deal with the many orders of magnitude that these probabilities can span. Conveniently, the natural logarithm of the likelihood term is simply equal to $-\frac{1}{2}\chi^2$, where

$$\chi^2 = \sum_i \left(\frac{y_i - \theta(x_i)}{\sigma_i} \right)^2, \tag{2.13}$$

leaving

$$\ln(P(\theta \mid D)) \propto -\frac{1}{2}\chi^2 + \ln(P(\theta)), \tag{2.14}$$

which is simple to calculate. It is clear from Equation 2.14 that with flat priors – i.e., the prior probability has no dependence on the model parameters – maximising $\ln(P(\theta \mid D))$ is equivalent to minimising the χ^2.

2.3.3 Markov Chain Monte Carlo

For many of the high-dimensional models in astrophysics, typical least-squares fitting methods such as the Levenberg-Marquardt algorithm (Levenberg, 1944; Marquardt, 1963) can struggle, or worse still, give incorrect results. In these cases, Markov Chain Monte Carlo (MCMC) methods are often used to sample from the posterior probability distribution. These get their name from their probabilistic nature, drawing random samples as a means of approximating a distribution (with the "Monte Carlo" part of the name referring to a casino in Monaco). MCMC methods are based around the idea of a Markov chain – a stochastic process where each new state only depends on the state preceding it, with no memory of the states before. In the case of MCMC fitting, one or more Markov chains (usually referred to as walkers) are initiated in the parameter space with their positions each having a certain probability. New positions in the parameter space are proposed, evaluated and then either accepted or rejected based on their relative improvement (or deterioration). In this way, samples can be drawn from a given probability distribution with the sample distribution being representative of the probability distribution. Since improvements are always accepted, the walkers will naturally migrate towards higher probability regions of the parameter space – although it must be noted that for distributions with local maxima, MCMC methods can still become stuck.

In these situations, other methods such as nested sampling (Feroz et al., 2009; Buchner et al., 2014) can be better suited for characterising multi-modal distributions. One of the main drawbacks of MCMC methods is the computational expense, with many draws from the probability distribution – and therefore model executions – being necessary to find and characterise the posterior distribution. This can sometimes mean that a fit can become prohibitively slow.

Metropolis-Hastings MCMC

The Metropolis-Hastings algorithm is the classic MCMC algorithm. In this method, each subsequent walker position, Y, is proposed from a probability distribution centred on the current position, $X(t)$. This probability distribution, $Q(Y; X(t))$, is known as the transition distribution and often takes the form of a multivariate Gaussian with a covariance that is optimised to give the best performance for the specific problem. The probability of accepting a proposed position is

$$\alpha = \min\left(1,\ \frac{P(Y|D)}{P(X(t)|D)}\frac{Q(X(t);Y)}{Q(Y;X(t))}\right), \tag{2.15}$$

which is decided by drawing a value between zero and one from a uniform random number generator – accepting a proposed position if the random number is less than α and rejecting it if it is more. If a proposed position is rejected then the current position is used again. Issues with this algorithm are that the transition distribution can require careful tuning to optimise the performance so that the fraction of accepted moves is not too high or too low (values between 0.2 and 0.5 are considered good). Additionally, as the transition distribution is set at the start of an MCMC fit, it cannot adapt as the covariance between parameters change in different regions of the parameter space. A more modern implementation of MCMC that solves some of these issues is known as affine-invariant MCMC.

Affine-invariant algorithms

Affine-invariant MCMC algorithms make use of the walker distribution as it stands on the current iteration to suggest subsequent proposed positions, with 'affine-invariance' meaning that the MCMC algorithm is immune to linear transformations and translations in the parameter space. By making use of the current ensemble of walkers, the covariance between parameters at the current position is naturally taken into account, significantly improving efficiency and requiring no careful tuning. It is this affine-invariant type of MCMC that I will use in later chapters to fit the light curves of eclipsing binaries.

The typical method of proposing a new position is known as the stretch-move. For each walker, X_k, in the ensemble, another complementary walker, X_j, is picked at random. The vector between these two walkers is then scaled by a number, Z, and applied to X_j giving the new proposed position. This can be written as

$$X_k(t) \longrightarrow Y = X_j + Z(X_k(t) - X_j), \tag{2.16}$$

where Z is chosen from the distribution

$$g(z) \propto \frac{1}{\sqrt{z}}, \tag{2.17}$$

and

$$z \in \left[\frac{1}{a}, a\right], \tag{2.18}$$

with $a = 2$ being a typical value. This distribution means that proposals are symmetric with an equal probability of a move and its reverse. This process is performed for each walker in turn with each position being accepted with the probability

$$\alpha = \min\left(1,\ Z^{n-1}\frac{P(Y|D)}{P(X_k(t)|D)}\right), \tag{2.19}$$

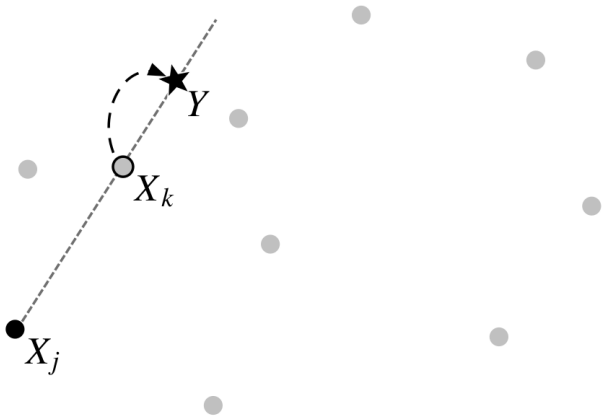

Figure 2.10: A 2-dimensional diagram of the stretch move proposed by (Goodman & Weare, 2010). For a walker, X_k, a new position in the parameter space, Y, is proposed by selecting another walker, X_j, from the ensemble at random and applying the vector between the two walker positions – multiplied by a stretch factor – to the position of X_j.

where n is the dimensions of the parameter space. As in the Metropolis-Hastings algorithm, if a position is rejected, the current position of that walker is carried forward.

As ensemble MCMC methods rely on the ensemble of walkers to naturally take the covariance into account when proposing new positions, a greater number of walkers is required than for the Metropolis-Hastings algorithm. Typically a few hundred walkers are used, but the number that is necessary can depend on the shape and complexity of the probability distribution.

As the result of an MCMC fit is a direct sampling of the probability distribution, the resulting best-fitting values and uncertainties for each parameter are often given by percentiles of the marginalised distributions rather than a mean and standard deviation. Typically, the 16th, 50th, and 84th percentiles are given for each parameter. For a marginalised distribution that is well-represented by a Gaussian, these percentiles correspond to the lower standard deviation, mean, and upper standard deviation of the Gaussian, but, importantly, remain useful even for non-Gaussian distributions.

2.3.4 Gaussian Processes

In certain modelling scenarios, there may be some correlated process affecting the data that needs to be accounted for in order to achieve a reasonable fit – for example, some correlated noise or variability that is not included in the base model. Gaussian processes are often a good way of accounting for this correlated noise or variability.

Gaussian processes are a powerful statistical tool that can be used to model data without requiring a functional form to be specified, instead being more of a probability distribution over an infinite number of possible functions defined by a multivariate Gaussian distribution. They are defined entirely by their covariance matrix with a mean of zero. Gaussian Processes are inherently Bayesian, with a chosen kernel, K, representing a prior distribution over functions. This kernel is evaluated pairwise for a sequence of x values to produce a covariance matrix, $\mathbf{K}$, which encodes the strength of the correlation between any two positions, e.g.,

$$\mathbf{K} = \mathbf{K}(\mathbf{x}, \mathbf{x}) = \begin{pmatrix} K(x_1, x_1) & K(x_1, x_2) & \dots & K(x_1, x_n) \\ K(x_2, x_1) & K(x_2, x_2) & \dots & K(x_2, x_n) \\ \vdots & \vdots & \ddots & \vdots \\ K(x_n, x_1) & K(x_n, x_2) & \dots & K(x_n, x_n) \end{pmatrix}. \tag{2.20}$$

A typical kernel is the squared-exponential (or Gaussian kernel),

$$K(x_i, x_j) = \sigma^2 \exp\left(-\frac{|x_i - x_j|^2}{2l^2}\right), \tag{2.21}$$

where x_i and x_j are two given points. With this kernel, the correlation between two points drops off as a function of the distance between them, meaning points closer together are more correlated than points further apart. The parameters (known as hyperparameters in this context) and form of the chosen kernel defines how the covariance changes with distance between points, and therefore the characteristics of the possible

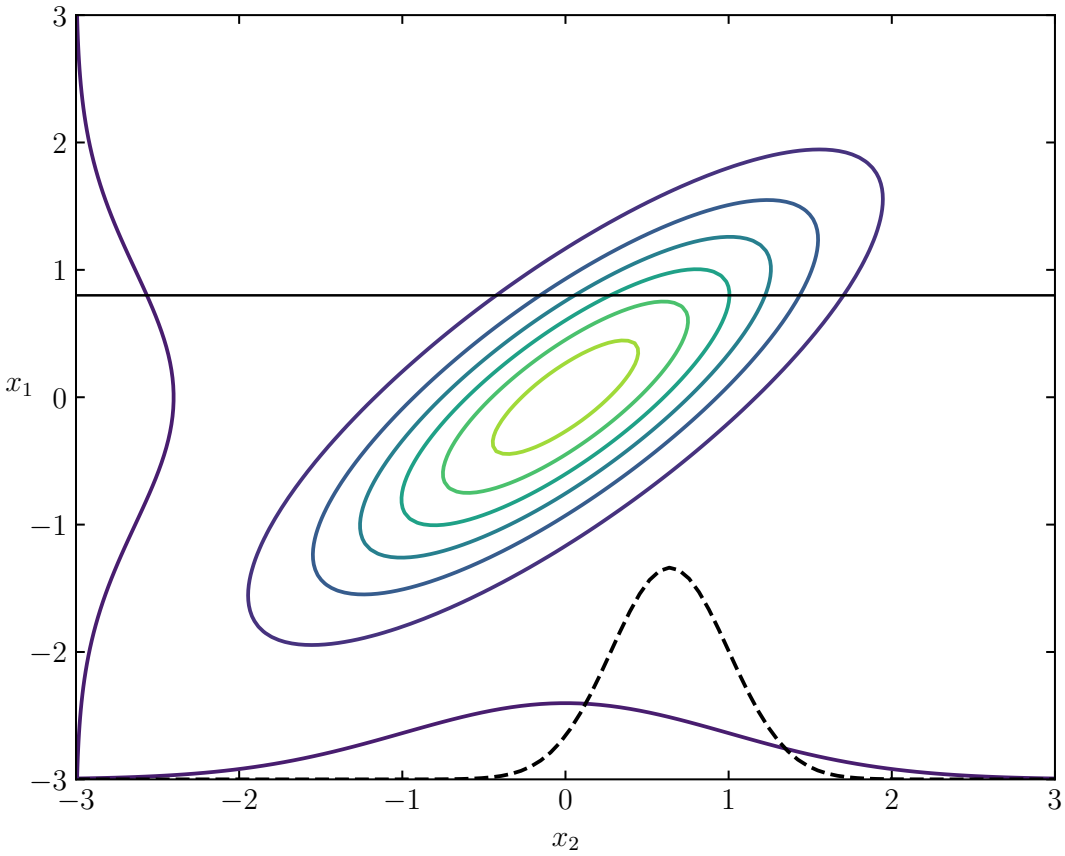

Figure 2.11: A diagram demonstrating the idea of conditioning. The contours show a joint distribution of two values, x_1 and x_2, each with a marginalised standard deviation of 1 (shown in purple on the corresponding axis) and with a covariance of 0.8 between them. The dashed line shows the distribution of x_2 given that x_1 is observed to have a value of 0.7 (shown by the horizontal black line).

function realisations that are most likely to be drawn from the Gaussian process. The hyper-parameters are often variables such as length scales (l in Equation 2.21) which correspond to smoother or sharper function realisations, or amplitudes (σ in Equation 2.21) which correspond to smaller or larger variations. By conditioning this prior covariance matrix on some given data (Figure 2.11), a posterior distribution over functions is obtained with a mean and corresponding standard deviation for each x position.

By optimising the hyper-parameters that maximise the likelihood of the Gaussian process given the observed data, the characteristics of the function realisations can be tuned to best represent the data, making them very useful for modelling functions with-

out requiring a physical model. The properties of Gaussian processes will be relied on in Chapter 4 to enable the light curve modelling of two binaries with pulsating components.

2.4 Summary

In this chapter I have provided an overview of Astronomical CCD detectors, their operation, and the reduction and calibration steps they require. I have also explained the light curve modelling software that I will use as well as the key data analysis techniques that underpin the chapters that follow.

Chapter 3

Characterising PCEBs from their eclipse light curves

3.1 Introduction

As discussed in Chapter 1, wide-field time-domain photometric sky surveys are now discovering large numbers of eclipsing WD + M systems and the rate at which these systems are found is only set to increase. An efficient method to follow these up will therefore be key to determining any population trends and finding any especially interesting systems.

In this chapter I will present and evaluate an MCMC code[1] developed to fit the parameters of these systems, namely the masses and temperatures of the WD and M dwarf components (referred to as the primary and secondary respectively), using only high-cadence multi-band photometry of the primary eclipse in combination with *Gaia* parallax measurements (Gaia Collaboration et al., 2016, 2021) and theoretical models. I will also use this code to characterise three new eclipsing WD + M binaries, finding two systems containing hot He-core WDs with low-mass companions (one near the brown dwarf transition regime) and a possible detached CV at the lower edge of the period gap.

3.2 Eclipse modelling

A multi-wavelength light curve model is constructed using the LROCHE routine (described in Chapter 2) to model the light curves of WD + M binaries immediately around the primary eclipse. The eclipse provides the strongest constraints on the system parameters and is sufficient to characterise the system without needing to observe a full orbit. Using this small region around the eclipse has the benefit of being much more efficient in terms of telescope time as well as being much faster to fit (due to a significant reduction in the number of light curve points that have to be modelled).

An LROCHE eclipse model for a detached binary is defined by 19 parameters:

1. The binary mass ratio, $q = \frac{M_2}{M_1}$.

[1] https://github.com/Alex-J-Brown/pylcurve

2. The radius of each star scaled by the orbital separation, R/a. These radii are measured from the centre of mass of the star towards the inner Lagrangian point, L_1.

3. The orbital inclination, i.

4. The equivalent blackbody temperature, $T_{\rm BB}$, of each star. These blackbody temperatures, together with the effective wavelength, $\lambda_{\rm eff}$, of the model define the monochromatic flux normal to the surface of the star via the Planck law.

5. The orbital ephemeris of the system, defined by the orbital period, P, and the time of mid-eclipse, T_0.

6. The limb-darkening coefficients of both stars, a_1, a_2, a_3, and a_4, using the four-parameter prescription (Claret, 2000).

7. The bandpass-specific gravity-darkening coefficient for the secondary star, y.

8. The fraction of incident flux from the WD that is absorbed by the M dwarf, $F_{\rm abs}$.

Many of these parameters vary with wavelength, most of which have little effect on the eclipse profile, resulting in an impractical number of free parameters when fitting multiple bands simultaneously. Additionally, degeneracies exist between some of these parameters, most notably the two scaled radii and the orbital inclination (as discussed in Chapter 2). Theoretical models and relations are therefore required in order to define these parameters from those that we are interested in – the WD masses and effective temperatures and the masses of their M dwarf companions – during the fitting procedure. In the following sections I will outline how this is achieved.

3.2.1 Mass-radius relations

The shape of the WD eclipse primarily constrains the scaled radii of the two stars and their orbital inclination. In order to retrieve masses from the eclipse photometry and break the degeneracy between scaled radii and inclination, mass-radius relations for both stars are required.

For the WD, I use the mass-radius relations of Panei et al. (2007) for He-core compositions, Bédard et al. (2020); Blouin et al. (2018); Tremblay et al. (2011) for CO-core compositions, and Althaus et al. (2005) for ONe-core compositions (Figure 3.1); this core composition must be selected for a particular fit[2]. Mass-radius relations for low- and intermediate-mass WDs have been well tested observationally and shown to be robust and accurate to within a couple of per cent (Parsons et al., 2017a). For higher mass WDs, the models are assumed to be similarly reliable, however, should this not be the case they will at least be sufficient to mark the system as containing a high mass WD worthy of further study. In terms of the outer hydrogen layer mass, models with thick hydrogen layers ($M_{\rm H}/M_{\rm WD} = 10^{-4}$) have been shown to represent WDs in PCEBs well in the majority of measured cases (Parsons et al., 2017a). I therefore use these thick layer models for all WD mass-radius relations.

M dwarfs, however, are often found to be inflated relative to theoretical models for their mass, with radii typically found to be around 5 to 15 per cent larger than models predict (López-Morales & Ribas, 2005; López-Morales, 2007; Parsons et al., 2018; Kesseli et al., 2018). In an attempt to minimise any systematic effects arising from inflation, I produce a semi-empirical mass-radius relation for M dwarfs. Masses are assigned to a sample of 15 279 M dwarfs with *Gaia* parallaxes, radii, and 2MASS K_S measurements (Morrell & Naylor, 2019) using the preferred fifth order ($n = 5$) polynomial representation of the $K_{\rm abs}$–M_* relation (Mann et al., 2019). The resulting M dwarf

[2]For systems with best-fit WD masses close to a border between core compositions it is worth running the fit again with the alternative core composition in order to determine which is most consistent with the expected mass range for the respective core composition (i.e. a 0.3 $M_\odot$ CO-core WD is unlikely).

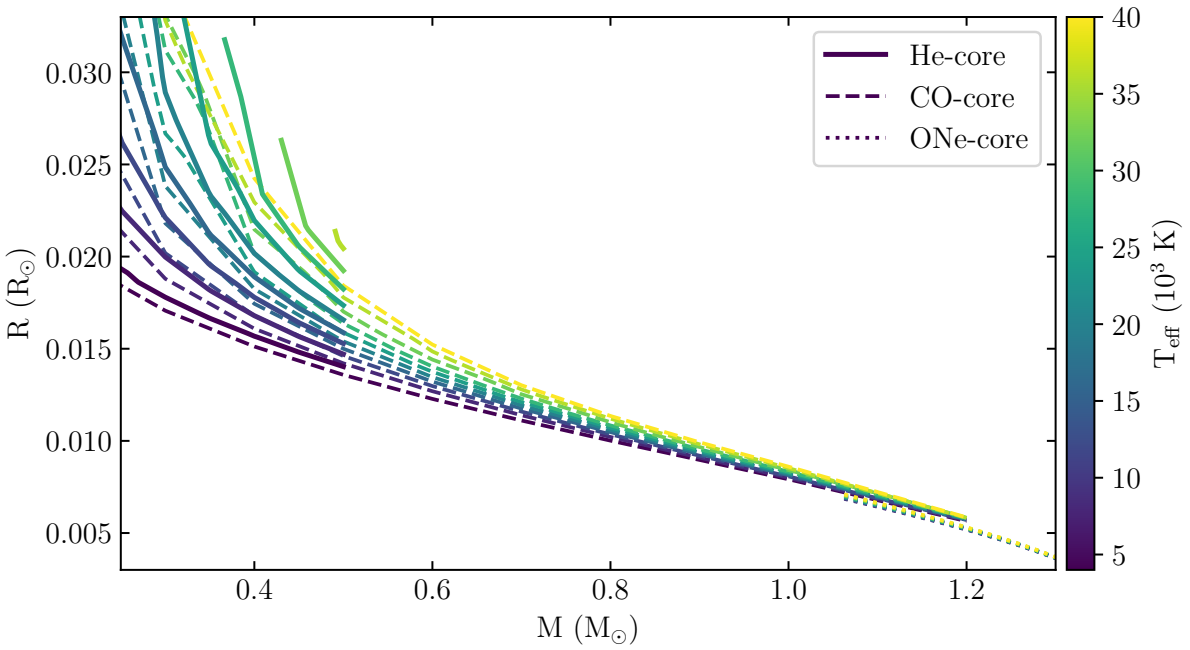

Figure 3.1: Mass-radius relations for WDs with He-, CO-, and ONe-core compositions (Panei et al., 2007; Bédard et al., 2020; Blouin et al., 2018; Tremblay et al., 2011; Althaus et al., 2005) as a function of effective temperature.

mass-radius relationship is shown in Figure 3.2. A population of stars exist above the main group. Checking these against the rest of the sample it is clear that they lie above the *Gaia* main sequence and are likely binaries or pre-main-sequence stars, explaining their anomalous radii measurements. Cross-matching these outlying points with Simbad confirms that a large proportion of these are indeed pre-main-sequence stars, variables, or binaries and can be discarded. I use an iterative sigma-clipping fitting procedure with a fifth-order polynomial to discard these points, removing around 3 per cent of the total sample. I then follow this up with an MCMC fit to retrieve the final relation (Table 3.1) while providing reliable uncertainty estimates on the polynomial coefficients. Due to the sparse nature of the sample in the low mass range and the convergence with the theoretical tracks of Baraffe et al. (2015), I switch to using the theoretical models below the mass where the semi-empirical fit crosses the models. This occurs at $M_* = 0.121$ M$_\odot$ (shown by the vertical dashed line in Figure 3.2). I use the 1 Gyr model from Baraffe et al.

(2015) below this point. Note that the apparent upturn in the sample above 0.65 $M_\odot$ is not real and is a result of the fitted effective temperatures in the M dwarf sample being limited to below 4400 K and therefore stars with slightly higher temperatures than this will require a larger radius to fit their observed luminosity. I therefore only consider the fitted relation valid below this point.

There is some scatter in the M dwarf sample around the best-fit semi-empirical relation. Much of this scatter is likely genuine variation in the radii of M dwarfs with similar masses. This is well known, with Parsons et al. (2018) measuring a scatter of ≈5 per cent in the radii of the M dwarfs in their sample. Additionally, some may be due to scatter in the fitted temperature for a given K_S-band magnitude. This is demonstrated by the gap in the sample due to the discontinuity in models at $T_{\rm eff}$ = 4000 K, described by Morrell & Naylor (2019), and which therefore describes a line of constant temperature. Additional contributions to the scatter come from K_S-band magnitude uncertainties (typically below 2 per cent) and metallicity dependence (estimated to be around 1.7 per cent, Morrell & Naylor 2019).

The fractional residuals have a standard deviation of 3 per cent. This is slightly higher when measuring the scatter as a function of radius with fractional residuals of around 3.5 per cent. This means that a fit with a hypothetical perfect determination of the secondary radius would translate into a secondary mass distribution with a standard deviation of 3.5 per cent and is therefore the maximum precision possible on the secondary mass using this relation. Additional errors in this mass-radius relation may be introduced through the $K_{\rm abs}$–M_* relation used to derive it, with Mann et al. (2019) predicting that it is able to determine the mass of a star to a precision of ≈2 per cent. They also mention that there exists a small (∼2 per cent) systematic offset for literature M dwarfs in eclipsing binary systems as compared to their predictions from the relation, possibly due to magnetic activity or rotation rates. Summing these contributions in quadrature with the 3.5 per cent scatter gives an estimated maximum precision in the

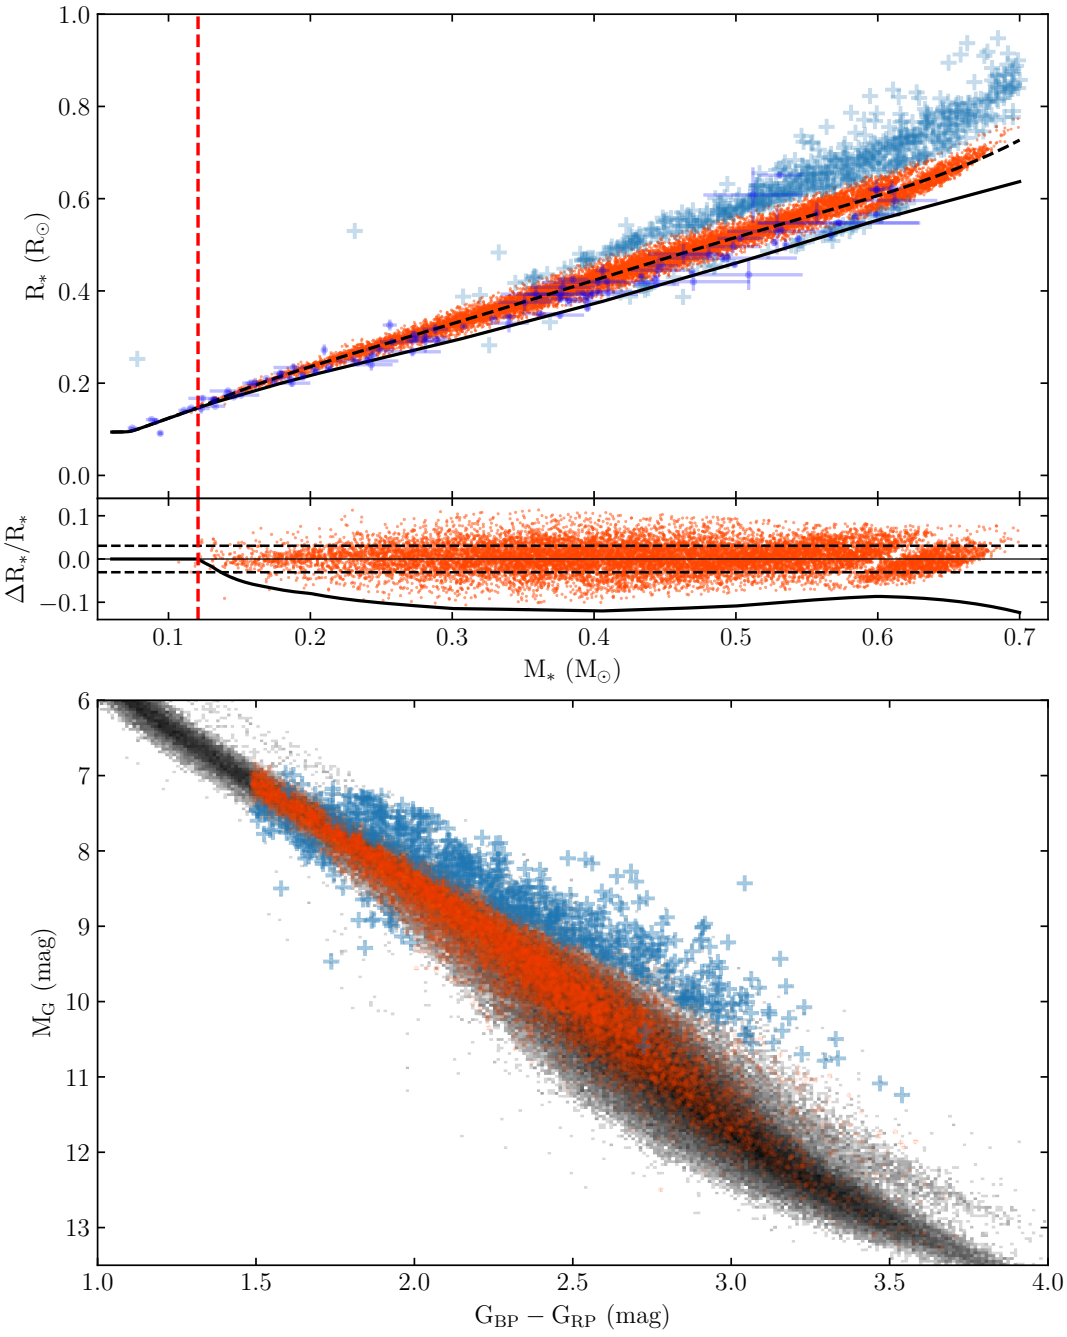

Figure 3.2: *Top*: Semi-empirical M dwarf mass-radius relation (dashed black line) alongside the 1 Gyr track of Baraffe et al. (2015) (solid black line). Red points are those that remain after the sigma clipping while the light blue crosses are those that are discarded. Dark blue points with error bars show M dwarfs with well constrained masses and radii (Parsons et al., 2018, Table A1). Fractional residuals relative to the semi-empirical relation are shown below with the dashed lines indicating $\pm 1\sigma$. The transition between the fitted relation and the Baraffe 1 Gyr model is indicated by the vertical dashed line. The gap in the data around $M_* = 0.6\ M_\odot$ is due to the discontinuity mentioned by Morrell & Naylor (2019) at $T_{\rm eff} = 4000\ K$. *Bottom*: *Gaia* Hertzprung-Russell diagram for stars within 100 pc. The M dwarf mass-radius sample is over-plotted using the same symbols as above demonstrating that the discarded points mostly lie above the main sequence.

	a_5	a_4	a_3	a_2	a_1	a_0
μ	27.4	-54.1	41.1	-14.9	3.53	-0.124
σ	1.8	3.6	2.7	0.9	0.16	0.010

Table 3.1: Best-fit coefficients and uncertainties for a fifth-order polynomial fit to the semi-empirical M dwarf mass and radius measurements of the form $\frac{R}{R_\odot} = \sum_{n=0}^{5} a_n \left(\frac{M}{M_\odot}\right)^n$.

secondary mass of ≈5 per cent. Any additional systematic contributions due to binarity (i.e. magnetic activity, rotation effects, or Roche distortion) are difficult to examine and I assume them to be small.

While this estimated uncertainty is straightforward to fold into the MCMC fitting routine, it increases the MCMC convergence time considerably, making it prohibitively long. I therefore choose to assume the best fit semi-empirical relation and account for the additional uncertainty in the relations at the end of the fitting, combining the formal errors from the MCMC in quadrature with the 5 per cent uncertainty for the secondary mass, and a 1 per cent uncertainty for the primary mass (as WD mass-radius relations have not yet been tested to higher precision than this).

3.2.2 Irradiation

Many of these WD + M binaries contain hot WDs ($T_{\rm eff} > 20\,000$ K). Given the small orbital separations in PCEBs this can result in high irradiating fluxes incident on the surface of the M dwarf, often many times greater than the typical outgoing flux from the secondary. This high irradiation can induce an inflation in the M dwarf companion and will therefore induce systematic errors in the modelling of such systems if it isn't taken into account. A simple model for this effect can be made by assuming that, for an element on the stellar surface, if the incoming irradiating flux is greater than or equal to the outgoing flux of that element then net flux emitted by it is zero. In order for the star to continue to radiate its nuclear luminosity and remain in thermal equilibrium, it must

therefore increase its unirradiated surface area in order to expel this excess luminosity (this model can increase a star's radius by, at most, ≈7 per cent; Ritter et al. 2000). Here, I assume that the effective temperature over the full surface of the secondary remains constant. The effective surface area, $s_{\rm eff}$, over which the outgoing flux is matched by an equal or greater irradiating flux is calculated according to

$$s_{\rm eff} = \frac{1}{2}\left\{1 - f_s - \int_0^{\theta_{max}} G(\theta)\sin\theta d\theta\right\} \tag{3.1}$$

(Ritter et al., 2000), where $f_s = \frac{R_2}{a}$, $\theta_{max} = \cos^{-1}(f_s)$, and

$$G(\theta) = \begin{cases} 1 - (F_{irr}(\theta)/F_0) & \text{if } F_{irr}(\theta) < F_0 \\ 0 & \text{otherwise.} \end{cases} \tag{3.2}$$

Here, $F_{irr}(\theta)$ is the irradiating flux incident on the secondary at a specified colatitude with respect to the substellar point, θ, and F_0 is the outgoing flux of the secondary, equal to σT_2^4 where σ is the Stefan-Boltzmann constant. F_{irr} is a calculated as

$$F_{irr} = \frac{\sigma R_1^2 T_1^4 h(\theta, f_s)}{a^2}, \tag{3.3}$$

where

$$h(\theta, f_s) = \frac{\cos\theta - f_s}{(1 - 2f_s\cos\theta + f_s^2)^{3/2}}. \tag{3.4}$$

$s_{\rm eff}$ can then be used to determine the inflated radius, $R_{\rm irr}$, using

$$R_{\rm irr} = R_0(1 - s_{\rm eff})^{-0.1} \tag{3.5}$$

(Ritter et al., 2000), where R_0 is the radius of the secondary without irradiation (i.e., the output from the semi-empirical mass-radius relation).

3.2.3 Roche distortion

The radii referred to in the previous sections are for isolated and therefore spherically symmetric stars. The compact nature of PCEBs mean that the Roche distortion of the secondary due to the WD can become significant and therefore needs to be corrected for. To do this I assume that the radius of a spherically symmetric star, used above, is equivalent to the volume-averaged radius of the star when experiencing Roche distortion. There is no analytical equation for calculating this correction. I therefore produce tables relating the scaled radius measured towards L_1, as used by LROCHE, to the scaled radius of a spherically symmetric star with the equivalent volume as a function of binary mass ratio. For a given mass ratio the scaled radius towards L_1 defines a Roche equipotential representing the surface of the star. I then determine the positions of points on this equipotential surface over a range of latitudes and longitudes and compute the volume of the convex hull defined by these points. The volume-averaged scaled radius can then be easily determined. The conversion then becomes a simple interpolation given the binary mass ratio and the volume-averaged scaled radius.

3.2.4 Blackbody temperatures

As previously mentioned, the temperatures used by LROCHE are a substitute for the monochromatic specific intensity normal to the surface of the star, i.e., at $\mu = \cos\theta = 1$ where θ is the angle between the line normal to the stellar surface and the line of sight. Claret et al. (2020) provide tables of specific intensities at $\mu = 1$, together with limb-darkening coefficients, a_k, for WDs in both the SDSS (Fukugita et al., 1996) and Super-SDSS (Dhillon et al., 2021) photometric systems.

For main sequence stars, no such tables exist for the Super-SDSS system. I therefore use the PHOENIX specific intensity model spectra from the Göttingen Spectral Library (Husser et al., 2013) to compute these. I calculate synthetic fluxes, $\langle f_\lambda \rangle_x$, in the Super-

SDSS system according to

$$\langle f_\lambda \rangle_x = \frac{\int f_\lambda(\lambda) S_x(\lambda) \lambda d\lambda}{\int S_x(\lambda) \lambda d\lambda}, \tag{3.6}$$

where $f_\lambda(\lambda)$ is the spectral flux density as a function of wavelength, λ, and $S_x(\lambda)$ is the throughput of the chosen filter. I do this at each value of μ supplied by the PHOENIX spectra. These can then be normalised to the flux at $\mu = 1$ and fit with the four-parameter law of Claret (2000),

$$\frac{I_\lambda(\mu)}{I_\lambda(1)} = 1 - \sum_{k=1}^{4} a_k (1 - \mu^{\frac{k}{2}}), \tag{3.7}$$

where $I_\lambda(\mu)$ is the specific intensity relative to that at $\mu = 1$. Rather than use the synthetic fluxes at $\mu = 1$ to determine the LROCHE temperatures, I calculate the total synthetic flux of the star, F_λ, by integrating the best-fit limb-darkening law around the disk of the star using

$$F_\lambda = 2\pi \int_0^1 I_\lambda(\mu) \mu d\mu. \tag{3.8}$$

I then take the central intensity required to match the total synthetic flux with the synthetic flux calculated for the equivalent HiRes PHOENIX spectrum. This ensures the absolute flux of a star modelled with these limb darkening parameters remains consistent with the full disk PHOENIX model.

Temperatures of blackbody spectra that give monochromatic specific intensities equal to these specific intensities at $\mu = 1$ are then computed at the pivot wavelength of each filter in the SDSS and Super-SDSS systems using Planck's law,

$$T_{BB} = \frac{hc}{k_B \lambda \ln\left(\frac{2hc}{I_\lambda(1)\lambda^3} + 1\right)}. \tag{3.9}$$

This process is performed for each PHOENIX spectral model in each desired filter,

producing tables of equivalent blackbody temperatures that can then be interpolated during the fitting process.

3.2.5 Model summary

In summary, multi-band eclipse light curves clearly resolve the SEDs[3] of both components, with the depths of the eclipses defining the flux contributed by the WD and the in-eclipse flux defining the contribution from the secondary. These SEDs from the eclipse light curves constrain the effective temperatures of both stars which, together with parallax information, places constraints on their radii. The shape of the eclipse strengthens this constraint whilst also restricting the orbital separation and therefore the masses of the stars when combined with mass-radius relations.

I use LROCHE to model these light curves, defining the LROCHE model from the parameters of interest – T_1, T_2, M_1, M_2, i, T_0, ϖ, and $E(B-V)$, where ϖ is the parallax – together with the orbital period, P, and the bandpass of the observation, via various theoretical models and relations. The LROCHE parameters, described previously, are defined as follows:

1. The mass ratio, q, is set from the masses as $\frac{M_2}{M_1}$.

2. The scaled radius of the primary, $\frac{R_1}{a}$, is defined by the WD mass-radius relation for the chosen core composition together with Kepler's third law and therefore depends on the WD mass and temperature, the secondary mass, and the orbital period. The scaled radius of the secondary, $\frac{R_2}{a}$, is primarily defined by the semi-empirical mass-radius relation together with Kepler's third law, with corrections for irradiation and Roche distortion. It is therefore dependent on the WD mass and temperature, the secondary mass and temperature, and the orbital period of the system.

[3]When referring to the method described in this work, the SED is from the eclipse light curves alone and does not include any additional photometric data

3. The orbital inclination, i, is a free parameter.

4. The equivalent blackbody temperature of each star, $T_{\rm BB}$, is defined by the mass and temperature of the respective star together with the chosen bandpass.

5. For the orbital ephemeris, T_0 is a free parameter in the fit while P is fixed at a previously determined value.

6. The limb-darkening coefficients of both stars, a_1, a_2, a_3, and a_4, like the blackbody temperatures, are defined by the mass and temperature of the respective star together with the chosen bandpass.

7. The gravity-darkening coefficient for the secondary star, y, is also defined by its mass and temperature along with the chosen bandpass.

8. The fraction of incident flux from the primary that is absorbed by the secondary can generally be ignored due to only fitting a small region surrounding the primary eclipse. I therefore leave it fixed at $F_{\rm abs} = 0.5$.

3.2.6 χ^2 calculation for flux calibrated light curves

When generating a model light curve, LROCHE can be supplied with a scale factor which sets the absolute flux level of the light curve. It is therefore possible to calculate the scale factor required to produce a true flux light curve model from the parallax, interstellar extinction, and orbital separation. The issue with this approach is that any small error in the flux calibration of the data will cause issues with the fitting. This is because the flux calibrated eclipse light curves are unlikely to correspond exactly to the model SEDs of both stars, resulting in the fit being unable to correctly model both the in-eclipse and out-of-eclipse flux simultaneously, preventing an accurate fit to the eclipse shape and therefore reliable parameter estimation. I instead allow LROCHE to automatically scale the model to the data, calculating the $\chi^2_{\rm eclipse}$ using this scaled model.

To include the absolute flux information, I take the WD flux contribution output by LROCHE for the scaled model (a reliable measure of the depth of the primary eclipse) and compare this to the theoretical WD flux for the given temperature, mass, parallax, and extinction, calculating the $\chi^2_{\rm SED}$ for this using the flux calibration uncertainty. This method allows for a more proper handling of the uncertainties, treating those from the flux calibration and from the differential photometry independently. I combine these two values of χ^2, repeating this over all observed bands, x, to obtain an overall value for the full flux calibrated model,

$$\chi^2_{\rm model} = \sum_x \left(\chi^2_{{\rm SED},x} + \chi^2_{{\rm eclipse},x}\right) . \tag{3.10}$$

The log likelihood used in the MCMC is then simply proportional to $-\frac{1}{2}\chi^2_{\rm model}$.

3.2.7 Fitting procedure

To fit the light curves of a system I use MCMC, implemented through the EMCEE Python package (Foreman-Mackey et al., 2013). For each walker position an LROCHE light curve model is generated for each observed bandpass, with the log likelihood calculated as described in Section 3.2.6. The log likelihood from this model is combined with a parallax prior. Here, I use a bounded Gaussian prior with a mean and standard deviation corresponding to the *Gaia* `parallax` and `parallax_error` (Gaia Collaboration et al., 2021) respectively with a limit at two standard deviations above and below.

The use of the *Gaia* parallax distribution as a prior in the MCMC prevents systematic issues that would be introduced by using the distance. All systems considered in this work have a `parallax_over_error` > 10 and, when combined with the photometric information, will be constrained sufficiently that the inclusion of the Galactic stellar density distribution is unnecessary.

For systems with larger uncertainty in their parallax measurement, or maybe even

no parallax information at all – as will be the case for many of the systems discovered by LSST – this method can still be successful in fitting the parameters. It is possible, however, that in the case of systems where the SED of either star does not match the models well (most likely due to irradiation effects or star spots on the secondary), then the fit will compensate using the parallax. For systems with good parallax measurements, this effect is relatively obvious and can be flagged. Assuming it is the secondary SED that is causing the problem (which is most likely) then allowing its temperature to be independent in each band can allow the fit to converge to values consistent with the measured parallax[4]. For those with high parallax uncertainties, though, it may go unnoticed, leading to incorrect parameters.

Priors on all other parameters are uniform with upper and lower limits defined by the range of the model grids. Each MCMC chain is run with 100 walkers for 20 000 steps. The chains are then inspected and a number of steps are discarded from the beginning as burn-in. The number of steps discarded is chosen to remove any steps before the final equilibrium position.

3.3 Observations

In order to validate this modelling, I test this fitting code on three previously well characterised WD + M PCEBs. I use archival photometry of these systems from ULTRACAM (Dhillon et al., 2007) and HiPERCAM (Dhillon et al., 2021) to test that the method is successful when using either three-band or five-band data. I then fit three, previously unpublished, systems, all observed with ULTRACAM (a journal of observations is provided in Table 3.2)

[4]Note that this increases the number of dimensions for the MCMC and so increases the convergence time as well as removing the effective temperature information for the secondary. It is therefore best left as a backup method in the case where the original fit is struggling.

Table 3.2: Journal of observations. Note that the exposure times listed here are the shortest exposure time used in an observation. The u-band exposure will typically be a few multiples of these.

Target	Date	Filters	Telescope-Instrument	Exp time (s)	Sky transparency	FWHM (")
NN Ser	2019-07-09	$u_s g_s r_s i_s z_s$	GTC-HiPERCAM	1.0	Some dust but stable	1.5
SDSS J0838+1914	2010-12-13	$u' g' r'$	NTT-ULTRACAM	4.8	Photometric	2
SDSS J1028+0931	2018-01-30	$u_s g_s r_s$	NTT-ULTRACAM	2.5	Photometric	1
2MASS J1358−3556	2018-06-01	$u_s g_s i_s$	NTT-ULTRACAM	4.0	Photometric	1.5
EC 12250−3026	2018-05-31	$u_s g_s i_s$	NTT-ULTRACAM	3.0	Photometric	1.5
SDSS J1642+0135	2019-03-04	$u_s g_s i_s$	NTT-ULTRACAM	5.0	Photometric	1.5

3.3.1 Reduction

The HiPERCAM pipeline (Dhillon et al., 2021) was used to debias, flat-field correct – and defringe in the case of HiPERCAM z_s data – and then extract aperture photometry. The radius of the target aperture was allowed to vary in line with the measured full-width at half-maximum of a reference star in each frame to minimise the effects of seeing variations. Here, a target aperture radius of $1.8\times FWHM$ was used. The counts from the target were measured relative to a brighter comparison star to remove any transparency variations and atmospheric extinction effects.

3.3.2 Flux calibration

Using the SED of the WD to determine its temperature requires precise flux calibration. This requirement is emphasised by the significant temperature dependence of the WD mass-radius relations on which this method relies heavily (Figure 3.1). Flux calibration is complicated by the significant departure of the HiPERCAM and ULTRACAM Super-SDSS filter sets from the standard SDSS system. This departure is most notable in the u_s band where it can be tenths of magnitudes. The typical solution would be to observe a selection of spectro-photometric standard stars with well known spectra – observed or theoretical – spanning the full wavelength range of the filter set in order to calibrate the observed light curves from synthetic photometry. This is the method that will be employed in Chapter 4 using the *Gaia* spectro-photometric standard stars (Pancino et al., 2012; Altavilla et al., 2015; Marinoni et al., 2016; Altavilla et al., 2021; Pancino et al., 2021). Synthetic AB magnitudes of these *Gaia* spectro-photometric standard stars computed for the ULTRACAM and HiPERCAM Super-SDSS systems are included in Appendix A along with a more thorough analysis of the differences between the Super-SDSS and SDSS photometric systems.

However, as this work is based on archival photometry, only flux standards with

USNO–40 photometry (Smith et al., 2002) are available for flux calibration. I therefore fit PHOENIX model spectra to the USNO–40 and *Gaia* photometry of these standards to calculate the magnitude offsets required to transform the USNO–40 photometry into the Super-SDSS system.

Using an MCMC method, implemented through the EMCEE Python package, I fit the effective temperature, surface gravity, radius, interstellar reddening, and parallax of the standard stars. At each walker step the corrections between the two filter systems are saved. This gives us the ability to propagate the effects of any non-trivial correlations through to the uncertainties in the desired offsets. 100 walkers are run for 10,000 steps and the first 2000 of these are discard as burn-in.

I then measure the atmospheric extinction in each filter by fitting a first order polynomial to the instrumental magnitudes as a function of airmass of any bright stars included in an observing run that covers a good airmass range on the same night as our target. I then use this, along with the transformed Super-SDSS standard star magnitudes, to calibrate the comparison star of the target. The target is therefore calibrated by performing differential photometry against this comparison star.

3.3.3 Comparison with previously published systems

NN Ser

NN Ser is a well-characterised eclipsing binary, first discovered by Haefner (1989), containing a hot WD and an M dwarf companion with more than a decade of archival high-speed multi-colour photometry. Parsons et al. (2010) combined ULTRACAM photometry with phase-resolved UVES spectroscopy to obtain precise parameters for the system, independent of any mass-radius relations. This makes it an ideal system with which to test the purely photometric approach presented in this paper. Additionally NN Ser has been observed with HiPERCAM allowing us to test the method with simul-

taneous $u_s g_s r_s i_s z_s$ data.

Initial fits to the HiPERCAM photometry struggled, attempting to change the parallax to values inconsistent with the *Gaia* measurement. This was determined to be due to the M dwarf SED (the in-eclipse HiPERCAM photometry) not matching the PHOENIX model for the published parameters, instead preferring a higher temperature model. The fit was compensating for this by altering the radius, and therefore mass of the secondary, having an effect on the rest of the system parameters. This likely reflects the fact that the high irradiation experienced by the M dwarf prevents it from being well represented by the PHOENIX models, even on the unirradiated face seen during the primary eclipse. Allowing the temperature of the secondary to be independent in each band solved this issue, allowing the MCMC to converge to a fit consistent with the *Gaia* parallax.

Our best fit to the HiPERCAM photometry (see Table 3.3 and Figure 3.3) achieves uncertainties (and therefore precisions) in the WD mass, effective temperature, and secondary mass of 1.8, 4.6, and 5.7 per cent respectively. Although our secondary mass is around 11 per cent more massive than the published value it is still consistent to better than 2σ. The other parameters all lie within 1σ of the high precision, published values (Parsons et al., 2017a, 2018). This demonstrates that, even for systems with highly irradiated companions, this method can still be successful.

SDSS J0838+1914

SDSS J0838+1914 (SDSS J083845.86+191416.5 in SIMBAD, also known as CSS 40190) is another well characterised PCEB and was discovered by Drake et al. (2010) in the Catalina Sky Survey. It was later characterised by Parsons et al. (2017a, 2018) using a combination of eclipse photometry and phase-resolved spectroscopy, finding it to contain a WD with a mass and temperature of 0.482 ± 0.008 $M_\odot$ and $14\,900 \pm 730$ K, respectively, in a 3.123 h orbit with a 3100 ± 100 K main sequence companion with a mass of 0.142 ± 0.013 $M_\odot$. It is therefore a fairly typical example of a PCEB. Good quality archival

ULTRACAM photometry is available for this system. Importantly, it has also been characterised using SDSS spectroscopy, making it a good system with which to compare our photometric fit against the spectroscopic one.

Fitting the ULTRACAM eclipse photometry of the system (see Table 3.3 and Figure 3.4) achieves a precision in the WD mass and effective temperature of 2.9 and 2.4 per cent respectively and the secondary mass and effective temperature of 6.1 and 0.3 per cent respectively. Comparing these to the SDSS spectroscopic values (Rebassa-Mansergas et al., 2012) – which determine the WD mass and effective temperature to be 0.39 ± 0.035 $M_\odot$ and $13\,904 \pm 424$ K with corresponding precisions of 9.0 and 3.0 per cent respectively, and the secondary mass to be 0.255 ± 0.124 $M_\odot$ (a precision of 49 per cent) – demonstrates that a higher precision can be reached using eclipse photometry. Additionally all my best fit values lie within 2σ of the high-precision, published values (Parsons et al., 2017a, 2018), showing that using the eclipse photometry can also yield accurate parameters.

SDSS J1028+0931

SDSS J1028+0931 (SDSS J102857.78+093129.8 in SIMBAD) was discovered to be an eclipsing WD + M system by Parsons et al. (2013b). Later and more precise characterisation was performed by Parsons et al. (2017a, 2018), who found the WD to have a mass and temperature of 0.4146 ± 0.0036 $M_\odot$ and $12\,221 \pm 765$ K respectively, with the fit to the secondary giving a mass and temperature of 0.403 ± 0.005 $M_\odot$ and $3\,500 \pm 100$ K. The higher mass secondary in this system means that the spectrum is dominated by the M dwarf redward of the g-band with measurable contribution in the u-band. This makes it a difficult system to characterise from low-resolution spectroscopy due to the dilution of the Balmer series of the WD by the brighter M dwarf. The significant M dwarf contribution therefore makes this system a good one with which to test the method.

The secondary of SDSS J1028+0931 appears to have at least one significant star

spot on its surface which adds an additional slope to the photometry. A linear term is therefore added to the g_s- and r_s-band models to account for this. The best fit including these two additional parameters (see Table 3.3 and Figure 3.5) achieves a precision of 1.4 per cent on the WD mass and 1.9 per cent on its temperature. For the M dwarf, the fit to the eclipse photometry manages precisions of 5.3 and 0.8 per cent respectively for the mass and temperature. Again, all of our best-fit parameters are within 2σ of the high-precision values of Parsons et al. (2017a, 2018). Comparing our fit with the spectroscopic values from SDSS of $T_1 = 15\,782 \pm 160$ K, $M_1 = 0.83 \pm 0.022$ $M_\odot$, and $M_2 = 0.406 \pm 0.086$ $M_\odot$ (Rebassa-Mansergas et al., 2012), demonstrates the benefit of eclipse photometry in systems where one component dominates – with the spectroscopic determination of the WD temperature being discrepant by over 4σ and the mass being discrepant by almost 10σ.

Table 3.3: Comparison of parameters determined using our purely photometric method against published, model-independent values (Parsons et al., 2017a, 2018). I show the deviation from the published values as a percentage and in units of standard deviation where the standard deviation is the uncertainty from our photometric fit summed in quadrature with the uncertainty of the published value. I also include a 1 per cent and 5 per cent systematic error contribution for the primary and secondary masses respectively.

	NN Ser		SDSS J0838+1914		SDSS J1028+0931	
	μ	Difference %(σ)	μ	Difference %(σ)	μ	Difference %(σ)
T_1(K)	60800^{+2200}_{-2800}	$-3.5(-0.6)$	14060^{+340}_{-340}	$-5.7(-1.0)$	13270^{+250}_{-140}	$+8.6(+1.3)$
T_2(K)	–	–	2910^{+10}_{-10}	$-6.2(-1.9)$	3550^{+30}_{-20}	$+1.3(+0.4)$
M_1($M_\odot$)	$0.548^{+0.010}_{-0.009}$	$+2.4(+0.9)$	$0.456^{+0.012}_{-0.013}$	$-5.2(-1.7)$	$0.428^{+0.006}_{-0.006}$	$+3.2(+1.9)$
M_2($M_\odot$)	$0.123^{+0.007}_{-0.007}$	$+10.8(+1.6)$	$0.148^{+0.009}_{-0.009}$	$+4.3(+0.4)$	$0.397^{+0.021}_{-0.021}$	$-1.4(-0.3)$

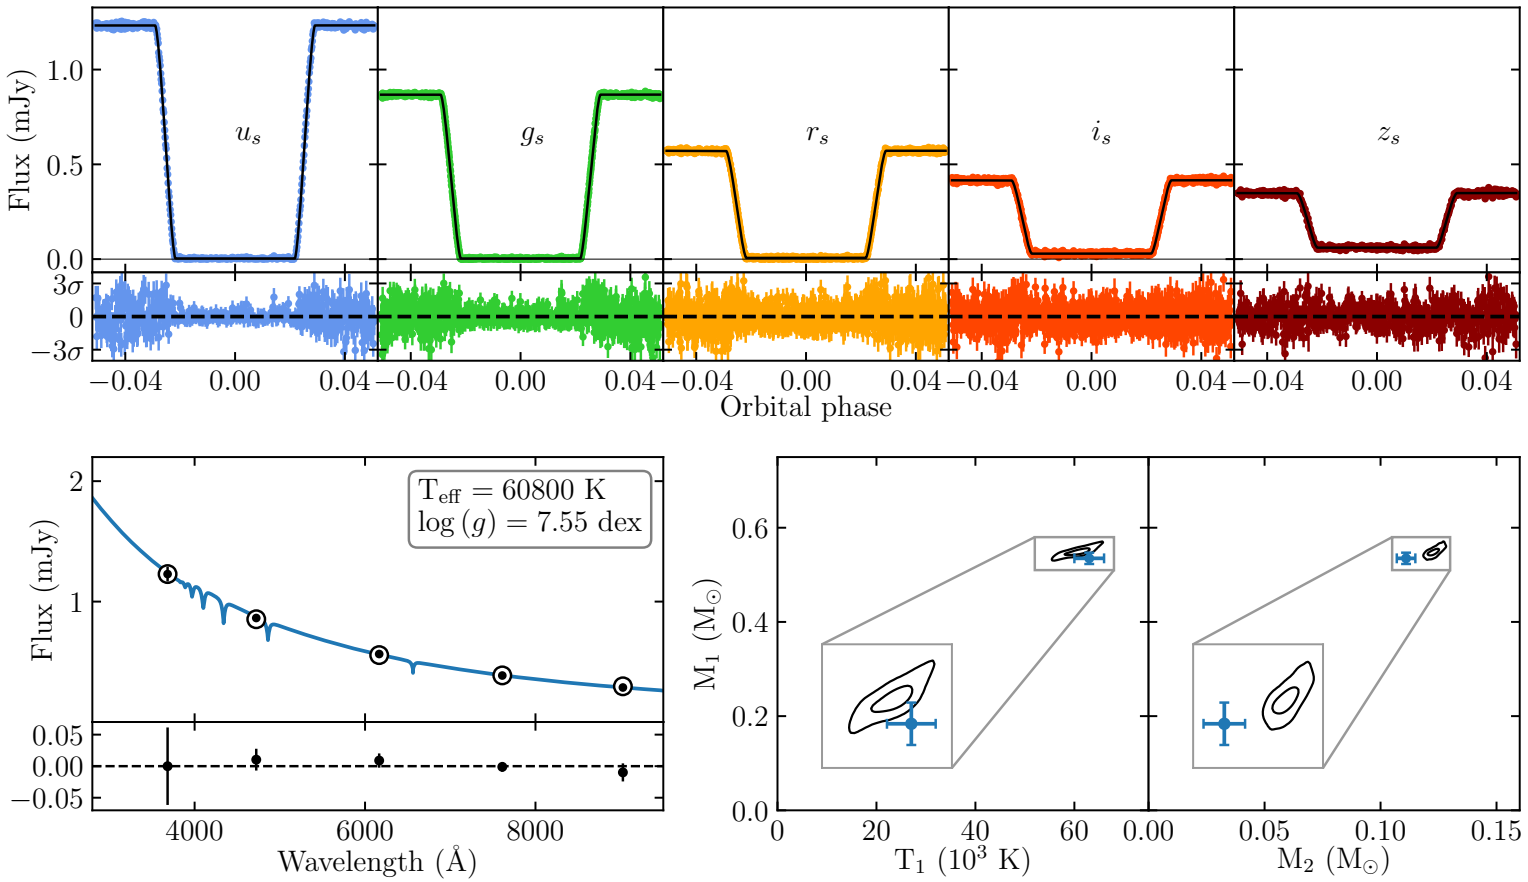

Figure 3.3: The best fit model for NN Ser. *Top*: The flux calibrated HiPERCAM $u_s g_s r_S i_s z_s$ eclipse light curves (coloured points) along with the best fit eclipse model (black line). A horizontal black line shows a flux of zero for reference with the residuals shown below. *Bottom left*: The SED of the WD i.e., the depths of the eclipses in each band (black points). These are shown against the Koester (2010) model spectrum for the best fit parameters (blue line) and the synthetic photometry from this model (open circles) with the residuals displayed below. *Bottom right*: The 1σ and 2σ contours from our MCMC fit alongside the model-independent parameters (blue points, Parsons et al. 2017a, 2018)

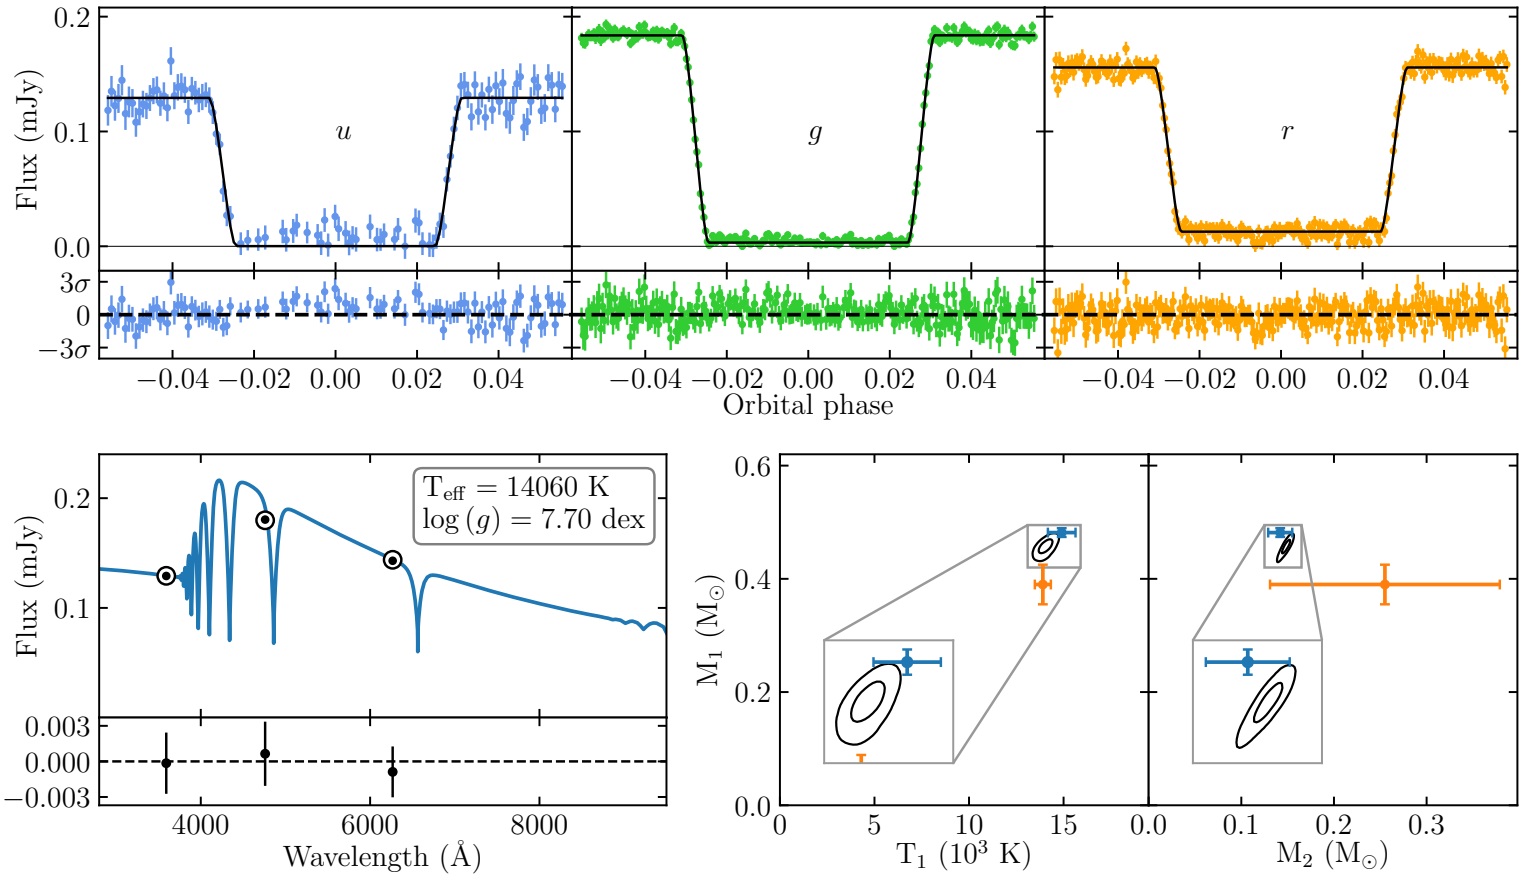

Figure 3.4: The best fit model for SDSS J0838+1914. *Top*: The flux calibrated ULTRACAM $u'g'r'$ eclipse light curves (coloured points) along with the best fit eclipse model (black line). A horizontal black line shows a flux of zero for reference with the residuals shown below. *Bottom left*: The SED of the WD i.e., the depths of the eclipses in each band (black points). These are shown against the Koester (2010) model spectrum for the best fit parameters (blue line) and the synthetic photometry from this model (open circles) with the residuals displayed below. *Bottom right*: The 1σ and 2σ contours from our MCMC fit alongside the model-independent parameters (blue points, Parsons et al. 2017a, 2018) and those derived from SDSS spectroscopy (orange points, Rebassa-Mansergas et al. 2012).

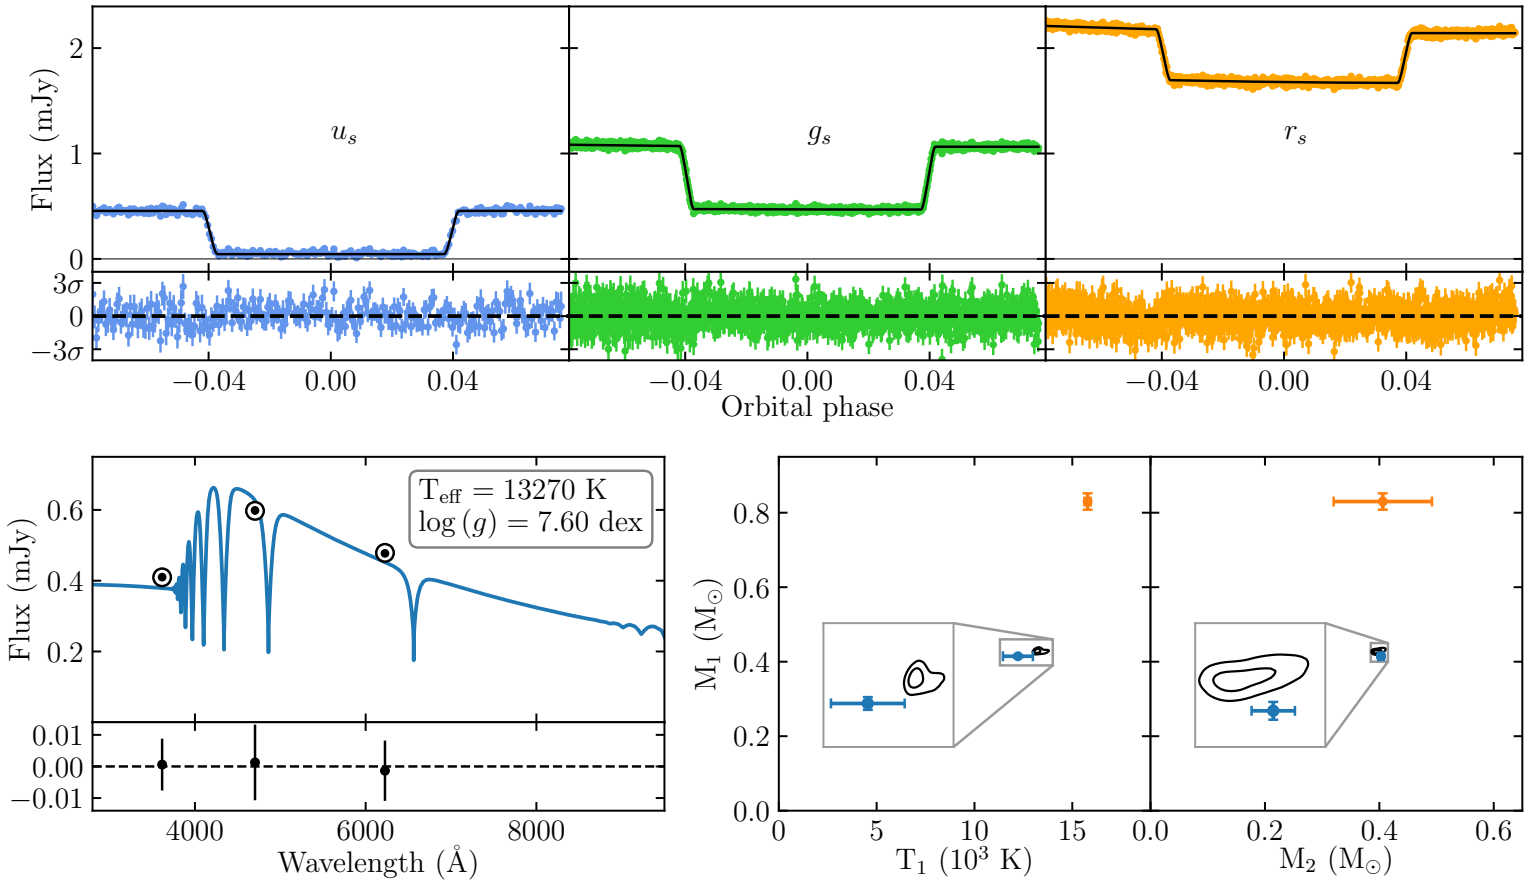

Figure 3.5: The best fit model for SDSS J1028+0931. *Top*: The flux calibrated ULTRACAM $u_s g_s r_s$ eclipse light curves (coloured points) along with the best fit eclipse model (black line). A horizontal black line shows a flux of zero for reference with the residuals shown below. *Bottom left*: The SED of the WD i.e., the depths of the eclipses in each band (black points). These are shown against the Koester (2010) model spectrum for the best fit parameters (blue line) and the synthetic photometry from this model (open circles) with the residuals displayed below. *Bottom right*: The 1σ and 2σ contours from our MCMC fit alongside the model-independent parameters (blue points, Parsons et al. 2017a, 2018) and those derived from SDSS spectroscopy (orange points, Rebassa-Mansergas et al. 2012).

3.3.4 New systems

With the code proving successful for these three well-characterised systems, I will now apply it to three previously uncharacterised PCEBs observed with ULTRACAM. Best-fit parameters for these new systems are listed in Table 3.4 with the light curves shown in Figure 3.6.

2MASS J1358−3556

2MASS J1358−3556 (2MASS J13581075−3556194) was found to be an eclipsing PCEB by combining *Gaia* measurements and data from the Catalina Real-Time Transient Survey (CRTS; Drake et al. 2009). Fitting the ULTRACAM eclipse photometry in u_s, g_s, and i_s gives a WD mass of 0.438±0.007 $M_\odot$ with an effective temperature of $40\,600^{+1600}_{-2200}$ K assuming a He-core composition. Rerunning the fit with a CO-core mass-radius relation favours a lower WD mass of 0.38 $M_\odot$. This mass is below what is expected for a WD with a CO-core composition and so the helium fit appears the most consistent with evolutionary models. The M dwarf in this system fits best with a mass of 0.118 ± 0.006 $M_\odot$ and an effective temperature of 2980^{+30}_{-40} K. The best fit parameters are shown in Table 3.4 with the light curve model shown in Figure 3.6. This system appears to be a fairly typical, if quite young, PCEB due to the high WD temperature.

EC 12250−3026

EC 12250−3026 was initially found in the Edinburgh-Cape blue object survey (Stobie et al., 1997) where it was thought to be a single hot subdwarf. It was later found to be an eclipsing WD + M system in CRTS. The fit to the ULTRACAM photometry (Table 3.4, Figure 3.6) gives a WD temperature of $33\,900^{+1000}_{-1300}$ K with a mass of $0.420^{+0.010}_{-0.009}$ $M_\odot$, again favouring a He-core composition. The secondary in this system – which is only detected in the i_s band – is quite low mass at 0.089 ± 0.005 $M_\odot$ although still stellar.

The low-mass of the secondary star makes this an interesting system for investigating the brown dwarf transition regime. For the secondary temperature, I find 2840^{+60}_{-30} K, although it is worth mentioning that, like NN Ser, the high temperature of this WD may mean that the best fit temperature of the M dwarf is influenced by the high irradiation and is not necessarily representative of the true unirradiated temperature.

SDSS J1642+0135

SDSS J1642+0135 (SDSS J164251.54+013554.9 in SIMBAD) was discovered by Denisenko & Larin (2018) who determined it to be a pre-CV with a cool ($T_{\rm eff} \lesssim 8000$ K) WD primary and strong ellipsoidal modulation. This system is of particular interest as its period of 2.31 h places it in the CV period gap (the range of orbital periods between 2 and 3 hours where there is a dramatically drop in the numbers of observed CVs).

Initial fits of this system were unable to match the extreme ellipsoidal modulation present in the g_s- and i_s-band photometry, even with a Roche lobe filling factor equal to one. This required adding the gravity-darkening exponent for the secondary star, β_1 (Claret & Bloemen, 2011), to the model as a free parameter, thereby allowing the fit to increase the effect of gravity darkening, preferring a value of $\beta_1 = 0.41 \pm 0.02$. This is roughly twice what would be expected for a star of the best fit temperature (Claret, 2003), however, observational validation of how β_1 varies with effective temperature seems to indicate a fairly large scatter around the theoretical values (Claret, 2003, see figure 3). With the addition of β_1 as a free parameter I confirm the predictions of Denisenko & Larin (2018), finding a 7650 ± 60 K WD with a mass of $0.69^{+0.010}_{-0.011}$ $M_\odot$ and a secondary that is on the verge of – if not already – filling its Roche lobe with a linear filling factor of 0.97 ± 0.02. I find a secondary mass of 0.198 ± 0.010 $M_\odot$ and effective temperature of 2897^{+5}_{-6} K (Figure 3.6 and Table 3.4).

This WD mass is higher than – although still very much consistent with – the average WD mass for PCEBs (0.58 ± 0.20 $M_\odot$, Zorotovic et al. 2011), but is also consistent with

that of the volume-limited CV population (0.83 ± 0.17 $M_\odot$, Pala et al. 2020). Additionally, the secondary mass matches the expected donor mass of $M_{\rm donor} = 0.2\pm0.02$ $M_\odot$ at which CV mass transfer ceases at the start of the period gap (Knigge, 2006). It therefore seems possible that SDSS J1642+0135 is a temporarily detached CV in the final stages of crossing the period gap with an orbital period ≈ 10 mins greater than the predicted period at the lower edge of the gap of $P = 2.15$ h (Knigge, 2006). A population of apparently gap-crossing CVs has been identified statistically using the SDSS sample of PCEBs (Zorotovic et al., 2016), however, this may be one of the first specific examples of an eclipsing, gap-crossing CV and is worth more detailed study.

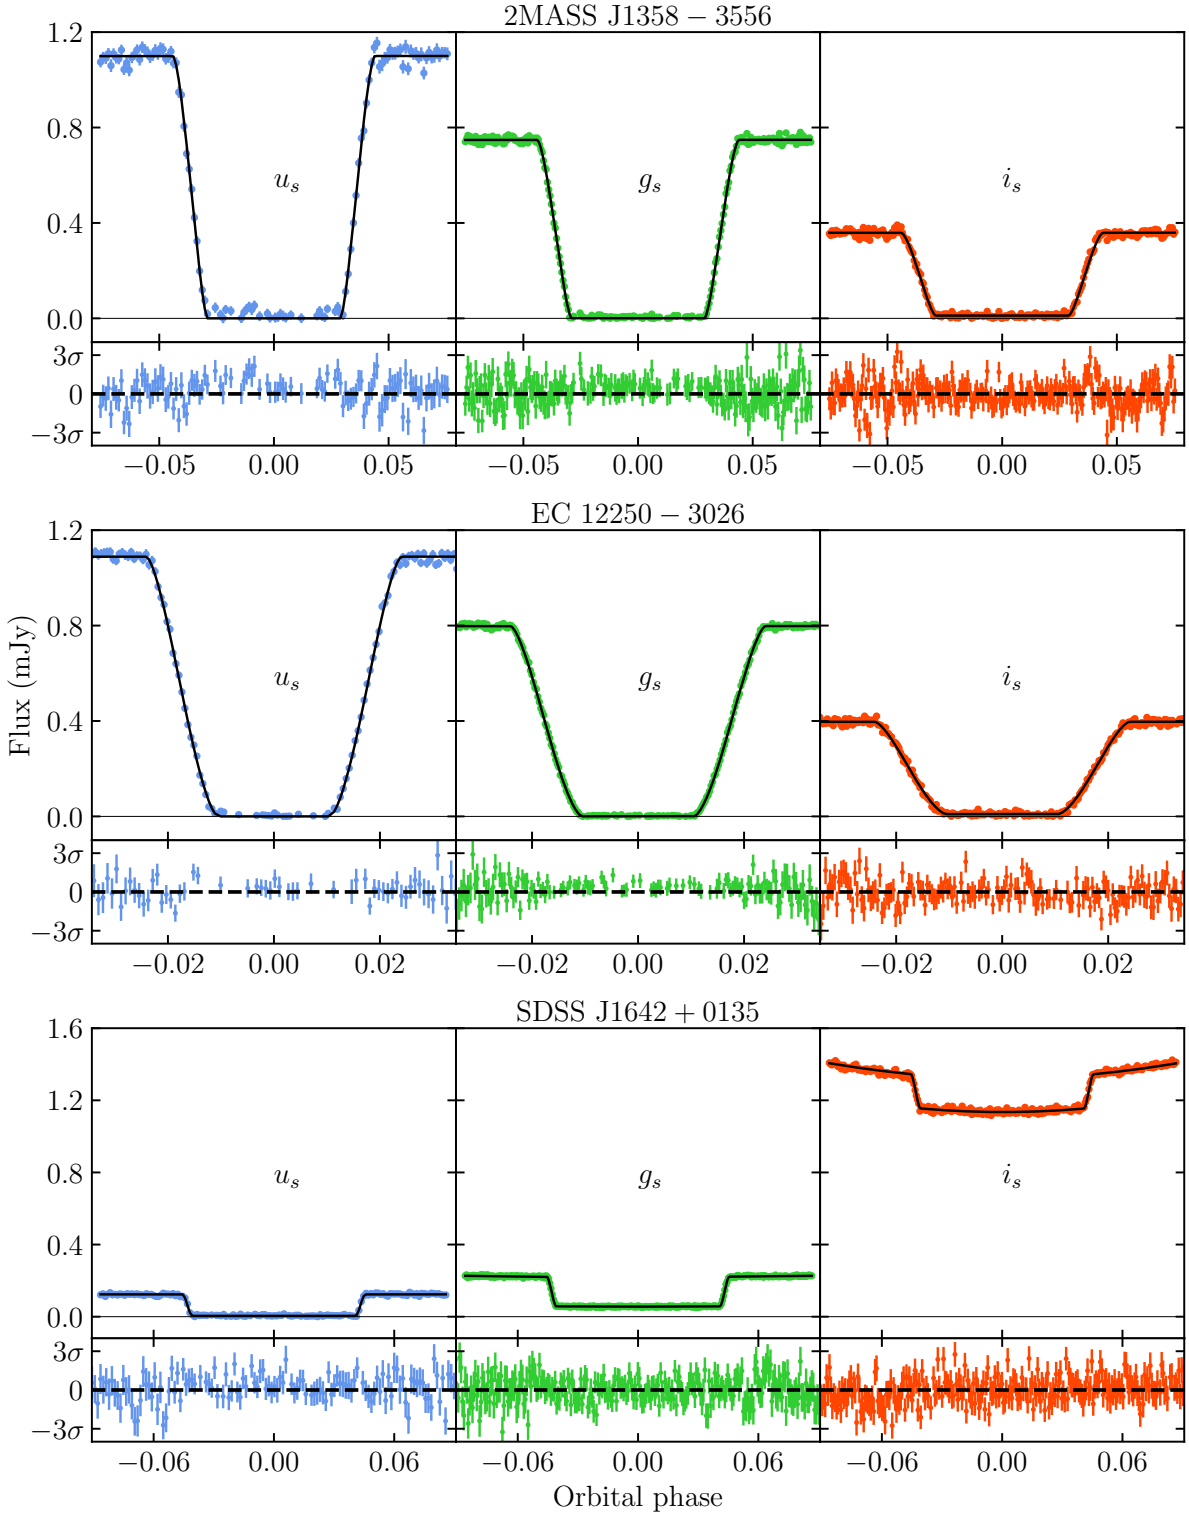

Figure 3.6: Best fits to ULTRACAM eclipse photometry of the three new systems. The best fit model is shown by the solid black line while a thin black line shows the zero level. Residuals are shown below with a dashed black line showing zero.

Table 3.4: Results of the MCMC fits to the three new WD + M systems.

Target	T_1 (K)	T_2 (K)	M_1 ($M_\odot$)	M_2 ($M_\odot$)	i (°)	$E(B-V)$	T_0 (BMJD)
2MASS J1358−3556	40600^{+1600}_{-2200}	2980^{+30}_{-40}	$0.438^{+0.007}_{-0.007}$	$0.118^{+0.006}_{-0.006}$	$88.7^{+0.9}_{-1.3}$	$0.04^{+0.01}_{-0.01}$	58270.1646820(12)
EC 12250−3026	33900^{+1000}_{-1300}	2840^{+60}_{-30}	$0.420^{+0.010}_{-0.009}$	$0.089^{+0.005}_{-0.005}$	$85.1^{+0.1}_{-0.2}$	$0.04^{+0.01}_{-0.01}$	58269.97928329(67)
SDSS J1642+0135	7650^{+60}_{-60}	2897^{+5}_{-6}	$0.693^{+0.010}_{-0.011}$	$0.198^{+0.010}_{-0.010}$	$89.1^{+0.6}_{-0.7}$	$0.01^{+0.01}_{-0.01}$	58546.3505016(17)

3.4 Discussion

Assessing the precision of our fits to the six systems, the median percentage uncertainty on the WD mass determination is 1.7 per cent with a maximum uncertainty of 2.9 per cent (SDSS J0838+1914). For the secondary mass these median and maximum uncertainties are 5.5 per cent and 6.1 per cent respectively which are dominated by the 5 per cent estimated contribution from the intrinsic scatter of the M dwarf mass-radius relationship. This precision is, therefore, at or below the aim of 5 per cent precision which is necessary to discern systems with interesting subtypes of either component. Comparing these percentage uncertainties with the two other methods of characterisation mentioned – spectral decomposition (Rebassa-Mansergas et al., 2012) and VOSA SED fitting (Rebassa-Mansergas et al., 2021) – demonstrates that using the eclipse gives similar or better precision. For the WD mass, the median percentage uncertainty of the SDSS sample (Rebassa-Mansergas et al., 2012) is 18 percent with a 16th percentile (comparable to one standard deviation below the mean for non-normal distributions) of 7 per cent. For the VOSA sample (Rebassa-Mansergas et al., 2021) the median is 11 per cent with a 16th percentile of 6 per cent. Our mass determinations are therefore significantly more precise than either of these methods. To illustrate this point further, our least precise measurement of WD mass is more precise than 98.5 per cent of WD mass measurements in the SDSS sample and better than all WD mass measurements in the VOSA sample. For the WD temperature our uncertainties are much more comparable, with a median uncertainty of 3 per cent and a maximum of 5 per cent. This is compared with a median value of 4 per cent for the SDSS sample.

Comparing our best fit parameters to the published, model-independent values (Table 3.3) additionally demonstrates that a purely photometric approach relying only on eclipse photometry can yield parameters with greater reliability and accuracy than from low-resolution spectroscopy. This improvement on the spectroscopic method is particu-

larly obvious in SDSS J1028+0931 where the contribution from the M dwarf companion is significant, causing the SDSS spectroscopic fit to significantly overestimate the mass and temperature of the WD. For all three systems, our parameters are consistent to within 2σ of the published parameters, with most values accurate to better than around 5 per cent. This level of accuracy is better than the original goal of 3σ.

Possible sources of systematic error that I have not accounted for include the assumption of thick, DA WD models. Although this is consistent with the findings of Parsons et al. (2017a), it will lead to systematic errors if used for a system containing a WD with a thin hydrogen atmosphere or a helium atmosphere. Additionally, many of the systems considered here contain WDs with masses in the range where theorised hybrid WDs lie (Zenati et al., 2019). Although these have not been confirmed observationally, with a tentative suggestion from Parsons et al. (2020) that has been supported by (Romero et al., 2022a), a hybrid core would introduce a similar error into the WD parameters due to an incorrect mass-radius relation. For the secondary, the best-fit effective temperature can be affected by the presence of star spots on the surface or from irradiation effects due to a hot WD, as seems to be the case in NN Ser. It is also worth mentioning that the statistical uncertainties on the secondary temperature resulting from the MCMC fit are very likely underestimated. The PHOENIX (Husser et al., 2013) model grid that is used here has a 100 K resolution in effective temperature and so any uncertainties much below this level are unlikely to represent the true error.

Given the success of eclipse photometry for initial characterisation of WD + M systems, this method will be applicable to LSST data. How well this works will depend on the quality of the absolute flux calibration as well as the final survey strategy, particularly whether the individual 15 second images or photometry are available. This is due to the need to resolve the sharp eclipse features in order to constrain the radii (and hence, masses) of the components. As previously mentioned, many systems discovered by LSST will have little to no parallax information initially (LSST will provide parallax

measurements to many of these as the survey progresses). Although fitting the eclipse photometry is still possible in this case, it is more difficult to flag when a fit is converging to erroneous values as a result of systematic errors.

3.5 Summary

In this chapter I have demonstrated that – when combined with *Gaia* parallax measurements – high-cadence, multi-colour eclipse photometry can be used to determine masses and temperatures of WD + M binaries more reliably than low-resolution spectroscopy, achieving a precision of better than 5 per cent on the WD parameters and better than 6 per cent for the M dwarf, making future follow up of these systems easier and more robust. The use of the primary eclipse also guarantees that the photometric SEDs are analysed at the same orbital phase, preventing any possible issues that may arise from using the Virtual Observatory SED Analyser to fit the system.

Additionally, as well as being able to be used for fainter systems than spectroscopic methods, the photometric nature of this method is better equipped to find systems of particular interest such as those displaying variability due to magnetic or pulsating WDs; high-mass WDs, from their sharp eclipse features; or systems with brown dwarf companions that would otherwise be washed out in the optical if not for the clear eclipses. The lack of need for low-resolution spectroscopy also makes this method more time efficient, with the high-cadence photometry being reusable for any high precision follow up work (unlike low-resolution spectroscopy which is often not useful beyond the initial identification).

I have also used this method to determine parameters for three new PCEBs, finding two to contain hot, He-core WDs with low mass companions (one of which is near the brown dwarf transition regime), and one to be a possibly detached CV close to coming back into contact.

Chapter 4

Photometric follow-up of eclipsing WD + M binaries from ZTF

4.1 Introduction

Having confirmed that high-cadence eclipse photometry can be used to accurately and precisely parametrise WD + M systems, I will now use this method to characterise a large number of eclipsing WD + M binaries discovered in the ZTF survey. I do this with the aim of significantly increasing the sample of well-characterised WD + M systems known, as well as discovering any particularly interesting systems within the sample, such as those containing pulsating, magnetic, or high-mass WDs and any with brown dwarf companions.

4.2 Observations

4.2.1 Target selection

The systems targeted for follow-up were selected from the ∼900 detached eclipsing WD + M systems discovered by van Roestel et al. (in preparation) using data from ZTF. This parent sample was created by searching for periodic outliers in the ZTF photometry, indicative of eclipses. The primary biases of the parent sample are therefore related to the probability of a given system eclipsing as viewed from Earth and the ability to detect an eclipse within the ZTF data. The former is dominated by the orbital period (with a very weak dependence on the secondary radius), while the latter is dominated by the signal-to-noise ratio of the eclipse, with a heavy dependence on the depth of the eclipse (and a much weaker dependence on the duration of the eclipse). A more detailed description of the full ZTF eclipsing WD + M sample identification method and the biases within it will be presented in van Roestel et al. (in preparation).

The target list was restricted to systems visible from the La Silla Observatory (Dec $<$ +25 deg) and brighter than $g = 19.5$ mag, leaving an observable sample of ∼300 systems, all possessing *Gaia* parallax measurements. We typically observed systems with eclipse

timings that made for the most efficient use of telescope time on a particular night, however we also tried to prioritise systems with longer periods where possible since the eclipses of these systems are more difficult to observe. Systems with ZTF light curves that indicated they may be of particular interest were also prioritised. This included systems with in-eclipse flux measurements at or below the detection threshold of ZTF (possibly indicating a brown dwarf companion) and systems with unusual ZTF light curves, showing variability inconsistent with typical binary variability mechanisms and indicating the presence of a magnetic WD. A journal of observations is included in Table 4.1.

Table 4.1: Journal of observations.

Target	Date at start of run	Telescope-Instrument	Filters	g_s Exposure time (s)	Number of exposures	Conditions (Transparency, seeing)
ZTF J0126+1210	2021-07-08	NTT-ULTRACAM	$u_s g_s i_s$	7.0	804	clear, ∼1 arcsec
ZTF J0220+6303	2021-08-07	GTC-HiPERCAM	$u_s g_s r_s i_s z_s$	0.7	3568	clear, ∼0.6 arcsec
ZTF J0406+0958	2021-02-08	NTT-ULTRACAM	$u_s g_s i_s$	7.5	269	clear, <1.5 arcsec
	2021-11-08	NTT-ULTRACAM	$u_s g_s i_s$	8.0	1426	clear, ∼1 arcsec
	2021-11-10	NTT-ULTRACAM	$u_s g_s r_s$	8.0	712	clear, <1.5 arcsec
	2021-09-13	GTC-HiPERCAM	$u_s g_s r_s i_s z_s$	1.2	3183	clear, <2.0 arcsec
ZTF J0410−0834	2022-03-08	NTT-ULTRACAM	$u_s g_s i_s$	3.8	528	clear, <1.8 arcsec
ZTF J0519+0925	2021-02-06	NTT-ULTRACAM	$u_s g_s i_s$	8.0	156	clear, <1.5 arcsec
ZTF J0528+2156	2021-02-08	NTT-ULTRACAM	$u_s g_s i_s$	4.0	632	clear, <1.4 arcsec
	2022-12-19	NTT-ULTRACAM	$u_s g_s i_s$	3.7	1806	clear, <2 arcsec
ZTF J0537−2450	2021-02-06	NTT-ULTRACAM	$u_s g_s i_s$	3.0	1386	clear, <1.3 arcsec
ZTF J0615+0510	2021-03-07	NTT-ULTRACAM	$u_s g_s i_s$	6.8	774	clear, <1.4 arcsec
ZTF J0618−0919	2022-12-19	NTT-ULTRACAM	$u_s g_s i_s$	4.5	690	clear, ∼1.3 arcsec
	2022-12-19	NTT-ULTRACAM	$u_s g_s i_s$	3.4	474	clear, ∼1.3 arcsec
ZTF J0638+0910	2021-02-07	NTT-ULTRACAM	$u_s g_s i_s$	7.0	656	clear, <1.2 arcsec
ZTF J0639+1919	2021-02-06	NTT-ULTRACAM	$u_s g_s i_s$	3.0	1300	clear, ∼1.3 arcsec
ZTF J0642+1314	2021-02-07	NTT-ULTRACAM	$u_s g_s i_s$	4.5	591	clear, ∼1 arcsec
ZTF J0651+1452	2021-02-07	NTT-ULTRACAM	$u_s g_s i_s$	7.0	359	clear, ∼1.5 arcsec
ZTF J0704−0201	2021-02-08	NTT-ULTRACAM	$u_s g_s i_s$	6.0	411	clear, <1.5 arcsec
ZTF J0717+1136	2021-02-06	NTT-ULTRACAM	$u_s g_s i_s$	4.0	666	clear, <1.3 arcsec
ZTF J0718−0852	2021-03-10	NTT-ULTRACAM	$u_s g_s i_s$	7.5	379	clear, <1.2 arcsec
ZTF J0804−0215	2022-03-08	NTT-ULTRACAM	$u_s g_s i_s$	3.1	736	clear, ∼1.7 arcsec
ZTF J0805−1430	2022-03-08	NTT-ULTRACAM	$u_s g_s i_s$	6.0	757	clear, ∼1 arcsec
ZTF J0948+2538	2021-01-25	NTT-ULTRACAM	$u_s g_s i_s$	6.5	602	clear, ∼1 arcsec
ZTF J1022−0803	2021-03-07	NTT-ULTRACAM	$u_s g_s i_s$	4.0	1224	clear, <1.4 arcsec
ZTF J1026−1013	2021-01-23	NTT-ULTRACAM	$u_s g_s i_s$	5.8	316	clear, ∼2 arcsec

Target	Date at start of run	Telescope-Instrument	Filters	g_s Exposure time (s)	Number of exposures	Conditions (Transparency, seeing)
ZTF J1049−1755	2021-01-24	NTT-ULTRACAM	$u_s g_s i_s$	4.5	866	clear, ∼1 arcsec
ZTF J1206+5100	2021-05-08	GTC-HiPERCAM	$u_s g_s r_s i_s z_s$	5.4	9509	thin cloud, ∼2 arcsec
ZTF J1220+0821	2021-02-07	NTT-ULTRACAM	$u_s g_s i_s$	7.0	536	clear, <1.6 arcsec
ZTF J1256+2117	2022-03-03	NTT-ULTRACAM	$u_s g_s i_s$	2.5	1440	clear, ∼1 arcsec
ZTF J1302−0032	2021-02-07	NTT-ULTRACAM	$u_s g_s i_s$	4.0	692	clear, ∼1 arcsec
ZTF J1341−0626	2021-01-23	NTT-ULTRACAM	$u_s g_s i_s$	3.5	1216	clear, ∼1.8 arcsec
ZTF J1400+0814	2021-02-08	NTT-ULTRACAM	$u_s g_s i_s$	10.0	434	clear, ∼1 arcsec
ZTF J1404+0655	2021-01-25	NTT-ULTRACAM	$u_s g_s i_s$	6.0	705	clear, ∼1 arcsec
ZTF J1405+1039	2021-02-06	NTT-ULTRACAM	$u_s g_s i_s$	10.0	276	clear, ∼1.2 arcsec
ZTF J1407+2115	2021-02-07	NTT-ULTRACAM	$u_s g_s i_s$	4.5	468	clear, ∼1.2 arcsec
	2022-03-05	NTT-ULTRACAM	$u_s g_s i_s$	6.0	1304	clear, ∼1 arcsec
	2022-03-27	NTT-ULTRACAM	$u_s g_s r_s$	5.0	3455	clear, ∼1.1 arcsec
ZTF J1458+1313	2021-02-07	NTT-ULTRACAM	$u_s g_s i_s$	6.5	284	clear, ∼1 arcsec
ZTF J1626−1018	2021-02-08	NTT-ULTRACAM	$u_s g_s i_s$	5.0	660	clear, ∼1.3 arcsec
ZTF J1634−27132	2021-02-08	NTT-ULTRACAM	$u_s g_s i_s$	5.4	420	clear, ∼1.2 arcsec
ZTF J1644+2434	2021-03-10	NTT-ULTRACAM	$u_s g_s i_s$	6.0	321	clear, ∼1.0 arcsec
ZTF J1802−0054	2022-03-08	NTT-ULTRACAM	$u_s g_s i_s$	4.0	728	clear, ∼1.2 arcsec
ZTF J1828+2308	2022-04-26	NTT-ULTRACAM	$u_s g_s i_s$	4.0	532	clear, ∼2 arcsec
ZTF J1922+1038	2022-06-07	NTT-ULTRACAM	$u_s g_s i_s$	3.9	803	thin cloud, ∼1 arcsec
ZTF J1954+1019	2022-04-28	NTT-ULTRACAM	$u_s g_s i_s$	1.4	3174	clear, <2.5 arcsec
ZTF J2142+4309	2021-09-06	GTC-HiPERCAM	$u_s g_s r_s i_s z_s$	2.0	1544	clear, ∼1 arcsec
ZTF J2220+0721	2021-09-11	GTC-HiPERCAM	$u_s g_s r_s i_s z_s$	0.7	4047	some clouds, ∼0.8 arcsec
ZTF J2353+4153	2021-09-09	GTC-HiPERCAM	$u_s g_s r_s i_s z_s$	1.0	2770	clear, ∼0.6 arcsec

4.2.2 High speed photometry

The photometric follow-up observations made use of the three-band frame-transfer camera, ULTRACAM (Dhillon et al., 2007), mounted on the 3.6 m New Technology Telescope (NTT) at the ESO La Silla Observatory in Chile, to obtain high-cadence multi-colour photometry of the primary eclipse of each system. For all targets observed with ULTRACAM, the higher throughput Super-SDSS $u_s\,g_s\,i_s$ filters were used (Dhillon et al., 2021), with the exception of one observation where the r_s filter was used in place of i_s. For a few of the systems thought to harbour magnetic WDs, I obtained high-speed photometry with the quintuple band frame-transfer camera, HiPERCAM (Dhillon et al., 2021), mounted on the 10.4 m Gran Telescopio Canarías (GTC) at the Roque de los Muchachos observatory in La Palma, again equipped with Super-SDSS $u_s\,g_s\,r_s\,i_s\,z_s$ filters.

All observations were bias-subtracted and flat-field corrected (and fringe corrected in the case of the HiPERCAM z_s band) using the HiPERCAM pipeline[1]. Differential aperture photometry was then extracted using a variable aperture radius set to scale with the measured full width at half-maximum (FWHM) in each frame in order to remove effects due to seeing and transparency variations. For this I use a target aperture radius of $1.8 \times$ FWHM. In observations with lower signal-to-noise ratios, optimal extraction (Naylor, 1998) was also performed, with the extraction method resulting in the highest signal-to-noise light curve being the one that was used.

Flux calibration was then performed by fitting the atmospheric extinction in each band using one or more observing runs taken on the same night as the target observations (each spanning a minimum of 0.2 airmasses). The atmospheric extinction measurements were combined with an observation of an ULTRACAM flux standard star (see Appendix A and Table A.3), reduced using a larger target aperture radius of $2.5 \times$ FWHM, in order to measure the instrumental zeropoint for the night. The calibrated flux of the comparison star was then determined using the same target aperture radius as for

[1] https://github.com/HiPERCAM/hipercam

the flux standard star, which was then used to flux calibrate the target. When using optimally extracted photometry, the flux calibration was still performed on the data reduced using a standard aperture photometry extraction. This calibration was then applied to the optimally extracted photometry to prevent systematic absolute flux errors between the two methods. These flux calibration steps were performed using the CAM_CAL[2] package.

4.3 Method

I fit the flux calibrated eclipse photometry using the PYLCURVE[3] package, a python wrapper for LCURVE'S LROCHE routine (Copperwheat et al., 2010). In general, I follow the method developed in Chapter 3 which involves fitting the eclipse photometry in multiple filters simultaneously with eight free parameters. These are the effective temperatures, T_1 and T_2, which define the spectral energy distributions (SEDs) of both stars through the use of stellar atmosphere models (Claret et al., 2020; Husser et al., 2013); the stellar masses, M_1 and M_2; the binary inclination, i; the parallax, ϖ; the interstellar reddening, $E(B-V)$; and the time of mid-eclipse, T_0. With the use of mass-radius relations and a given (fixed) orbital period, P, the radii of both stars and the orbital separation of the binary can be defined, allowing model light curves to be generated for each filter.

For this work, however, I implement two changes to the methodology mentioned above, both regarding the spectral modelling of the secondary star:

1. In Chapter 3, PHOENIX stellar atmospheres (Husser et al., 2013) were used to model the SED of the secondary star. However, these models are limited to a minimum effective temperature of 2300 K, preventing the modelling of systems with brown dwarf companions. I have therefore switched to using the BT-Settl CIFIST

[2] https://github.com/Alex-J-Brown/cam_cal
[3] https://github.com/Alex-J-Brown/pylcurve

stellar atmosphere grid (Allard et al., 2012) which goes as low as 1200 K, allowing for a seamless transition to the brown dwarf regime and keeping the modelling consistent throughout.

2. It is well known that there are significant differences in the synthetic photometry of low mass stars calculated using different spectral models for a given effective temperature and surface gravity. This is most apparent for lower effective temperatures (< 3500 K), with models struggling to reproduce the transitions from M dwarfs to L dwarfs to T dwarfs (Saumon & Marley, 2008; Allard et al., 2012; Best et al., 2021). Rigidly defining the SED of the secondary from these spectral models could therefore introduce problems where the model photometry cannot reproduce the observed SED of the star in question to the precision of our observations. I counter this by allowing the secondary to have a separate effective temperature in each observed bandpass. Despite being allowed to vary, these individual filter-specific effective temperatures should be consistent with each other at a certain level. I implement this consistency requirement using priors to favour solutions where these effective temperatures are similar across the different filters.

In order to inform the priors on the filter-specific secondary temperatures mentioned in item 2, I again use the sample of 15 279 well-characterised M dwarfs (Morrell & Naylor, 2019). Cross-matching this sample with SDSS DR13 returns a sample of 5 222 M dwarfs, on which I then make colour cuts informed by synthetic photometry of the BT-SETTL-CIFIST model atmospheres ($4.0 < (u' - i') < 6.4$ and $1.5 < (g' - i') < 3.4$) to remove many of the extreme outliers. This leaves 4 158 M dwarfs with SDSS photometry. I then fit fifth-order polynomials to the measured effective temperature as a function of $u' - i'$ and $g' - i'$ colours individually, using an iterative sigma clipping procedure with a 3σ cut to remove any outliers that remain after the initial colour cuts (Figure 4.1). The standard deviations of the residuals of the remaining points are 80 K for a $u' - i'$ colour

and 30 K for a $g' - i'$ colour. I therefore implement Gaussian priors on the difference in effective temperature between the u' and i', and g' and i' bands of 80 K and 30 K respectively, both centred at zero. As, with this method, there are as many temperature measurements available for the secondary as filters used, I take the i_s-band measurement as being representative of the true secondary temperature. I make this choice based on it being the the band where the secondary is brightest and is therefore the most strongly constrained by the photometry.

As in Chapter 3, I use an MCMC method to fit each light curve, implemented through the PYTHON package, EMCEE[4] (Foreman-Mackey et al., 2013). I run each fit for a minimum of 10 000 steps using 100 walkers and inspect each fit manually for convergence and stability. Each system is first fit using a CO-core WD mass-radius relation (Bédard et al., 2020; Blouin et al., 2018; Tremblay et al., 2011) with the fit then being repeated using a He-core model (Panei et al., 2007) if the best-fit CO-core WD mass is below 0.5 $M_\odot$. If this subsequent fit using the He-core model is restricted by the upper mass limit of the He-core models – 0.5 $M_\odot$ – then we consider the WD to have a CO core-composition, if not then we assume the WD to possess a He core.

[4] https://emcee.readthedocs.io/en/stable/

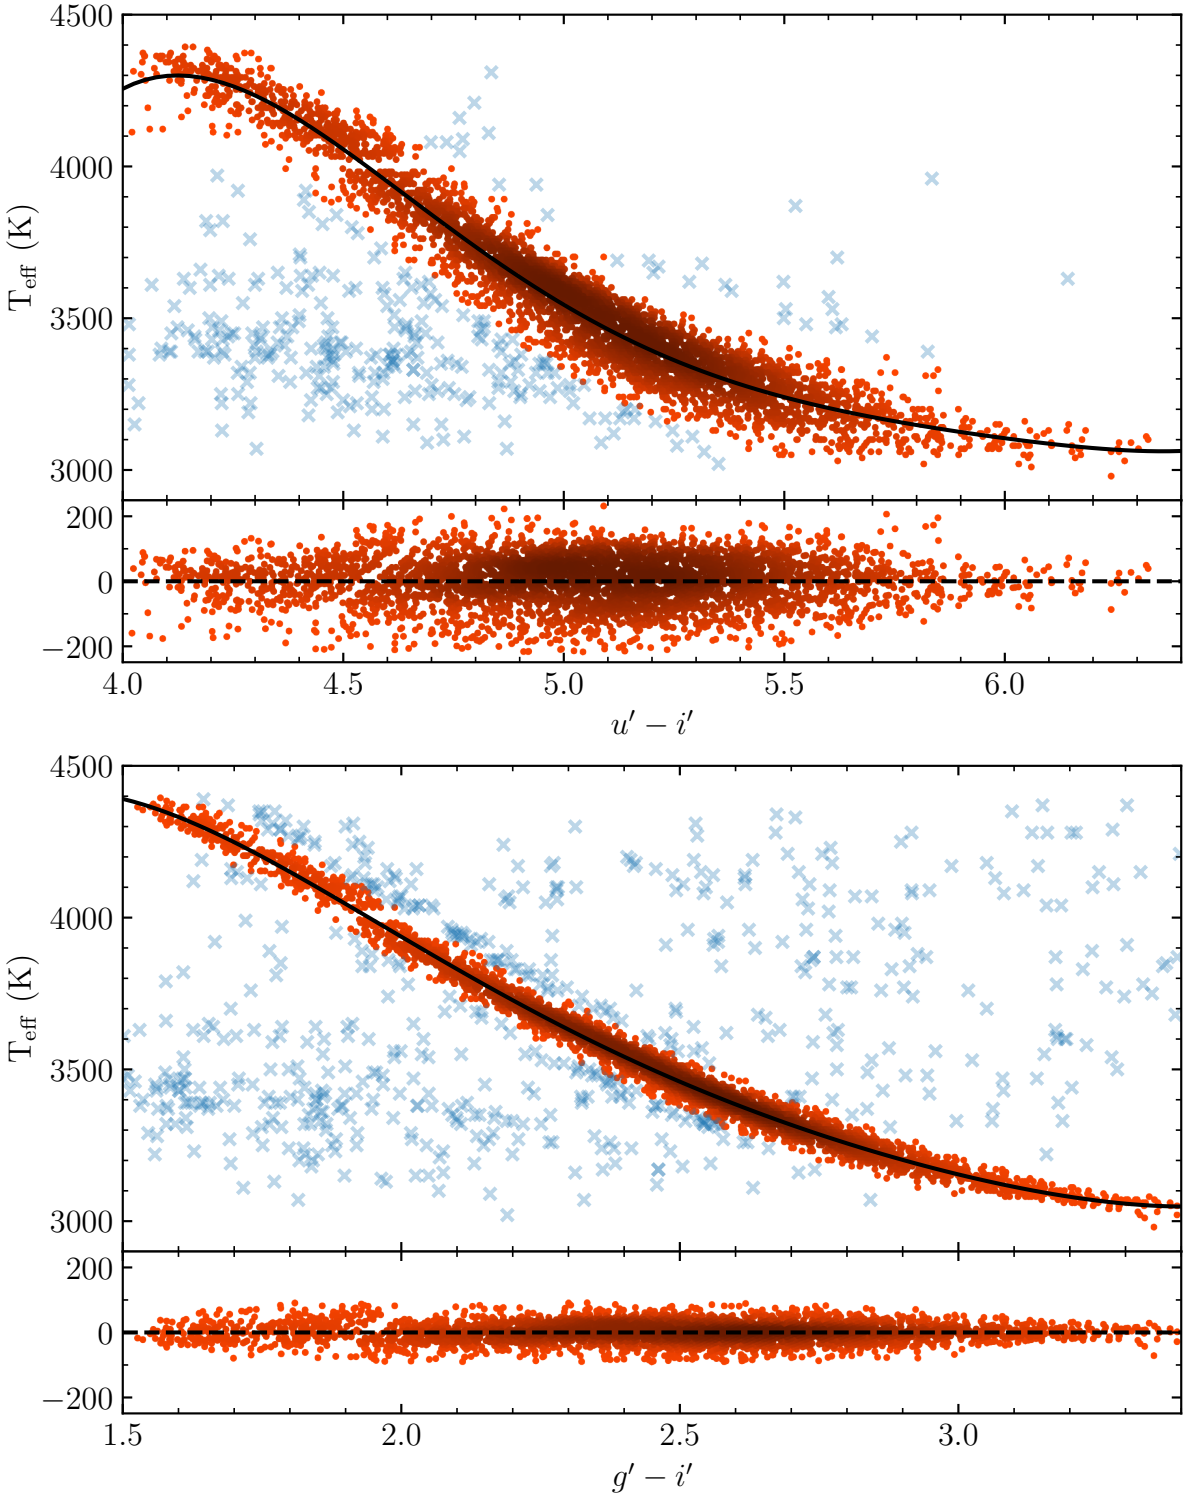

Figure 4.1: Effective temperatures of M dwarfs measured by Morrell & Naylor (2019) against their SDSS colours. Blue crosses show points discarded by the sigma clipping procedure and the solid black lines show the final polynomial fits to these sigma-clipped distributions. The residuals of these fits, from which we calculate the standard deviations, are shown in the panels below, with the dashed black lines showing the zero-level. The gap in the sample at an effective temperature of 4000 K is due to a discontinuity in the model grid used by Morrell & Naylor (2019).

4.4 Results

The results of the light curve fits are presented in Table 4.2 and Table 4.3. Note that of the 42 systems that I have followed-up, 9 do not have measured parameters because they either harbour magnetic WDs, or are strong candidates to harbour magnetic WDs (see section 4.5.4). The best-fit values are taken to be the median of the posterior distributions of the MCMC with lower and upper uncertainties taken as the 16th and 84th percentiles respectively. As in Chapter 3, the formal uncertainties from the MCMC do not include contributions from systematic errors and so I attempt to take this into account by adding estimated systematic uncertainties in quadrature with the formal uncertainties of the MCMC. I add 1.5 per cent in quadrature with the uncertainties on the primary temperature, T_1 (Gianninas et al., 2011), and 100 K in quadrature with the secondary temperature, T_2. I also add 1 per cent in quadrature with the WD mass, M_1, and 5 per cent in quadrature with the secondary mass, M_2 (for the reasons explained in Chapter 3). These contributions are included in the uncertainties shown in Table 4.2 and in all figures. An example ULTRACAM eclipse light curve and best-fit model is shown in Figure 4.2 with all best-fit light curves shown in Appendix B. Cornerplots for these MCMC fits are also presented in Appendix D.

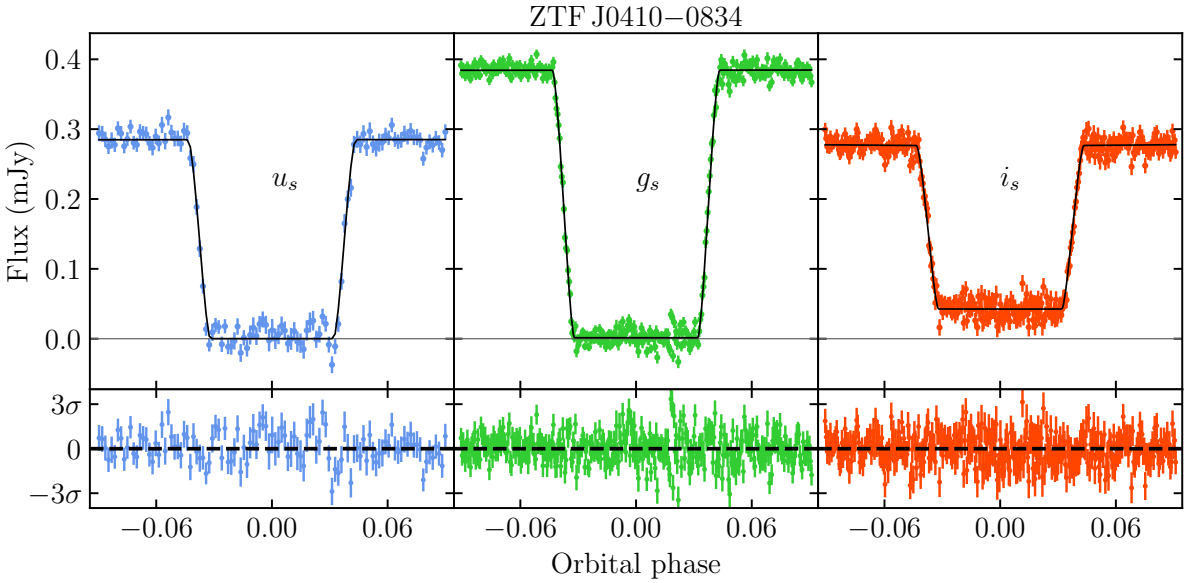

Figure 4.2: ULTRACAM $u_s\, g_s\, i_s$ eclipse light curve (coloured points) of ZTF J041016.82−083419.5 with the best-fit light curve model over-plotted in black and the residuals of this fit shown below. The zero-flux level is shown by the horizontal grey line.

Table 4.2: Best fit stellar parameters to the ULTRACAM eclipse photometry. Uncertainties include estimated systematic errors added in quadrature with the formal uncertainties of the MCMC. These estimated systematics are 1.5 per cent on T_1 (Gianninas et al., 2011), 100 K on T_2, 1 per cent for M_1, and 5 per cent for M_2.

Target	He/CO	T_1 (K)	M_1 ($M_\odot$)	R_1 ($R_\odot$)	$\log(g_1)$	T_2 (K)	M_2 ($M_\odot$)	R_2 ($R_\odot$)	R_2/R_{L1}
ZTF J0410−0834	He	14690^{+560}_{-550}	$0.355^{+0.015}_{-0.011}$	$0.0204^{+0.0004}_{-0.0006}$	$7.37^{+0.04}_{-0.03}$	2840^{+110}_{-110}	$0.123^{+0.009}_{-0.008}$	$0.151^{+0.008}_{-0.006}$	$0.680^{+0.033}_{-0.021}$
ZTF J0519+0925	He	10750^{+770}_{-580}	$0.391^{+0.019}_{-0.029}$	$0.0178^{+0.0009}_{-0.0005}$	$7.53^{+0.04}_{-0.08}$	2800^{+140}_{-110}	$0.177^{+0.014}_{-0.019}$	$0.214^{+0.012}_{-0.018}$	$0.842^{+0.079}_{-0.084}$
ZTF J0528+2156	CO	12100^{+700}_{-630}	$0.787^{+0.025}_{-0.025}$	$0.0105^{+0.0003}_{-0.0003}$	$8.29^{+0.04}_{-0.04}$	3130^{+110}_{-110}	$0.184^{+0.014}_{-0.013}$	$0.220^{+0.011}_{-0.009}$	$0.408^{+0.014}_{-0.011}$
ZTF J0537−2450	He	16100^{+440}_{-410}	$0.397^{+0.009}_{-0.007}$	$0.0191^{+0.0002}_{-0.0002}$	$7.48^{+0.02}_{-0.02}$	2970^{+100}_{-100}	$0.204^{+0.012}_{-0.011}$	$0.241^{+0.007}_{-0.005}$	$0.333^{+0.006}_{-0.004}$
ZTF J0615+0510	CO	15220^{+600}_{-510}	$0.560^{+0.011}_{-0.011}$	$0.0139^{+0.0002}_{-0.0002}$	$7.90^{+0.02}_{-0.02}$	3380^{+110}_{-110}	$0.533^{+0.030}_{-0.029}$	$0.547^{+0.013}_{-0.011}$	$0.531^{+0.008}_{-0.008}$
ZTF J0638+0910	CO	22500^{+1200}_{-1000}	$0.604^{+0.013}_{-0.011}$	$0.0136^{+0.0001}_{-0.0002}$	$7.95^{+0.02}_{-0.02}$	3320^{+110}_{-110}	$0.410^{+0.024}_{-0.022}$	$0.432^{+0.012}_{-0.008}$	$0.295^{+0.005}_{-0.004}$
ZTF J0639+1919	CO	15980^{+520}_{-520}	$0.701^{+0.011}_{-0.009}$	$0.0117^{+0.0001}_{-0.0001}$	$8.15^{+0.01}_{-0.01}$	3200^{+100}_{-100}	$0.210^{+0.011}_{-0.011}$	$0.246^{+0.004}_{-0.002}$	$0.398^{+0.004}_{-0.002}$
ZTF J0642+1314	CO	14560^{+540}_{-500}	$0.633^{+0.011}_{-0.008}$	$0.0127^{+0.0001}_{-0.0001}$	$8.03^{+0.02}_{-0.01}$	3110^{+100}_{-100}	$0.150^{+0.008}_{-0.008}$	$0.183^{+0.004}_{-0.001}$	$0.438^{+0.006}_{-0.002}$
ZTF J0651+1452	CO	13140^{+560}_{-670}	$0.515^{+0.019}_{-0.020}$	$0.0145^{+0.0003}_{-0.0003}$	$7.83^{+0.03}_{-0.04}$	3170^{+120}_{-110}	$0.242^{+0.018}_{-0.019}$	$0.276^{+0.012}_{-0.013}$	$0.589^{+0.018}_{-0.019}$
ZTF J0704−0201	CO	9280^{+230}_{-250}	$0.500^{+0.012}_{-0.015}$	$0.0143^{+0.0003}_{-0.0002}$	$7.82^{+0.02}_{-0.03}$	3300^{+100}_{-100}	$0.344^{+0.018}_{-0.020}$	$0.370^{+0.006}_{-0.010}$	$0.915^{+0.040}_{-0.043}$
ZTF J0717+1136	CO	21110^{+920}_{-750}	$0.528^{+0.016}_{-0.017}$	$0.0149^{+0.0003}_{-0.0003}$	$7.81^{+0.03}_{-0.03}$	3150^{+120}_{-110}	$0.296^{+0.020}_{-0.022}$	$0.326^{+0.013}_{-0.015}$	$0.320^{+0.008}_{-0.009}$
ZTF J0718−0852	CO	18940^{+870}_{-880}	$0.794^{+0.019}_{-0.018}$	$0.0106^{+0.0002}_{-0.0002}$	$8.28^{+0.03}_{-0.03}$	3120^{+110}_{-110}	$0.306^{+0.020}_{-0.019}$	$0.335^{+0.012}_{-0.011}$	$0.555^{+0.014}_{-0.012}$
ZTF J0804−0215	CO	13430^{+560}_{-550}	$0.577^{+0.010}_{-0.009}$	$0.0134^{+0.0001}_{-0.0001}$	$7.94^{+0.01}_{-0.01}$	$< 1510^{+260}_{-200}$	$< 0.069^{+0.007}_{-0.007}$	$0.098^{+0.002}_{-0.001}$	$0.377^{+0.008}_{-0.006}$
ZTF J0805−1430	He	26500^{+1200}_{-9000}	$0.393^{+0.013}_{-0.013}$	$0.0239^{+0.0007}_{-0.0006}$	$7.28^{+0.03}_{-0.04}$	3250^{+120}_{-110}	$0.291^{+0.020}_{-0.023}$	$0.331^{+0.013}_{-0.017}$	$0.586^{+0.016}_{-0.021}$
ZTF J0948+2538	CO	11290^{+480}_{-450}	$0.504^{+0.026}_{-0.024}$	$0.0145^{+0.0004}_{-0.0004}$	$7.82^{+0.05}_{-0.05}$	3120^{+120}_{-120}	$0.169^{+0.015}_{-0.014}$	$0.205^{+0.013}_{-0.012}$	$0.546^{+0.024}_{-0.024}$
ZTF J1022−0803	CO	8330^{+260}_{-250}	$0.605^{+0.027}_{-0.025}$	$0.0127^{+0.0003}_{-0.0003}$	$8.01^{+0.04}_{-0.04}$	3170^{+110}_{-110}	$0.405^{+0.030}_{-0.029}$	$0.428^{+0.021}_{-0.020}$	$0.620^{+0.023}_{-0.021}$
ZTF J1026−1013	He	19320^{+710}_{-670}	$0.376^{+0.012}_{-0.010}$	$0.0214^{+0.0004}_{-0.0007}$	$7.35^{+0.04}_{-0.02}$	2840^{+110}_{-110}	$0.105^{+0.008}_{-0.006}$	$0.134^{+0.007}_{-0.004}$	$0.558^{+0.021}_{-0.012}$

Target	He/CO	T_1 (K)	M_1 ($M_\odot$)	R_1 ($R_\odot$)	$\log(g_1)$	T_2 (K)	M_2 ($M_\odot$)	R_2 ($R_\odot$)	R_2/R_{L1}
ZTF J1049−1755	He	13000^{+440}_{-460}	$0.426^{+0.010}_{-0.007}$	$0.0173^{+0.0001}_{-0.0002}$	$7.59^{+0.02}_{-0.01}$	3170^{+100}_{-110}	$0.198^{+0.012}_{-0.010}$	$0.235^{+0.007}_{-0.003}$	$0.402^{+0.008}_{-0.003}$
ZTF J1220+0821	CO	10170^{+270}_{-260}	$0.580^{+0.017}_{-0.018}$	$0.0132^{+0.0003}_{-0.0002}$	$7.96^{+0.03}_{-0.03}$	3140^{+110}_{-110}	$0.275^{+0.019}_{-0.020}$	$0.306^{+0.012}_{-0.013}$	$0.157^{+0.004}_{-0.004}$
ZTF J1256+2117	CO	5073^{+79}_{-79}	$0.479^{+0.010}_{-0.009}$	$0.0141^{+0.0001}_{-0.0001}$	$7.82^{+0.02}_{-0.01}$	2950^{+100}_{-100}	$0.101^{+0.005}_{-0.005}$	$0.125^{+0.001}_{-0.001}$	$0.152^{+0.001}_{-0.001}$
ZTF J1302−0032	CO	11790^{+400}_{-330}	$0.811^{+0.021}_{-0.016}$	$0.0102^{+0.0002}_{-0.0002}$	$8.33^{+0.03}_{-0.02}$	3030^{+100}_{-100}	$0.179^{+0.012}_{-0.010}$	$0.216^{+0.008}_{-0.005}$	$0.502^{+0.013}_{-0.009}$
ZTF J1341−0626	CO	58300^{+8400}_{-8700}	$0.509^{+0.038}_{-0.035}$	$0.0225^{+0.0009}_{-0.0016}$	$7.43^{+0.09}_{-0.04}$	2800^{+210}_{-220}	$0.126^{+0.015}_{-0.009}$	$0.159^{+0.018}_{-0.007}$	$0.617^{+0.062}_{-0.021}$
ZTF J1400+0814	CO	13340^{+650}_{-610}	$0.563^{+0.009}_{-0.008}$	$0.0137^{+0.0001}_{-0.0001}$	$7.92^{+0.01}_{-0.01}$	2970^{+100}_{-100}	$0.232^{+0.012}_{-0.012}$	$0.268^{+0.003}_{-0.001}$	$0.418^{+0.003}_{-0.001}$
ZTF J1404+0655	CO	14980^{+470}_{-460}	$0.736^{+0.016}_{-0.015}$	$0.0113^{+0.0002}_{-0.0002}$	$8.20^{+0.02}_{-0.02}$	3100^{+100}_{-100}	$0.409^{+0.023}_{-0.023}$	$0.432^{+0.010}_{-0.010}$	$0.884^{+0.045}_{-0.031}$
ZTF J1405+1039	He	29900^{+9000}_{-1100}	$0.404^{+0.008}_{-0.008}$	$0.0279^{+0.0006}_{-0.0006}$	$7.15^{+0.02}_{-0.02}$	3430^{+130}_{-140}	$0.085^{+0.005}_{-0.005}$	$0.112^{+0.003}_{-0.003}$	$0.234^{+0.004}_{-0.004}$
ZTF J1407+2115	He	10870^{+350}_{-350}	$0.406^{+0.018}_{-0.014}$	$0.0173^{+0.0004}_{-0.0004}$	$7.57^{+0.04}_{-0.03}$	3160^{+110}_{-110}	$0.263^{+0.021}_{-0.016}$	$0.296^{+0.015}_{-0.009}$	$0.702^{+0.029}_{-0.016}$
ZTF J1458+1313	CO	9420^{+260}_{-260}	$0.581^{+0.010}_{-0.010}$	$0.0131^{+0.0001}_{-0.0001}$	$7.97^{+0.01}_{-0.01}$	$< 1730^{+240}_{-270}$	$< 0.067^{+0.006}_{-0.006}$	$0.095^{+0.001}_{-0.000}$	$0.446^{+0.011}_{-0.006}$
ZTF J1626−1018	CO	36700^{+2700}_{-2700}	$0.499^{+0.015}_{-0.012}$	$0.0180^{+0.0002}_{-0.0003}$	$7.62^{+0.02}_{-0.01}$	3180^{+110}_{-110}	$0.212^{+0.013}_{-0.011}$	$0.259^{+0.008}_{-0.003}$	$0.425^{+0.008}_{-0.004}$
ZTF J1634−2713	He	10680^{+790}_{-630}	$0.436^{+0.042}_{-0.054}$	$0.0166^{+0.0013}_{-0.0009}$	$7.64^{+0.09}_{-0.12}$	2400^{+130}_{-120}	$0.134^{+0.016}_{-0.020}$	$0.163^{+0.019}_{-0.022}$	$0.759^{+0.128}_{-0.099}$
ZTF J1644+2434	He	13270^{+520}_{-460}	$0.382^{+0.020}_{-0.018}$	$0.0188^{+0.0007}_{-0.0007}$	$7.47^{+0.05}_{-0.05}$	2500^{+110}_{-110}	$0.103^{+0.009}_{-0.009}$	$0.129^{+0.009}_{-0.008}$	$0.607^{+0.033}_{-0.028}$
ZTF J1802−0054	He	10770^{+630}_{-500}	$0.458^{+0.019}_{-0.021}$	$0.0160^{+0.0004}_{-0.0003}$	$7.69^{+0.03}_{-0.04}$	3150^{+110}_{-110}	$0.150^{+0.010}_{-0.011}$	$0.182^{+0.008}_{-0.010}$	$0.319^{+0.010}_{-0.012}$
ZTF J1828+2308	CO	16620^{+560}_{-650}	$0.594^{+0.009}_{-0.008}$	$0.0134^{+0.0001}_{-0.0001}$	$7.96^{+0.01}_{-0.01}$	$< 2290^{+110}_{-120}$	$< 0.068^{+0.007}_{-0.006}$	$0.096^{+0.002}_{-0.000}$	$0.392^{+0.009}_{-0.005}$
ZTF J1954+1019	CO	21500^{+1000}_{-1100}	$0.509^{+0.015}_{-0.012}$	$0.0154^{+0.0002}_{-0.0002}$	$7.77^{+0.03}_{-0.02}$	3480^{+110}_{-110}	$0.449^{+0.028}_{-0.026}$	$0.470^{+0.016}_{-0.013}$	$0.523^{+0.012}_{-0.010}$

Table 4.3: Best fit binary parameters to the ULTRACAM eclipse photometry. The orbital periods are listed here for reference but are not fitted parameters and so do not have corresponding uncertainties. The *Gaia* DR3 parallax measurements are included for comparison.

Target	i (°)	a ($R_\odot$)	$E(B-V)$	$\varpi_{\rm UCAM}$	$\varpi_{\rm Gaia}$	T_0 (BMJD(TDB))	P (d)
ZTF J0410−0834	$86.6^{+2.1}_{-1.7}$	$0.616^{+0.009}_{-0.006}$	$0.031^{+0.017}_{-0.017}$	$3.863^{+0.091}_{-0.078}$	4.07 ± 0.11	59646.0489782(16)	0.0811093
ZTF J0519+0925	$76.3^{+1.1}_{-0.6}$	$0.715^{+0.012}_{-0.020}$	$0.112^{+0.028}_{-0.023}$	$2.835^{+0.140}_{-0.140}$	2.92 ± 0.30	59251.0519387(57)	0.0929131
ZTF J0528+2156	$87.7^{+1.4}_{-1.0}$	$1.546^{+0.017}_{-0.016}$	$0.090^{+0.020}_{-0.021}$	$5.666^{+0.104}_{-0.111}$	5.59 ± 0.13	59932.215321(52)	0.2259952
ZTF J0537−2450	$88.1^{+0.7}_{-0.6}$	$1.688^{+0.014}_{-0.010}$	$0.015^{+0.011}_{-0.010}$	$4.580^{+0.044}_{-0.047}$	4.574 ± 0.049	59251.2246115(52)	0.3277936
ZTF J0615+0510	$85.0^{+0.7}_{-0.7}$	$2.146^{+0.015}_{-0.014}$	$0.019^{+0.019}_{-0.013}$	$3.163^{+0.060}_{-0.051}$	3.166 ± 0.081	59280.12536567(83)	0.3481742
ZTF J0638+0910	$88.2^{+0.5}_{-0.6}$	$3.197^{+0.024}_{-0.019}$	$0.021^{+0.018}_{-0.015}$	$1.709^{+0.047}_{-0.047}$	1.65 ± 0.14	59252.1564861(10)	0.6576453
ZTF J0639+1919	$88.9^{+0.7}_{-0.7}$	$1.659^{+0.008}_{-0.005}$	$0.028^{+0.013}_{-0.015}$	$5.394^{+0.070}_{-0.075}$	5.387 ± 0.085	59251.17799186(52)	0.2593556
ZTF J0642+1314	$89.1^{+0.7}_{-0.8}$	$1.195^{+0.006}_{-0.003}$	$0.022^{+0.016}_{-0.013}$	$3.583^{+0.075}_{-0.073}$	3.77 ± 0.20	59252.10345653(59)	0.1710542
ZTF J0651+1452	$85.3^{+1.5}_{-1.0}$	$1.166^{+0.016}_{-0.017}$	$0.037^{+0.016}_{-0.018}$	$2.567^{+0.073}_{-0.076}$	2.70 ± 0.17	59252.2124933(16)	0.1677075
ZTF J0704−0201	$74.3^{+0.4}_{-0.2}$	$1.079^{+0.006}_{-0.010}$	$0.052^{+0.011}_{-0.013}$	$3.715^{+0.076}_{-0.075}$	3.643 ± 0.088	59253.2216462(43)	0.1413708
ZTF J0717+1136	$84.9^{+0.4}_{-0.3}$	$2.326^{+0.027}_{-0.030}$	$0.018^{+0.017}_{-0.012}$	$2.812^{+0.065}_{-0.072}$	2.74 ± 0.13	59251.1312794(93)	0.4527638
ZTF J0718−0852	$84.6^{+0.7}_{-0.7}$	$1.563^{+0.014}_{-0.013}$	$0.064^{+0.017}_{-0.019}$	$2.157^{+0.062}_{-0.058}$	2.39 ± 0.22	59283.1026109(12)	0.2158113
ZTF J0804−0215	$85.3^{+0.1}_{-0.1}$	$0.889^{+0.007}_{-0.004}$	$0.027^{+0.015}_{-0.014}$	$5.631^{+0.092}_{-0.089}$	5.47 ± 0.11	59646.09050723(97)	0.1209762
ZTF J0805−1430	$81.0^{+1.0}_{-0.7}$	$1.260^{+0.015}_{-0.019}$	$0.010^{+0.012}_{-0.008}$	$1.102^{+0.034}_{-0.039}$	1.39 ± 0.16	59646.18599526(74)	0.1981669
ZTF J0948+2538	$79.9^{+0.6}_{-0.6}$	$1.003^{+0.017}_{-0.017}$	$0.018^{+0.008}_{-0.009}$	$2.911^{+0.116}_{-0.108}$	2.94 ± 0.26	59239.2668295(95)	0.1418270
ZTF J1022−0803	$76.9^{+0.6}_{-0.6}$	$1.592^{+0.024}_{-0.023}$	$0.019^{+0.013}_{-0.012}$	$5.750^{+0.158}_{-0.160}$	5.58 ± 0.21	59280.2576703(43)	0.2314179
ZTF J1026−1013	$87.4^{+1.6}_{-1.5}$	$0.677^{+0.008}_{-0.006}$	$0.033^{+0.010}_{-0.010}$	$2.027^{+0.075}_{-0.063}$	1.65 ± 0.19	59237.2453759(14)	0.0929868

Target	i (°)	a ($R_\odot$)	$E(B-V)$	$\varpi_{\rm UCAM}$	$\varpi_{\rm Gaia}$	T_0 (BMJD(TDB))	P (d)
ZTF J1049−1755	$88.6^{+1.1}_{-1.1}$	$1.407^{+0.012}_{-0.006}$	$0.029^{+0.006}_{-0.008}$	$2.579^{+0.070}_{-0.070}$	2.47 ± 0.17	59238.3654607(11)	0.2447332
ZTF J1220+0821	$87.5^{+0.2}_{-0.2}$	$4.592^{+0.052}_{-0.056}$	$0.024^{+0.004}_{-0.008}$	$3.874^{+0.093}_{-0.106}$	3.53 ± 0.17	59252.2559984(25)	1.2329254
ZTF J1256+2117	$89.8^{+0.2}_{-0.2}$	$2.374^{+0.013}_{-0.012}$	$0.005^{+0.006}_{-0.004}$	$22.221^{+0.095}_{-0.094}$	22.171 ± 0.096	59641.3540758(33)	0.5560572
ZTF J1302−0032	$86.0^{+0.5}_{-0.6}$	$1.268^{+0.012}_{-0.008}$	$0.016^{+0.013}_{-0.011}$	$8.554^{+0.071}_{-0.068}$	8.555 ± 0.073	59252.387889(43)	0.1661310
ZTF J1341−0626	$86.8^{+2.4}_{-2.7}$	$0.764^{+0.018}_{-0.017}$	$0.030^{+0.009}_{-0.009}$	$0.894^{+0.089}_{-0.093}$	0.97 ± 0.12	59237.3073649(28)	0.0969505
ZTF J1400+0814	$89.2^{+0.6}_{-0.7}$	$1.589^{+0.006}_{-0.005}$	$0.005^{+0.008}_{-0.004}$	$2.155^{+0.070}_{-0.075}$	1.58 ± 0.30	59253.2966645(14)	0.2602766
ZTF J1404+0655	$84.5^{+1.0}_{-0.9}$	$1.342^{+0.010}_{-0.009}$	$0.025^{+0.007}_{-0.009}$	$2.538^{+0.059}_{-0.057}$	2.24 ± 0.14	59239.3665054(12)	0.1683096
ZTF J1405+1039	$88.5^{+0.4}_{-0.3}$	$1.389^{+0.009}_{-0.009}$	$0.016^{+0.010}_{-0.009}$	$0.752^{+0.031}_{-0.024}$	0.78 ± 0.26	59251.334651(12)	0.2714122
ZTF J1407+2115	$86.5^{+2.4}_{-2.0}$	$1.008^{+0.016}_{-0.011}$	$0.051^{+0.010}_{-0.013}$	$4.077^{+0.072}_{-0.070}$	4.079 ± 0.091	59643.3349542(63)	0.1432802
ZTF J1458+1313	$86.8^{+0.1}_{-0.2}$	$0.742^{+0.004}_{-0.003}$	$0.025^{+0.009}_{-0.012}$	$5.067^{+0.152}_{-0.154}$	4.86 ± 0.21	59252.3531663(17)	0.0920516
ZTF J1626−1018	$88.7^{+0.9}_{-1.1}$	$1.503^{+0.013}_{-0.010}$	$0.291^{+0.007}_{-0.013}$	$1.733^{+0.087}_{-0.085}$	1.91 ± 0.20	59253.3679108(15)	0.2530067
ZTF J1634−2713	$80.6^{+2.1}_{-1.5}$	$0.637^{+0.020}_{-0.028}$	$0.182^{+0.029}_{-0.026}$	$4.127^{+0.238}_{-0.230}$	4.24 ± 0.26	59253.3310632(36)	0.0780396
ZTF J1644+2434	$80.3^{+0.6}_{-0.7}$	$0.614^{+0.011}_{-0.011}$	$0.031^{+0.017}_{-0.015}$	$2.197^{+0.087}_{-0.086}$	2.43 ± 0.22	59283.3945858(11)	0.0801054
ZTF J1802−0054	$84.2^{+0.3}_{-0.3}$	$1.485^{+0.020}_{-0.023}$	$0.114^{+0.028}_{-0.024}$	$4.700^{+0.077}_{-0.101}$	4.38 ± 0.15	59646.3585478(21)	0.2690033
ZTF J1828+2308	$88.7^{+0.1}_{-0.3}$	$0.852^{+0.005}_{-0.002}$	$0.088^{+0.014}_{-0.018}$	$4.955^{+0.079}_{-0.077}$	4.914 ± 0.097	59695.37741036(84)	0.1120067
ZTF J1954+1019	$84.5^{+0.7}_{-0.8}$	$1.901^{+0.020}_{-0.016}$	$0.078^{+0.022}_{-0.023}$	$3.495^{+0.051}_{-0.050}$	3.449 ± 0.057	59697.3389707(12)	0.3102884

4.5 Discussion

4.5.1 Comparison with previous parameters

Initial parameter estimates for these systems were made by fitting the ZTF time-series photometry alongside photometric measurements from other surveys, where available, covering a wide wavelength range (van Roestel et al. in preparation). Comparing the parameters determined from the three-band eclipse photometry against these initial estimates demonstrates general agreement between the two methods (Figure 4.3). The WD temperatures and masses, in particular, show good agreement but there are some significant differences in the measured temperatures and masses for certain systems. The standard deviations of the histograms in Figure 4.3, however, demonstrate that the uncertainties are underestimated in some way. As the method I have used in this work has been shown to retrieve accurate parameters (see Chapter 3), this seems to imply an underestimation in the uncertainties determined by combining SED fitting with ZTF photometry. This may be due, in part, to the survey SED data used by van Roestel et al. (in preparation) being taken at a range of different orbital phases and therefore suffering from increased systematics due to ellipsoidal modulation or reflection effect.

In terms of precision, the parameters determined using the high-speed eclipse photometry are typically more precise than those measured by van Roestel et al. (in preparation). This is most apparent for the primary and secondary masses with a median uncertainty in the WD mass from the ULTRACAM photometry of 2.6 per cent, and 7.2 per cent for the secondary mass. These values are 6.0 per cent and 13.7 per cent respectively from the ZTF photometry for the same systems and so the ULTRACAM measurements are typically a factor of 2 more precise. This is very likely due to the high time resolution of the ULTRACAM photometry, enabling the duration of the eclipse as well as the ingress and egress to be measured very precisely – particularly important for constraining the radii of the two stars and therefore their masses.

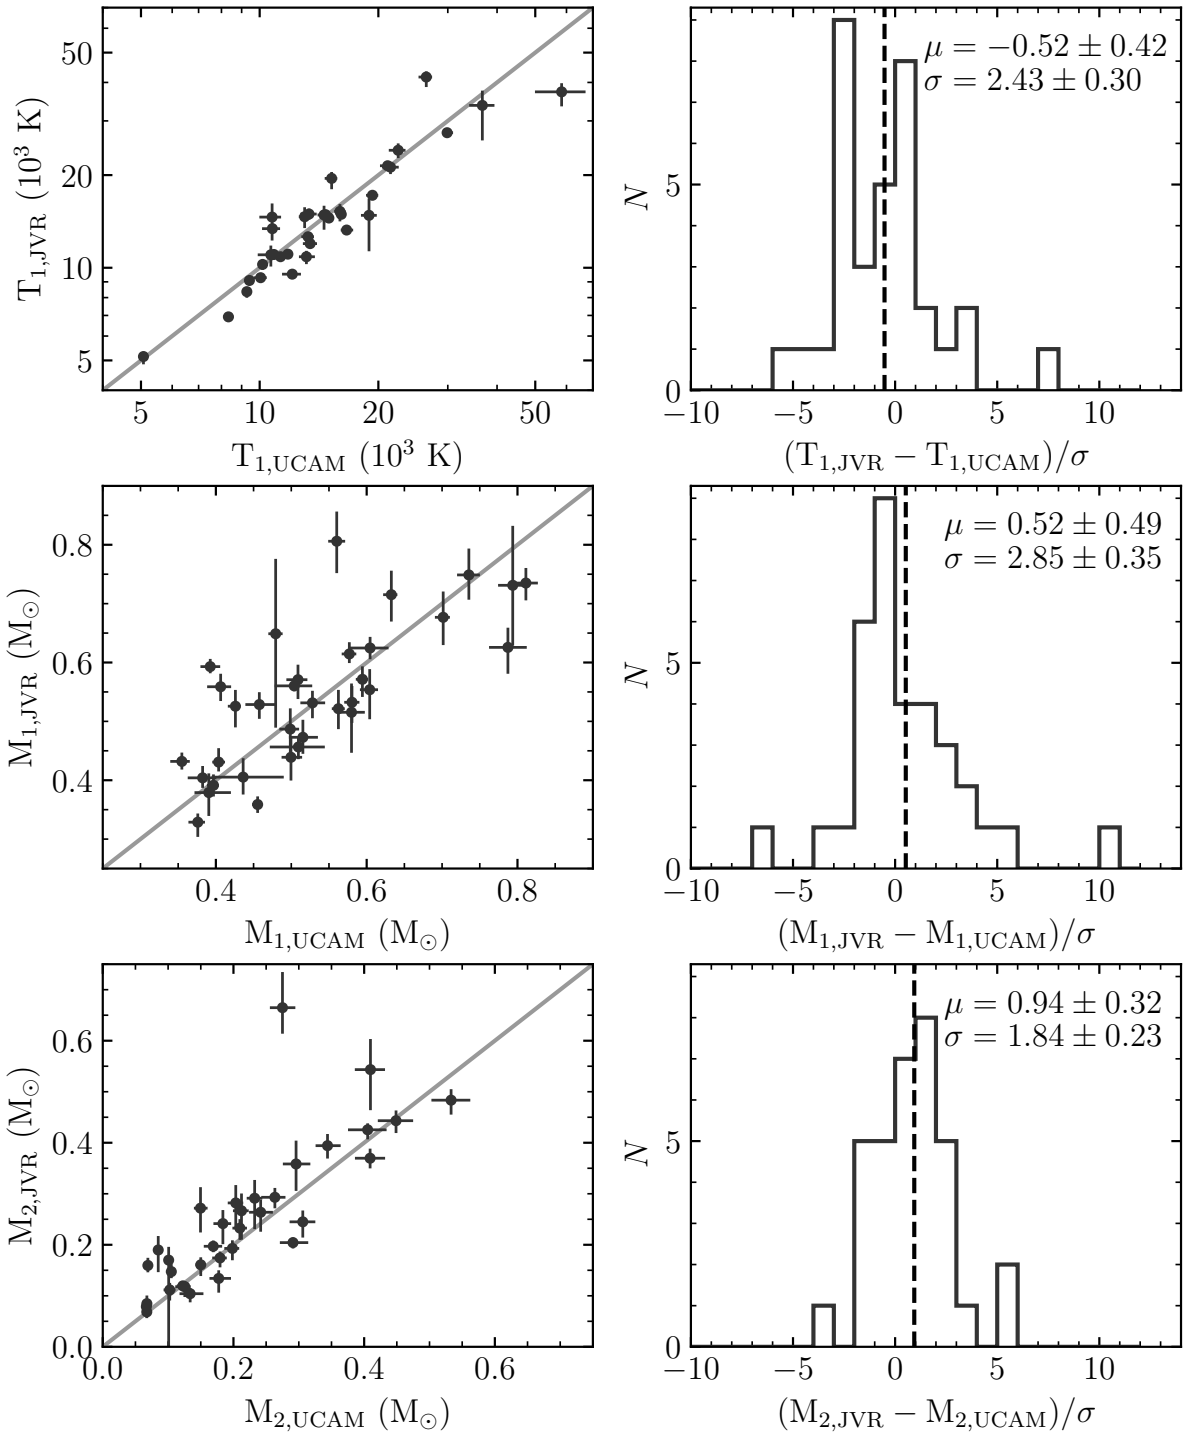

Figure 4.3: *Left*: Comparison of the parameters determined from the NTT-ULTRACAM photometry (subscript UCAM) against the initial parameters of van Roestel et al. (in preparation; subscript JVR) using ZTF photometry. *Right*: Histograms showing the differences in determined values between the two methods, normalised by their combined standard deviation , σ, where $\sigma = \sqrt{\sigma_{JVR}^2 + \sigma_{UCAM}^2}$. The mean and standard deviation of each histogram is shown in the plot. Overall, the two methods are in good agreement with similar mean values, however the standard deviations are all greater than unity and so there is likely some contribution to the uncertainties that is not taken into account.

In addition to the initial parameter estimates discussed above, two of the systems fit in this work have been included in previously published analyses – ZTF J1256+2117 and ZTF J1644+2434. Comparisons with these previous works are made below.

ZTF J1256+2117

ZTF J1256+2117 was previously fitted by Rebassa-Mansergas et al. (2021), using the Virtual Observatory SED Analyser (VOSA) to fit the available survey photometry. Out of the 112 systems that they analysed, 13 systems were determined to possess a WD with a mass below 0.2 $M_\odot$. It is not known how such low mass WDs could form in PCEBs with low-mass main sequence companions – with any mass transfer initiating a common envelope phase in which the envelope would most likely not gain sufficient energy to be ejected, leading to a merger scenario (Rebassa-Mansergas et al., 2021). This system is – as far as we know – the only one of these 13 systems that eclipses, enabling a valuable check on the system parameters. The fit to the eclipse photometry determines the WD mass to be 0.48 ± 0.01 $M_\odot$, discrepant with the 0.155 ± 0.02 $M_\odot$ obtained from VOSA by over 14σ. The WD effective temperatures determined from the two methods are also significantly different, with the VOSA fit finding a temperature of 4250 ± 150 K and the fit to the eclipse photometry finding a higher temperature of 5073 ± 79 K – discrepant by around 4.8σ. One of the causes for this discrepancy could be that there is a strong temperature–radius (and therefore mass) degeneracy in the VOSA SED fitting, such that a small preference for a particular WD temperature entirely dictates the radius required to match the observed luminosity at that temperature as $R \propto T^{-2}$. In this low-temperature range, WD SEDs – particularly from survey data where the two binary components can not be disentangled – will likely not provide strong enough constraints on the WD temperature to accurately determine the WD radius and therefore mass. Additionally, the nature of this system – a faint WD that is dominated by the flux of the secondary – means that any relatively small change in the parameters of the M dwarf

will need to be compensated for by a relatively large change in the parameters of the WD. Together, these could explain the extreme discrepancy between the two methods, but spectroscopic follow-up of this system will be required to confirm the cause along with the true WD mass.

ZTF J1644+2434

ZTF J1644+2434 was one of the four deeply eclipsing PCEBs found and fitted by Kosakowski et al. (2022). For this target in particular they did not detect the eclipse minimum and so their parameters from the light curve fit represent limits rather than specific values. As would be expected, the light curve fit to the ULTRACAM photometry is consistent with these parameter limits. As well as fitting the eclipse light curve, Kosakowski et al. (2022) performed a spectroscopic fit to the WD, determining the effective temperature, surface gravity, and mass (determined from the surface gravity using CO-core composition models). From the fit to the ULTRACAM photometry I find an effective temperature of 13270 ± 490 K, cooler than the 14900 ± 760 K determined by their spectroscopic fit but still consistent to within 2σ. For the WD mass there is a little more deviation, with my fit finding a WD mass of 0.38 ± 0.02 $M_{\odot}$, 2.3σ below the 0.55 ± 0.07 $M_{\odot}$ found from their spectroscopic fit and suggesting a He-core composition rather than a CO-core. For the companion, Kosakowski et al. (2022) estimate a mass of 0.084 ± 0.004 $M_{\odot}$ by fitting the Pan-STARRS SED with a composite model, placing it close to the hydrogen-burning limit. Here, I find a higher mass of 0.103 ± 0.009 $M_{\odot}$ from the light curve fit, taking it into more typically stellar territory. Again though, these two values are consistent to within 2σ. Overall, the fit to the ULTRACAM photometry is fully consistent with their light curve fit and consistent with their spectroscopic and Pan-STARRS SED fits at around the 2σ level.

4.5.2 Brown dwarf companions

WDs with brown dwarf (BD) companions are rare, with as few as around 0.5 per cent of WDs expected to have substellar partners (Steele et al., 2011). Eclipsing examples are, predictably, even rarer with only four systems currently confirmed (Beuermann et al., 2013; Littlefair et al., 2014; Parsons et al., 2017b; Casewell et al., 2020b; van Roestel et al., 2021). These eclipsing WD + BD binaries are valuable as they are one of the few places where both the BD's radii and mass can be measured precisely and are therefore important benchmarks for BD models. Additionally, as some of the lowest mass objects thought to survive the common-envelope (Casewell et al., 2018), BDs in PCEBs occupy an important area of the parameter space when studying common-envelope evolution, with the study of the common-envelope phase in this low-mass regime having implications for systems with planetary mass companions (Vanderburg et al., 2020).

In the ULTRACAM follow-up photometry, I have found three systems so far that the light curve fits suggest as having BD companions. These are ZTF J0804−0215, ZTF J1458+1313, and ZTF J1828+2308. As my mass-radius relation for M dwarfs Figure 3.2 is horizontal below 0.07 $M_\odot$ – and therefore uninformative in this regime – the best fit secondary masses can only be regarded as upper limits. Additionally, as none of the secondaries for these systems are detected in-eclipse, only an upper limit can be given for their effective temperatures. One of these systems, ZTF J1828+2308, has a high secondary temperature for a brown dwarf. In order to rule out problems with the photometry, I stack the in-eclipse images (Figure 4.4). This reveals a faint ($G = 20.88$ mag) source 2.79 arcsec away from the target which results in an erroneous slight 'detection' in eclipse and therefore a higher than expected temperature. The true upper limit for the secondary temperature will be lower than given by the fit.

In addition to these three systems with substellar companions, I have measured one with a companion mass just above the hydrogen-burning limit, ZTF J1405+1039. The

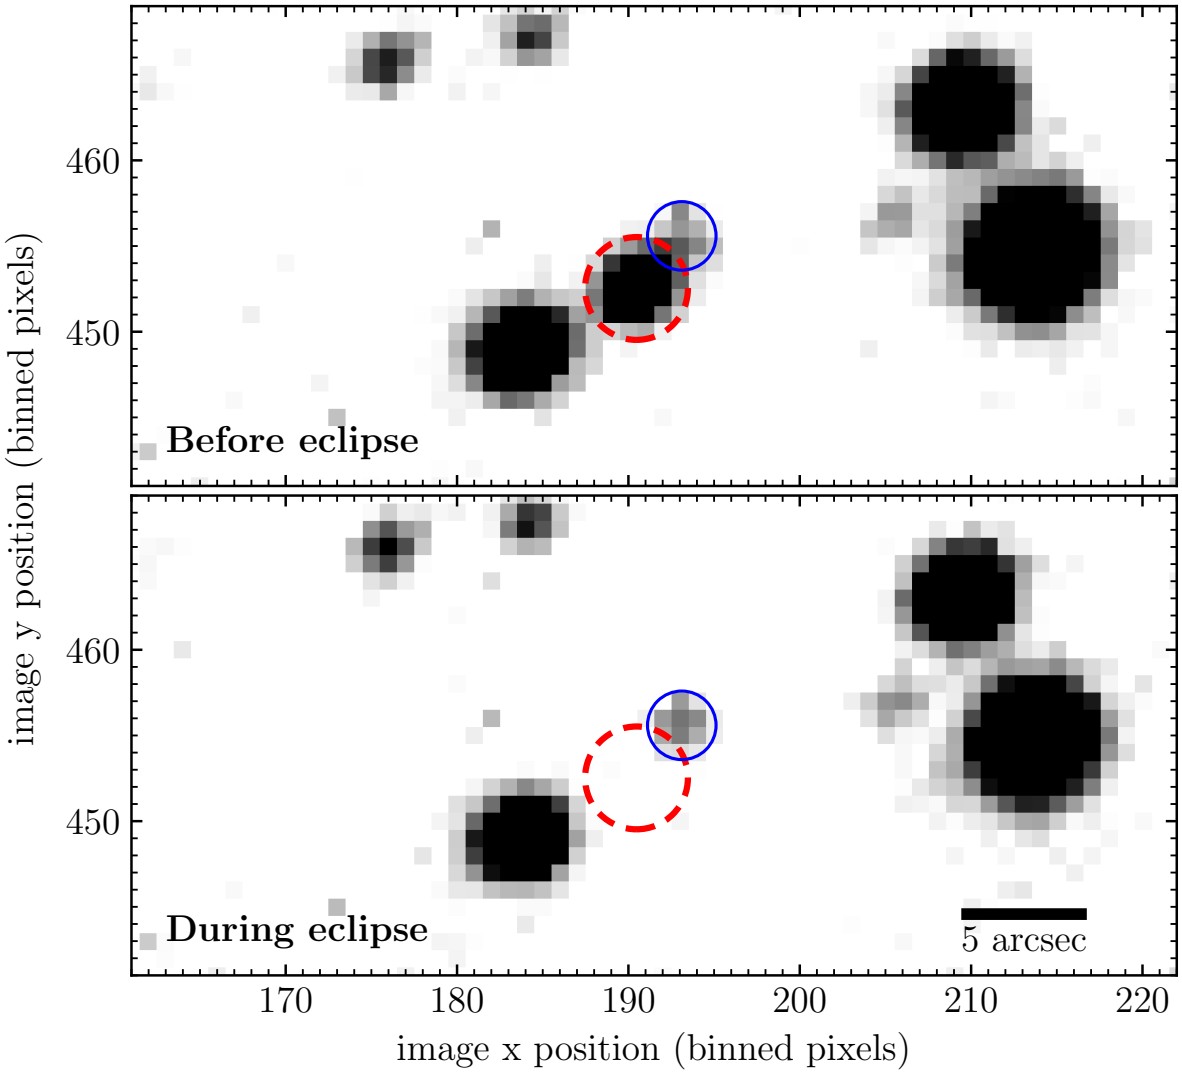

Figure 4.4: Stacked images of ZTF J1828+2308 taken with ULTRACAM in the i_s-band before and during the eclipse. The red dashed aperture shows the location of ZTF J1828+2308 itself while the solid blue aperture shows the fainter background source 2.79" away (Gaia DR3 4529477702982880512) that is marginally affecting the in-eclipse photometry.

best-fit parameters for this system suggest that the secondary is significantly hotter than would be expected for its mass (shown as the blue point in Figure 4.5). Again, I stack the in-eclipse images to rule out problems in the photometry (Figure 4.6), demonstrating that the source is indeed detected in-eclipse, with an i_s-band magnitude of around 22.4 mag. The most likely explanation for this seems to be that ZTF J1405+1039 is actually a triple system, with a tertiary companion contributing a significant fraction of the in-eclipse flux. Unfortunately, its distance is such that any astrometric signal due to a tertiary companion will be far below the detection limit of *Gaia* making confirmation of this hypothesis is difficult.

4.5.3 ZZ Ceti WDs

ZZ Cetis are pulsating WDs, possessing hydrogen atmospheres and pulsation periods ranging from tens of seconds to tens of minutes (Fontaine & Brassard, 2008; Winget & Kepler, 2008; Romero et al., 2022b). The presence of pulsations enable asteroseismological analyses to be performed, providing insight into the internal structure of the WD which is otherwise concealed by their highly stratified nature. In PCEBs, the possibility of measuring the internal structure of the WD is especially interesting as it can reveal how the WD itself is affected by the common-envelope phase (Hermes et al., 2015).

Previously, only one ZZ Ceti WD in a detached eclipsing binary was known (Parsons et al., 2020). This system is a double WD binary, however, and as such its evolutionary history is less well defined, with the number of common-envelope events it has passed through being uncertain. ZZ Cetis found in WD-main sequence PCEBs do not have this problem with their evolutionary past known to comprise of a single common-envelope phase. These systems are therefore potentially very interesting systems to find. Currently there are only two known ZZ Ceti WDs in detached PCEBs, SDSS J1136+0409 (Pyrzas et al., 2015) and GD 1400 (Fontaine et al., 2003; Farihi & Christopher, 2004)

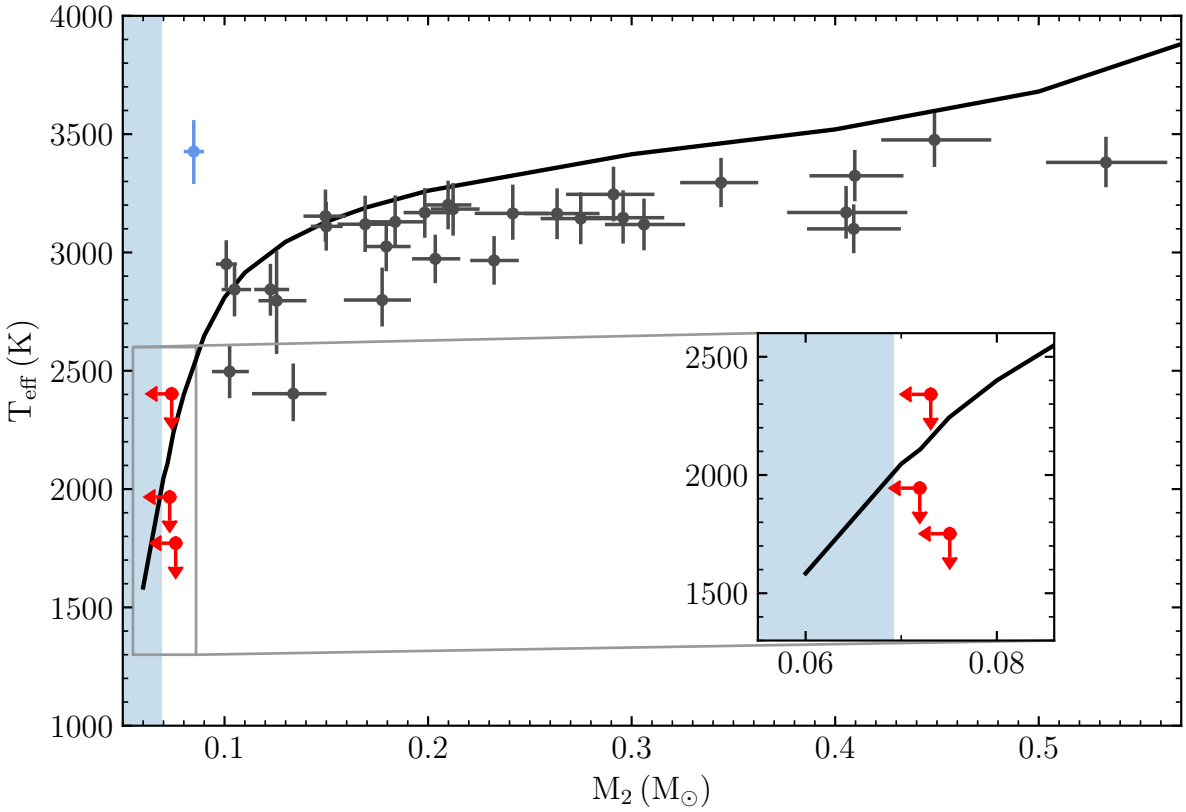

Figure 4.5: Measured masses and effective temperatures of the M dwarf components with an inset plot zoomed in around the brown dwarfs (which are shown in red). The solid black line shows the 1 Gyr track from Baraffe et al. (2015) and the shaded blue area denotes the region where our mass-radius relation is horizontal (i.e. the radius is constant in this mass range). For the brown dwarfs I plot the masses and temperatures as upper limits centred on the 84th percentile of the fit. The blue point denotes ZTF J1405+1039 which has a best-fit secondary temperature that is much hotter than expected for its mass.

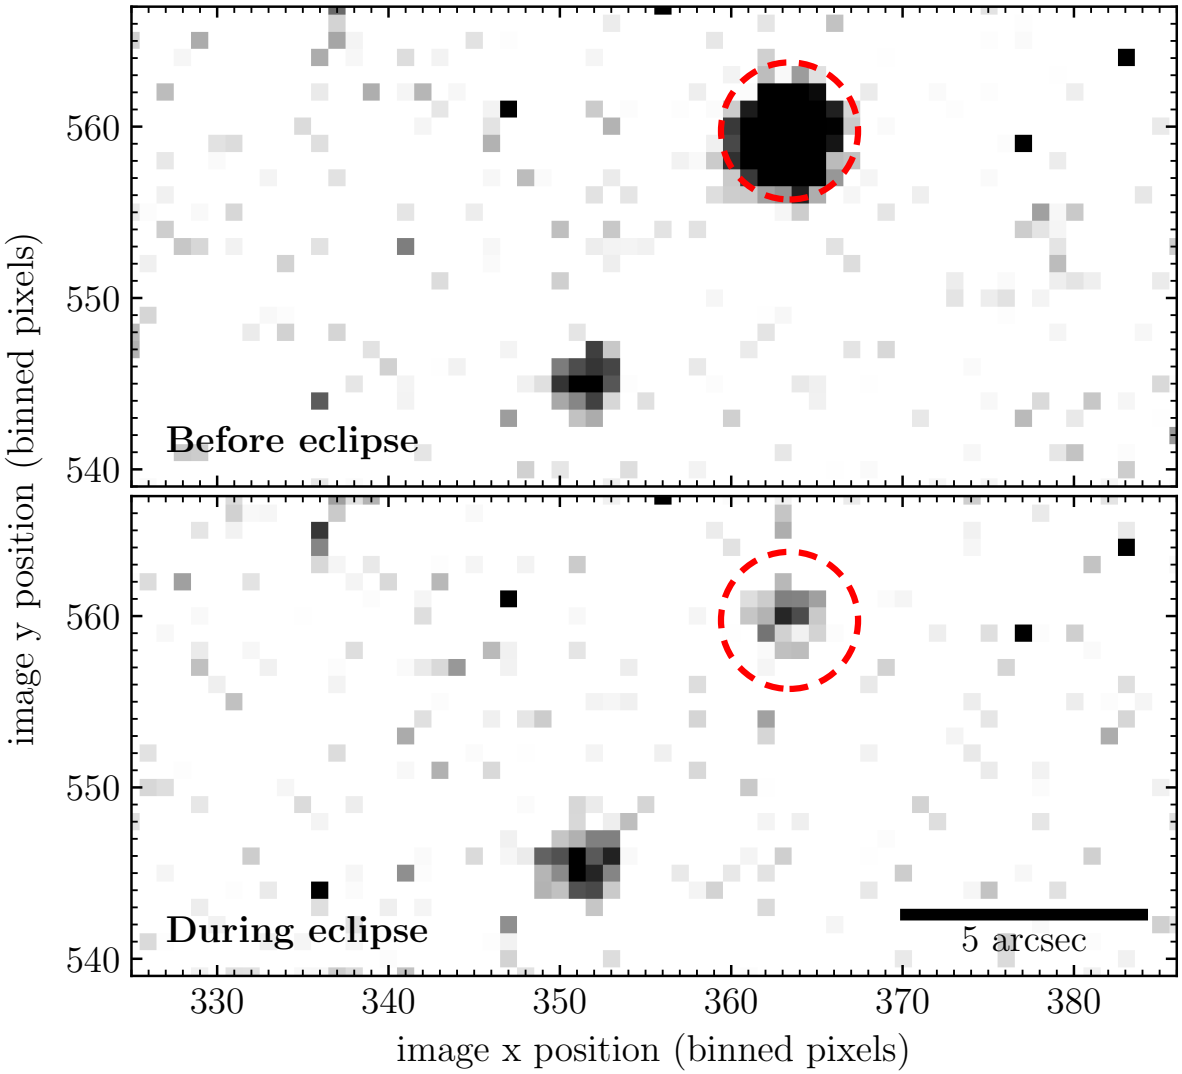

Figure 4.6: Stacked images of ZTF J1405+1039 taken with ULTRACAM in the i_s filter before and during the eclipse. The red dashed aperture shows the location of ZTF J1405+1039. It is clear that the source is still detected in-eclipse.

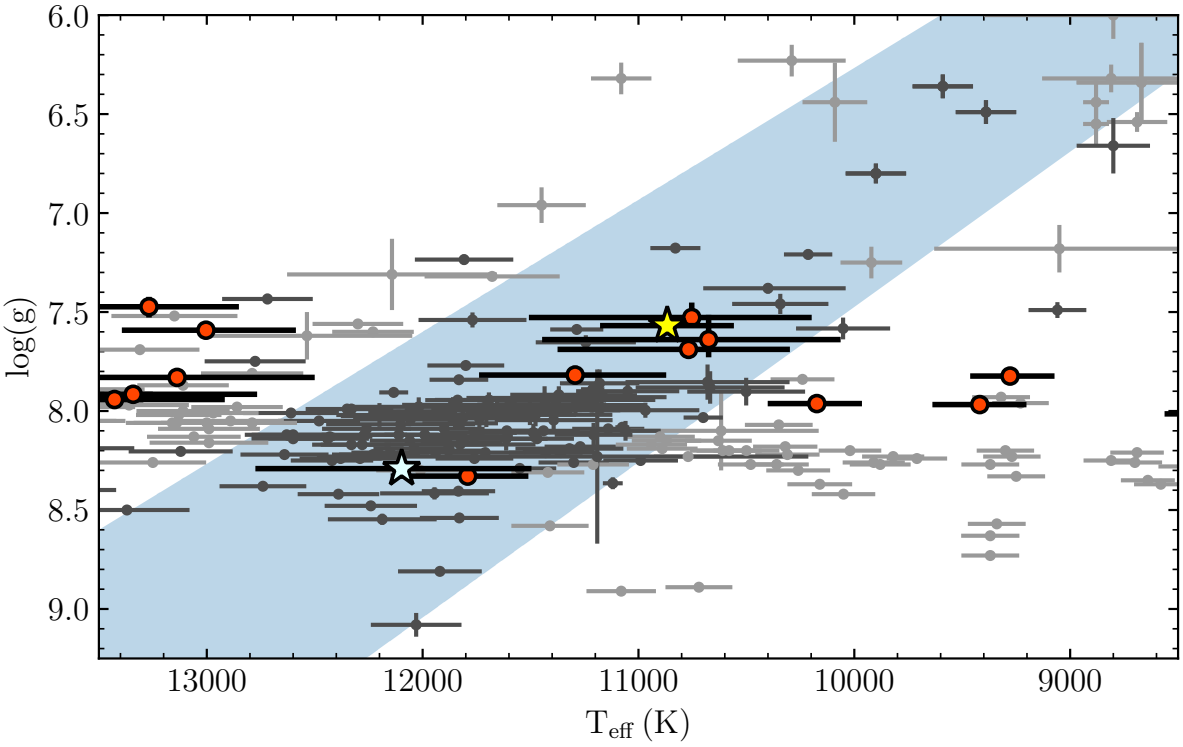

Figure 4.7: The ZZ Ceti instability strip (blue region) with known pulsating (dark grey) and non-pulsating (light grey) WDs from Gianninas et al. (2011); Steinfadt et al. (2012); Hermes et al. (2012, 2013a,c,b); Romero et al. (2022b). Points in red show the measured parameters of the WD components of binaries fit in this work, with the confirmed pulsators, ZTF J1407+2115 and ZTF J0528+2156, shown by the yellow and cyan stars respectively.

with the latter possessing a brown dwarf secondary, however neither of these systems are known to be eclipsing. Although these are important finds, Hermes et al. (2015) noted that there were a lot of unconstrained parameters in their fit to SDSS J1136+0409, limiting the precision of their asteroseismological analysis. Eclipsing examples of such systems would provide tight constraints on the stellar and binary parameters, enabling much more precise asteroseismological analyses.

Comparing the best fit parameters for the WD components to the ZZ Ceti instability strip (Figure 4.7), I find that seven of the systems I have fit have WDs that lie within 1σ of the instability strip (shown in Table 4.4). Closer inspection of their light curves do not reveal any clear photometric variability indicative of pulsations in five of the systems,

however, the out-of-eclipse data for many of these typically lasts less than 30 minutes and so is not enough to rule out pulsations either. Note that the WD temperatures are not necessarily precise enough to say with certainty whether a particular WD lies within the instability strip or not. Of these seven systems with WDs that lie in the instability strip, I have found two that show clear variability due to pulsations. These represent the first two ZZ Ceti WDs found in eclipsing WD + M PCEBs.

ZTF J1407+2115

ZTF J1407+2115 was first observed with ULTRACAM in February 2021. Unusual out-of eclipse variation was noticed but the data taken in this run were insufficient to confirm pulsations. We observed ZTF J1407+2115 again for 1 h on the 2nd of March 2022, detecting 3 clear pulsations and confirming it as the first eclipsing detached PCEB containing a ZZ Ceti WD. With this confirmation, we observed ZTF J1407+2115 in two long observing runs on the 4th and 26th of March 2022 using the $u_s\,g_s\,i_s$ and $u_s\,g_s\,r_s$ filters and lasting ~ 2 h and ~ 5 h respectively (Lomb-scargle periodograms of these two long runs are shown in Figure 4.10). It is the photometry from the observing run on the 4th of March that I use to fit the system parameters. I choose this observation primarily for consistency with the modelling performed on the other systems in this chapter, but also as the wider wavelength range provided by the i_s-band strengthens the constraints on the WD temperature. Additionally, chromospheric variability in the Hα feature can lead to higher scatter of M dwarf fluxes in the r_s-band.

In order to fit the eclipse photometry of this system, the pulsations need to be included in the light curve model to prevent them introducing large systematic errors in the best-fit parameters. I do this using a Gaussian process (GP) implemented through the PYTHON package, GEORGE[5] (Ambikasaran et al., 2015). The GP is applied to the residuals of the PYLCURVE model at each MCMC walker position, with the posterior

[5] https://george.readthedocs.io/en/latest/

log probability calculated as the sum of the GP marginalised log likelihood, the log likelihood from comparing the model WD SED with the measured eclipse depths, and the log priors (parallax and interstellar reddening).

I use the `ExpSquaredKernel`, defined by an amplitude, temperature, and scale-length – with the temperature scaling the pulsation amplitude between the light curves in different filters according to a blackbody law. These three GP parameters are included as free parameters in the fit. The GP is switched off between the second and third contact points where the WD is totally eclipsed by its M dwarf companion, with the contact points being calculated for every walker position.

EMCEE (Foreman-Mackey et al., 2013) is then used to sample from the posterior probability distribution and determine the best-fit parameters. This best-fit model is shown in Figure 4.8. I find the WD to have an effective temperature of $10\,900 \pm 300$ K and a mass of 0.41 ± 0.01 $M_\odot$, suggesting a core composed primarily of helium. This mass and temperature corresponds to a surface gravity of 7.57 ± 0.04 dex, placing it in a relatively sparsely sampled region in the middle of the instability strip (Figure 4.7).

Subtracting the best-fit eclipse light curve model from the g_s-band photometry of the longer run on the 26th of March leaves just the pulsation signal. Running a periodogram on this determines the main pulsation mode to have a frequency of 1.11 mHz (898 s) with an amplitude of around 47 parts per thousand (ppt) (Figure 4.10). The 3σ significance threshold for this is calculated to be 8 ppt following the method of Greiss et al. (2014), shuffling the flux values 10 000 times and taking the amplitude of the 99.7th percentile highest peak.

ZTF J0528+2156

ZTF J0528+2156, like ZTF J1407+2115, was first observed with ULTRACAM in February 2021. Attempts at fitting the eclipse light curve showed some possible structure in the residuals, prompting us to observe it again to search for pulsations. We observed

Target	RA	Dec	G	Pulsating
ZTF J0519+0925	05:19:02.1	+09:25:26.38	19.0	Candidate
ZTF J0528+2156	05:28:48.2	+21:56:28.94	17.7	Confirmed
ZTF J0948+2538	09:48:26.4	+25:38:10.68	18.7	Candidate
ZTF J1302−0032	13:02:28.3	-00:32:00.11	16.8	Candidate
ZTF J1407+2115	14:07:02.6	+21:15:59.75	17.4	Confirmed
ZTF J1634−2713	16:34:21.0	-27:13:21.54	18.8	Candidate
ZTF J1802−0054	18:02:56.4	-00:54:58.47	18.0	Candidate

Table 4.4: eclipsing PCEBs with – either confirmed or candidate – ZZ Ceti WDs

ZTF J0528+2156 again on the 18th of December 2022 for 1.8 h, detecting pulsations with a period of around 11 minutes and amplitude of around 5 per cent.

I fit the ULTRACAM photometry in the same way as for ZTF J1407+2115 – using a GP to model the pulsations. I find the WD to have an effective temperature of $11\,900 \pm 600$ K and a mass of 0.78 ± 0.02 $M_\odot$, corresponding to a surface gravity of 8.27 ± 0.044 dex and placing it comfortably within the instability strip (Figure 4.7). This best-fit model is shown in Figure 4.9. Computing the periodogram of the residuals of the eclipse light curve model in the same way as for ZTF J1407+2115, I find the main mode to have a frequency of 1.5 mHz (670 s) and amplitude of around 19 ppt with a 3σ significance threshold of 7 ppt (Figure 4.10).

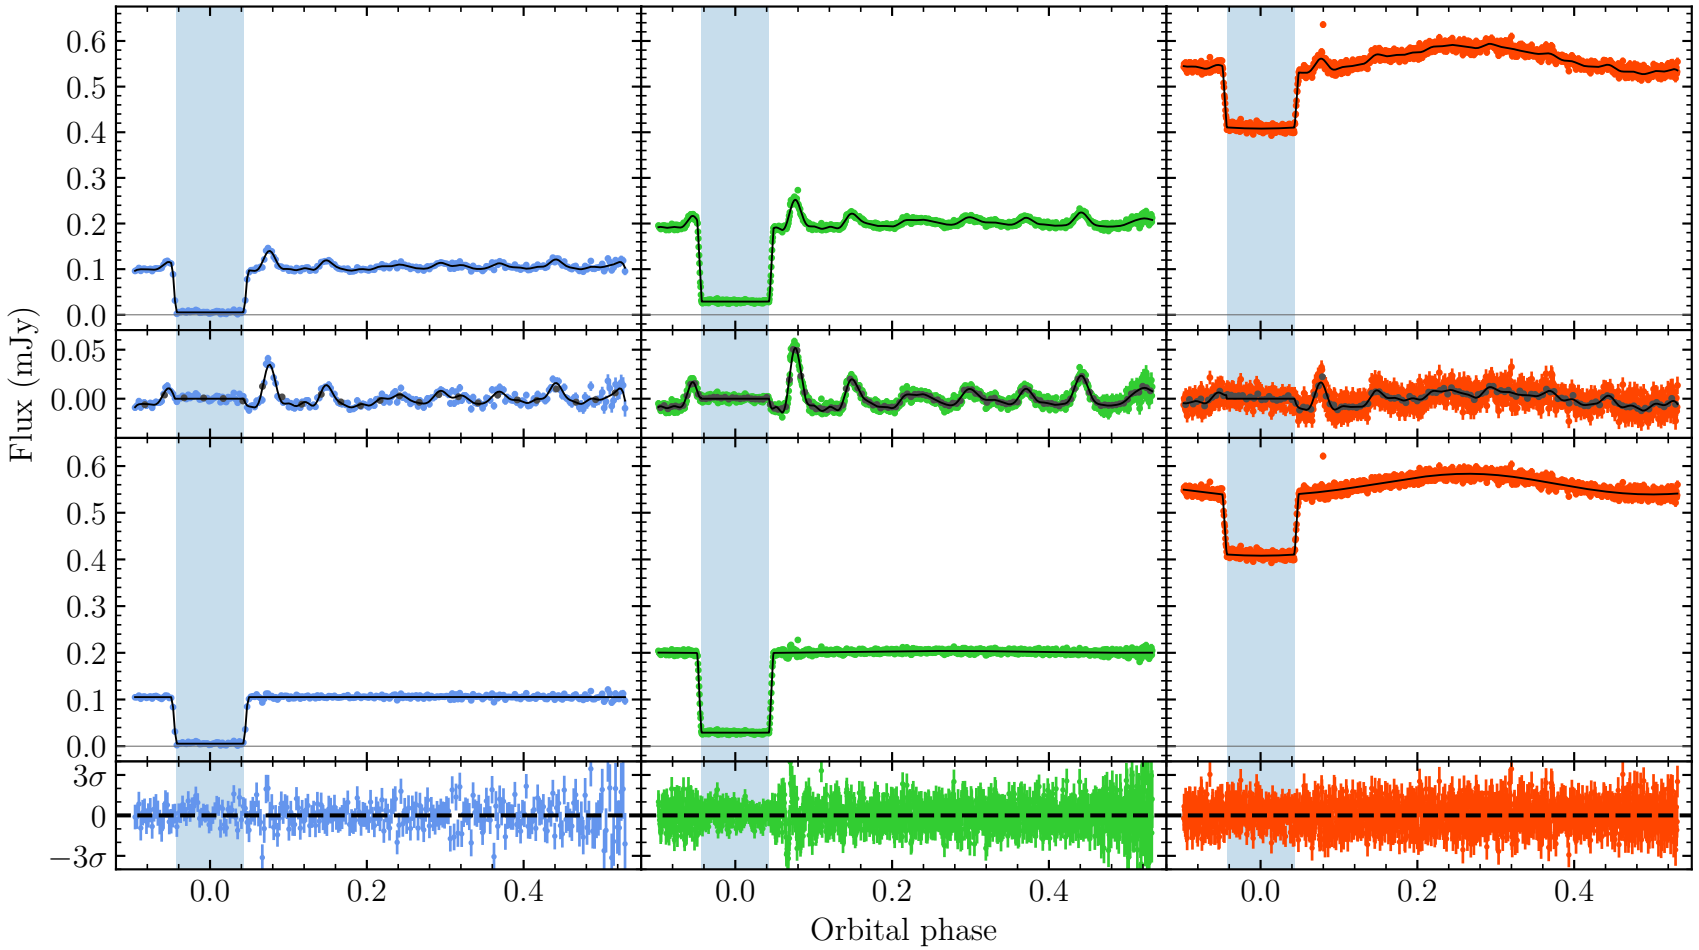

Figure 4.8: ULTRACAM $u_s\, g_s\, i_s$ light curves of ZTF J1407+2115. The top row of each plot shows the observed light curve (coloured points) with the combined eclipse plus mean GP pulsation model (black line). The second row shows the observed light curve with the eclipse model subtracted (coloured points) as well as the same data binned up by a factor of ten (dark grey points) with the mean GP model (black line). The third row shows the observed light curve with the mean GP subtracted with the black line showing the eclipse model. The bottom row shows the residuals of the full light curve model. The filled region shows the phase range where the GP is switched off (between the second and third eclipse contact points).

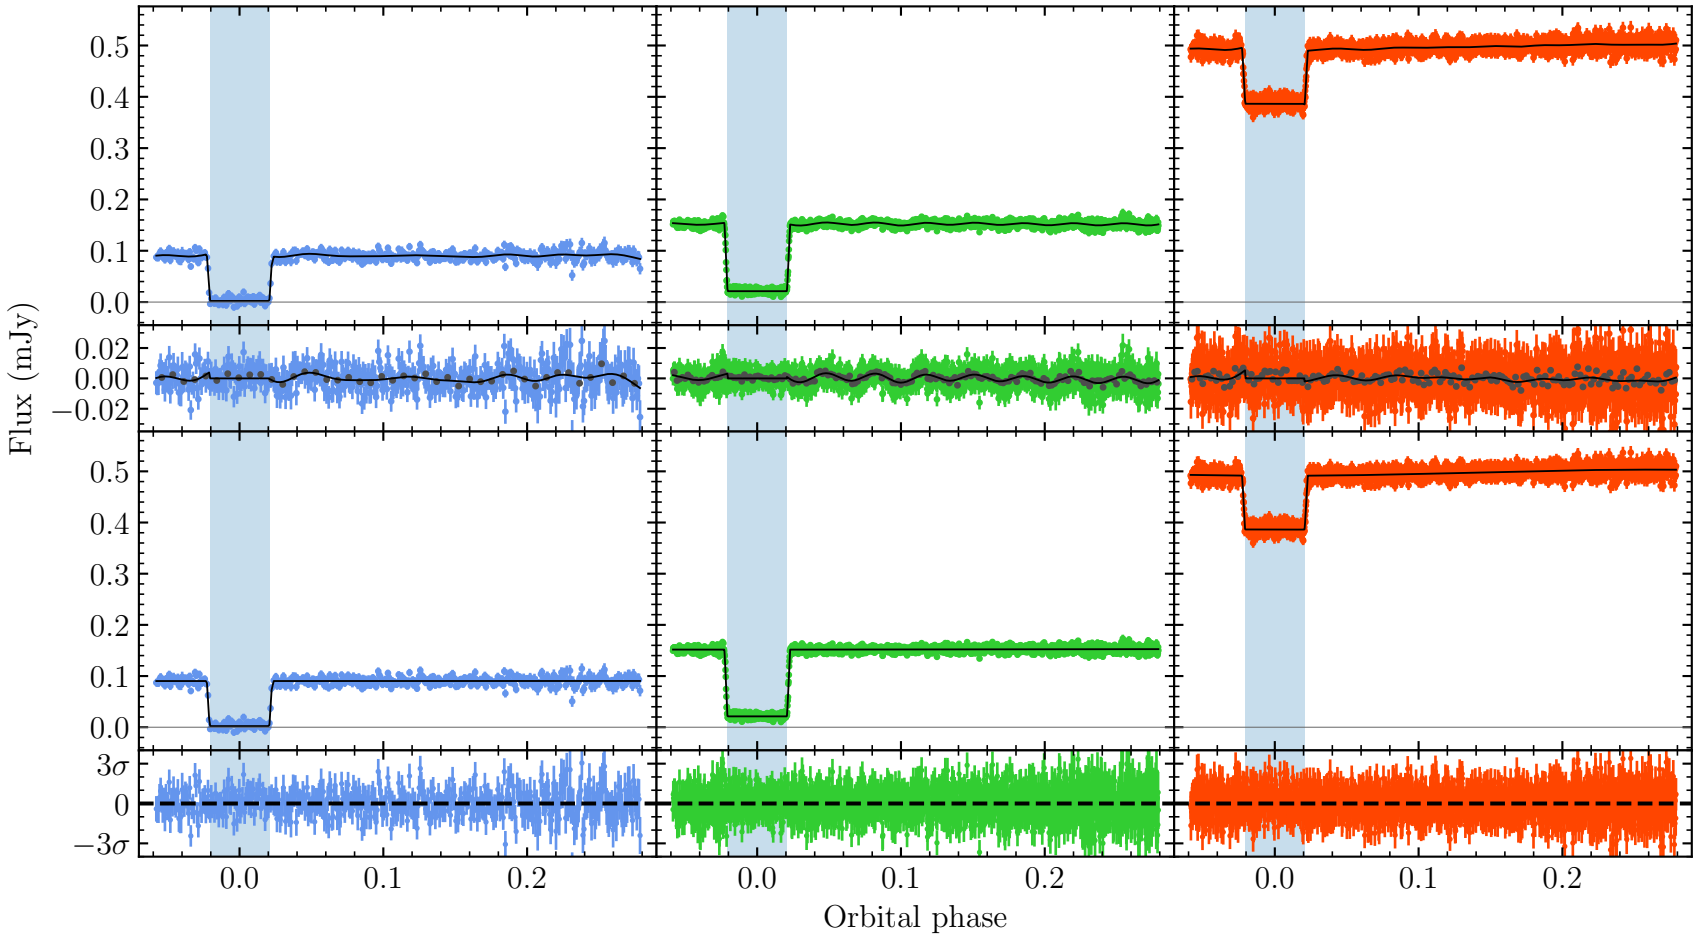

Figure 4.9: ULTRACAM $u_s\, g_s\, i_s$ light curves of ZTF J0528+2156. The top row of each plot shows the observed light curve (coloured points) with the combined eclipse plus mean GP pulsation model (black line). The second row shows the observed light curve with the eclipse model subtracted (coloured points) as well as the same data binned up by a factor of ten (dark grey points) with the mean GP model (black line). The third row shows the observed light curve with the mean GP subtracted with the black line showing the eclipse model. The bottom row shows the residuals of the full light curve model. The filled region shows the phase range where the GP is switched off (between the second and third eclipse contact points).

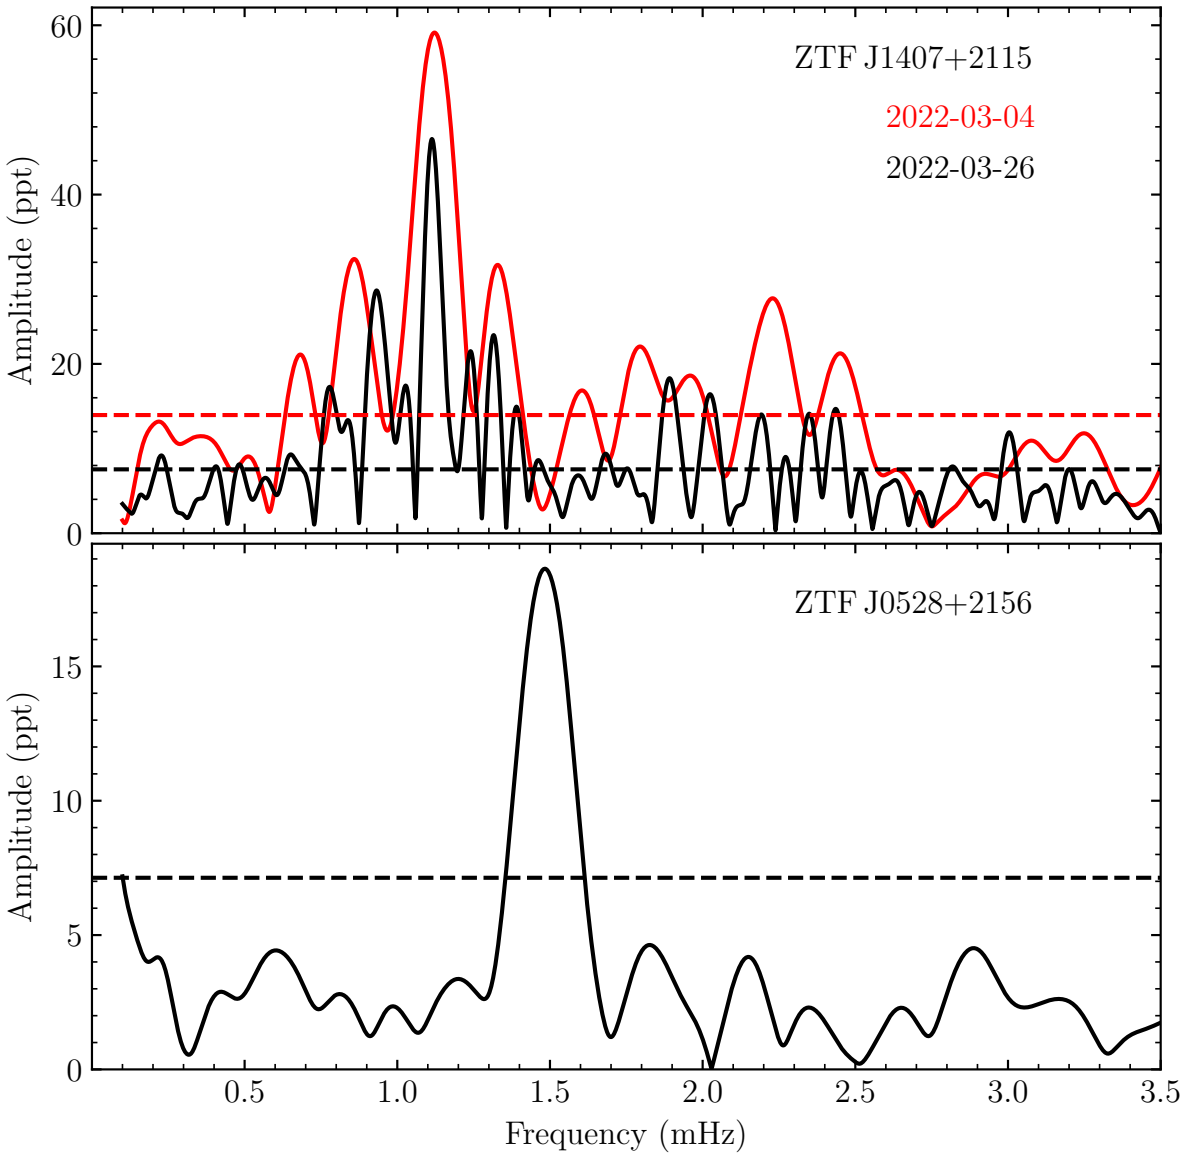

Figure 4.10: Lomb-Scargle periodograms (shown in parts per thousand relative to the flux of the WD) of the ULTRACAM g_s light curves of ZTF J1407+2115 and ZTF J0528+2156 with their respective eclipse light curve models subtracted. Horizontal dashed lines show the 3σ significance levels calculated using the bootstrapping method described by Greiss et al. (2014, Section 4.1).

4.5.4 Magnetic WDs

Around 36 per cent of WDs in CVs are observed to be strongly magnetic (Pala et al., 2020). This is in stark contrast with their progenitor population – the detached PCEBs –

of which only a handful possess WDs with strong magnetic fields. Schreiber et al. (2021) have proposed an evolutionary channel (shown in Figure 4.11) between the magnetic CVs and the detached magnetic population to explain this discrepancy. This relies on a rotation-driven dynamo in which a crystallising WD, spun up due to accretion during the CV phase, can generate the strong magnetic fields that we observe in CVs. Interactions between the newly-formed magnetic field of the WD and the magnetic field of the M dwarf then act to detach the binary, halting mass transfer and causing the binary to appear as a strongly magnetic detached PCEB for a period of time before angular momentum loss due to magnetic braking and gravitational wave radiation brings the two stars back into a mass-transferring state as a polar or intermediate polar.

A test of this model was performed by Parsons et al. (2021), using spectroscopic observations of detached magnetic PCEBs to constrain their evolutionary history, attempting to assess whether or not they are consistent with having undergone a mass-transferring phase in the past. All systems studied were found to be consistent with a previous CV phase but spectroscopic observations alone were not powerful enough to draw strong conclusions. More powerful constraints can be made if such systems are found to be eclipsing, enabling more precise measurements to be made from the eclipse photometry and therefore a robust test of the model.

As part of the ULTRACAM follow-up program I have discovered 6 new eclipsing PCEBs (Table 4.5) that have been confirmed from high-speed photometry as possessing magnetic WDs – showing clear evidence of a bright magnetic pole in the eclipse ingress/egress, with one previously known as a magnetic system but not known to be eclipsing (Figure 4.12). I have additionally found 3 candidate systems that show out-of-eclipse variation that disappears when the WD is eclipsed but for which the ingress/egress of the eclipse do not confirm a bright magnetic pole (Figure 4.13). These systems have been found by searching for unusual out-of-eclipse variation in their ZTF light curves (Figure 4.12), inconsistent with the ellipsoidal modulation or reflection effect that is com-

mon in PCEBs. This unusual out-of-eclipse variability was noted in the pre-intermediate polar, SDSS J0303+0054, (Parsons et al., 2013a) and is due to additional emission in the form of cyclotron radiation from the magnetic poles of the WD. The effect of the cyclotron emission on the eclipse profiles – introducing steps in the ingress and egress due to the eclipse of the small, bright magnetic pole (Figure 4.12) – makes the light curves of the magnetic systems much more complicated to fit and so the analysis of these systems will be left for future work.

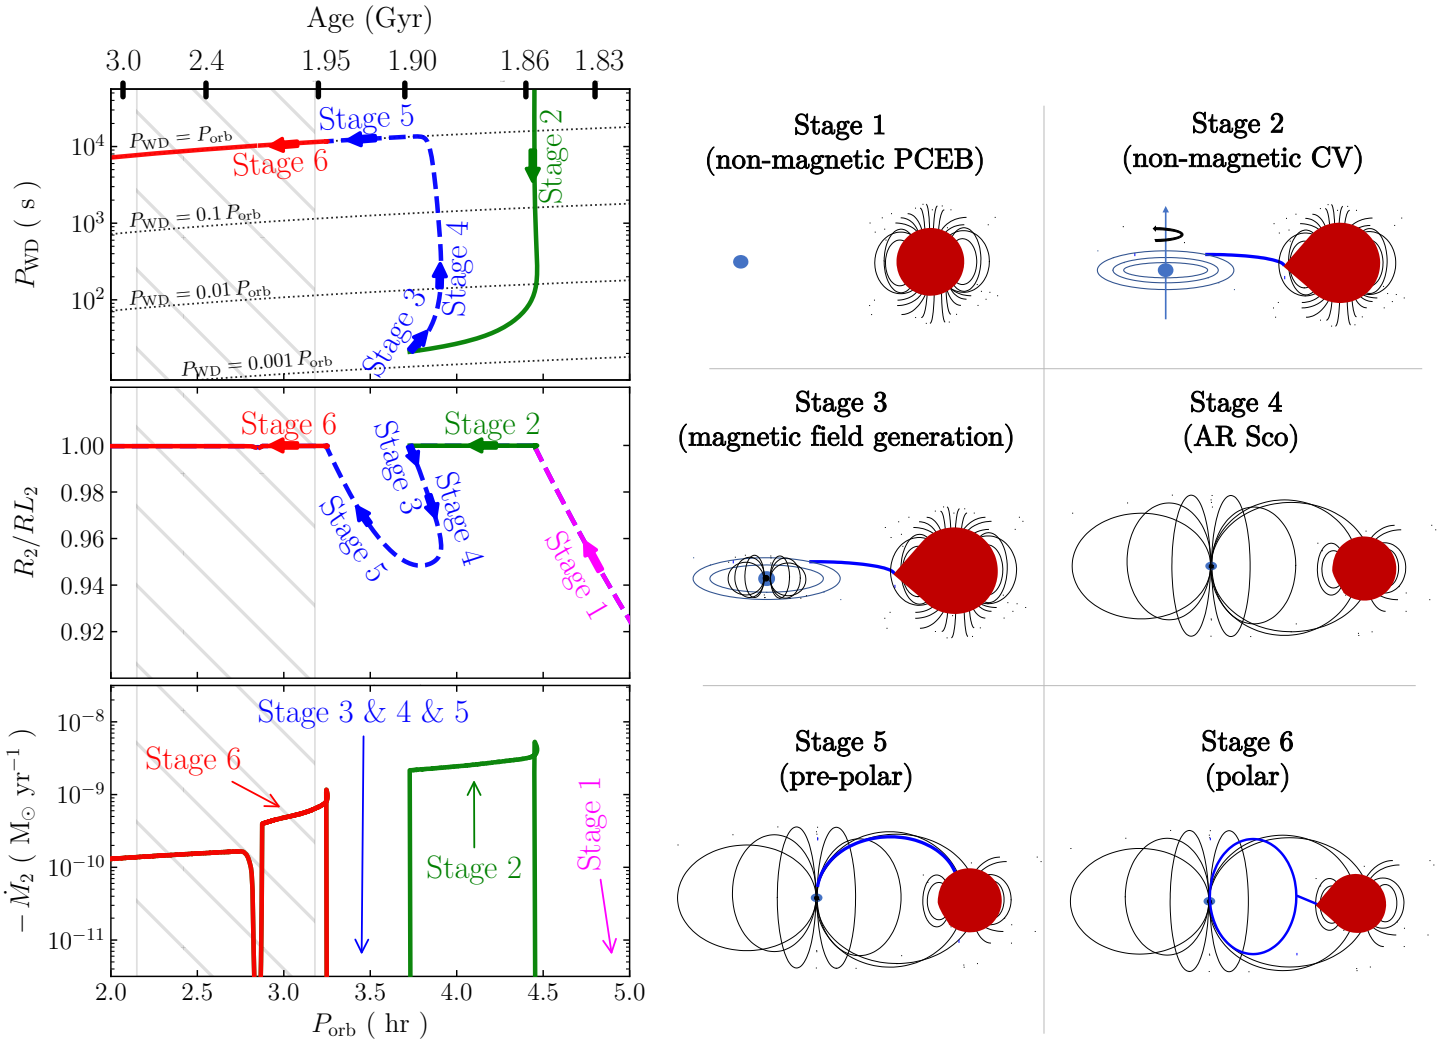

Figure 4.11: *Left*: WD spin period, Roche lobe filling factor, and mass-transfer rate as a function of orbital period. *Right*: Schematic diagram of the proposed evolutionary phases that lead to the formation of a PCEB containing a magnetic WD. Figure from Schreiber et al. (2021).

Target	RA	Dec	G	Magnetic
ZTF J0126+1210	01:26:07.8	+12:10:49.14	18.8	Candidate
ZTF J0220+6303	02:20:04.6	+63:03:59.63	17.4	Candidate
ZTF J0406+0958	04:06:27.2	+09:58:26.97	18.1	Candidate
ZTF J0618−0919	06:18:09.9	−09:19:04.28	16.5	Confirmed
ZTF J1206+5100	12:06:15.7	+51:00:46.77	18.6	Confirmed
ZTF J1922+1038	19:22:15.3	+10:38:38.13	16.0	Confirmed
ZTF J2142+4309	21:42:32.0	+43:09:28.97	19.3	Confirmed
ZTF J2220+0721	22:20:07.5	+07:21:29.74	18.3	Confirmed
ZTF J2353+4153	23:53:55.0	+41:53:04.40	18.7	Confirmed

Table 4.5: eclipsing PCEBs with – either confirmed or candidate – magnetic WD components.

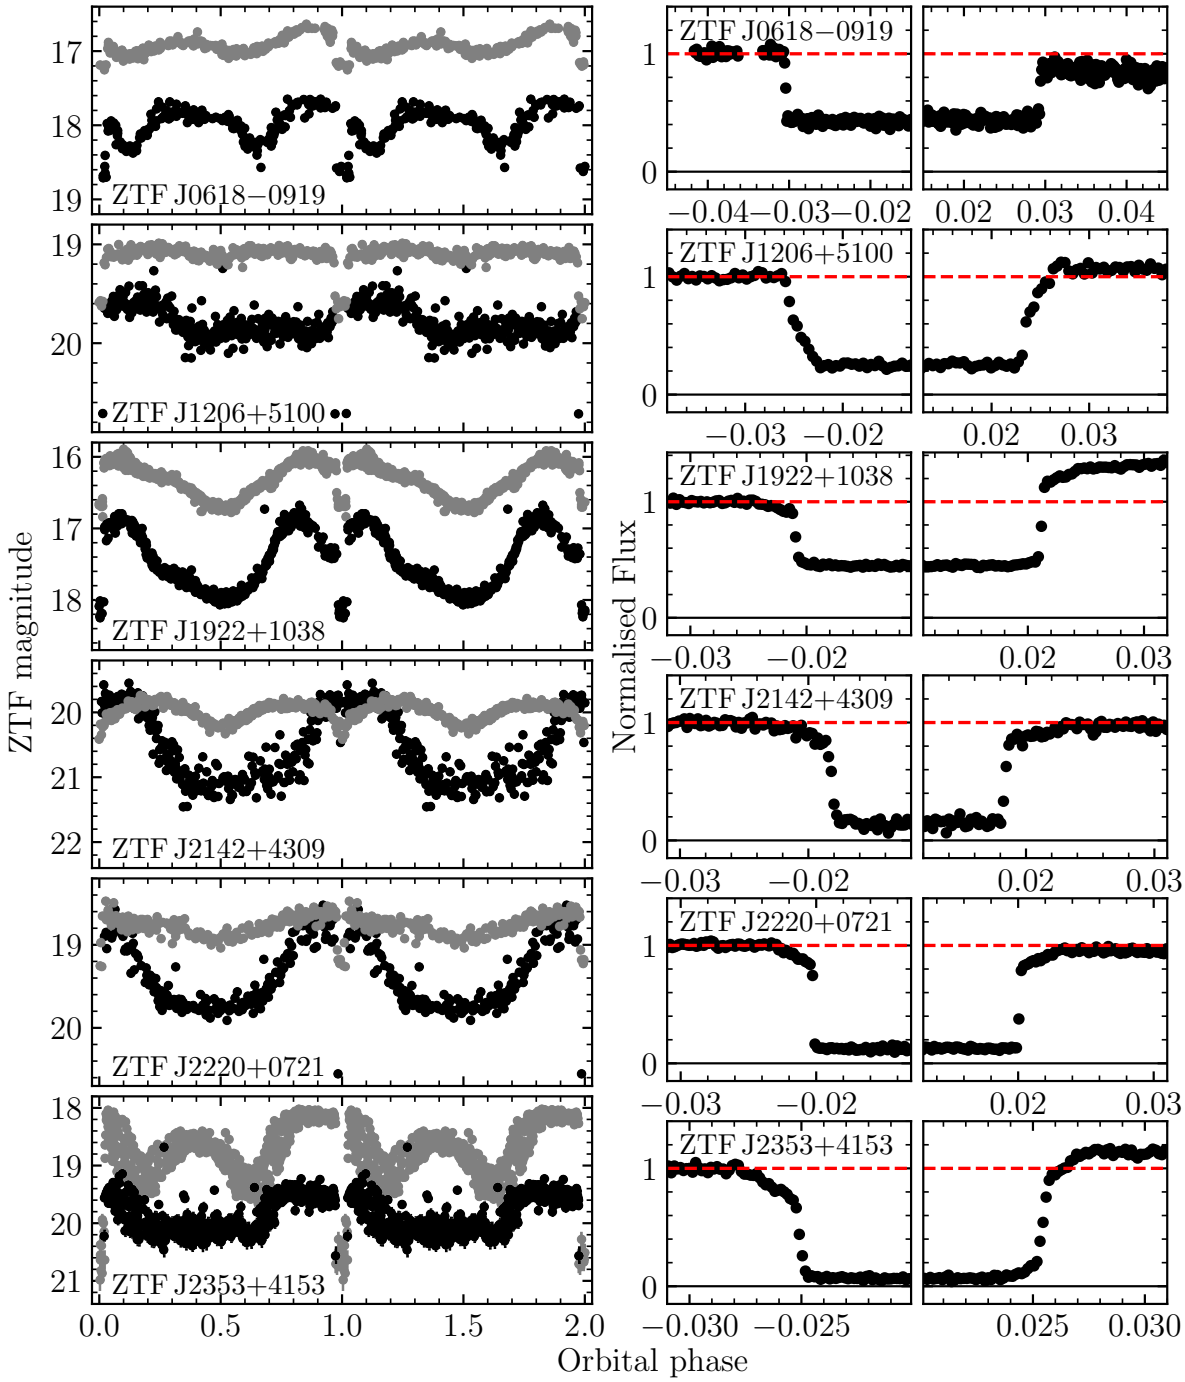

Figure 4.12: *Left*: ZTF g-band (black) and r-band (grey) light curves of the 6 new confirmed eclipsing PCEBs with magnetic WDs. All show out-of-eclipse variation inconsistent with reflection effect or ellipsoidal modulation in at least one filter. Some light curves have been binned for clarity. *Right*: Normalised ULTRACAM/HiPERCAM g_s-band primary eclipse light curves (zoomed in on the ingress and egress) of the 6 new confirmed eclipsing PCEBs with magnetic WDs. The solid grey line shows a flux of zero while the red dashed line shows the mean flux of the first 10 points shown.

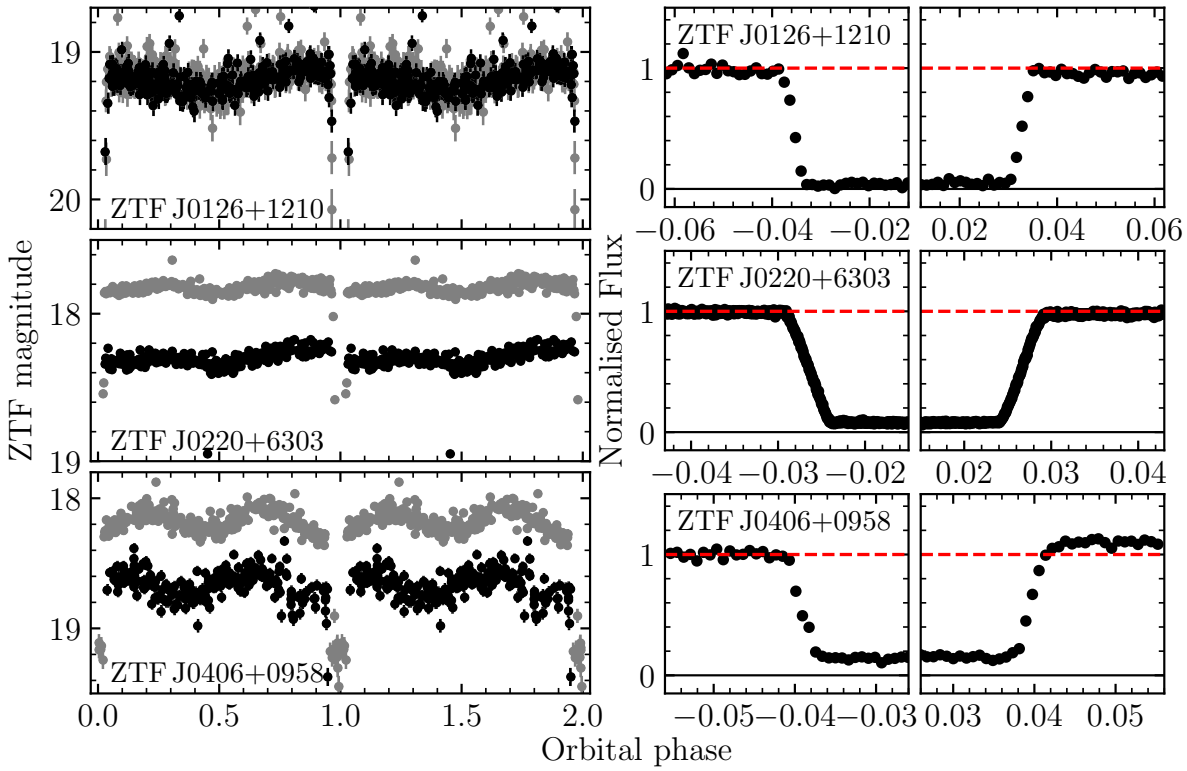

Figure 4.13: *Left*: ZTF g-band (black) and r-band (grey) light curves of the 3 new candidate eclipsing PCEBs with magnetic WDs. All show out-of-eclipse variation inconsistent with reflection effect or ellipsoidal modulation in at least one filter. Some light curves have been binned for clarity. *Right*: Normalised ULTRACAM/HiPERCAM g_s-band primary eclipse light curves (zoomed in on the ingress and egress) of the 3 new candidate eclipsing PCEBs with magnetic WDs. The solid grey line shows a flux of zero while the red dashed line shows the mean flux of the first 10 points shown.

4.6 Summary

Through a dedicated program of high-speed photometric follow-up I have obtained multi-band eclipse light curves for 42 new PCEBs found using ZTF. I have characterised 33 of these systems from their eclipse light curves alone – finding three that contain substellar companions, almost doubling the number of eclipsing examples known, and two with pulsating WDs representing the first ZZ Ceti WDs known in eclipsing WD + M binaries. Of the remaining nine systems, I have found six to contain strongly magnetic WDs from their eclipse photometry with three further candidates. These will be invaluable to the study of magnetic field generation in binary WDs. These results demonstrate that a photometric approach to the follow-up of eclipsing systems can effectively discern interesting sub-types of PCEBs, including those that would be otherwise missed by spectroscopic follow-up.

Chapter 5

Future work and Summary

5.1 Future work

Having demonstrated in this thesis that high-speed eclipse photometry alone can successfully be used to characterise WD + M binaries as well as to find particularly interesting systems, there are some avenues of research that naturally follow on from the work presented in the previous chapters.

5.1.1 A volume-limited survey of eclipsing WD + M binaries

With the success of the *Gaia* space mission, volume-limited samples of various astronomical objects are becoming commonplace. These allow for population studies that suffer from significantly fewer biases than previous works that relied on magnitude-limited samples or samples with hard-to-account-for selection effects.

A volume-limited survey of eclipsing WD + M binaries from ZTF would be complementary to that of the CV population (Pala et al., 2020) and those of single field WDs (Hollands et al., 2018; Bagnulo & Landstreet, 2021) which have made significant contributions to the field. A reliable measurement of the WD mass distribution in PCEBs will be one of the key results of this. When compared with that of the CV population this places constraints on the enhanced angular momentum losses which have been invoked in order explain the higher than expected masses of WDs in CVs (Schreiber et al., 2016), as well as explaining their orbital period distribution. Previous measurements of the PCEB WD mass distribution have used either low-resolution spectroscopy or SED fitting (Rebassa-Mansergas et al., 2021) to determine stellar parameters, with both methods suffering from a lack of precision and accuracy in the measurement technique as well as difficult-to-determine observational biases. Again, this issue is solved by using an eclipsing sample together with the fitting technique I have developed and used in this thesis.

This sample of PCEBs will also enable a comparison with the population of wide

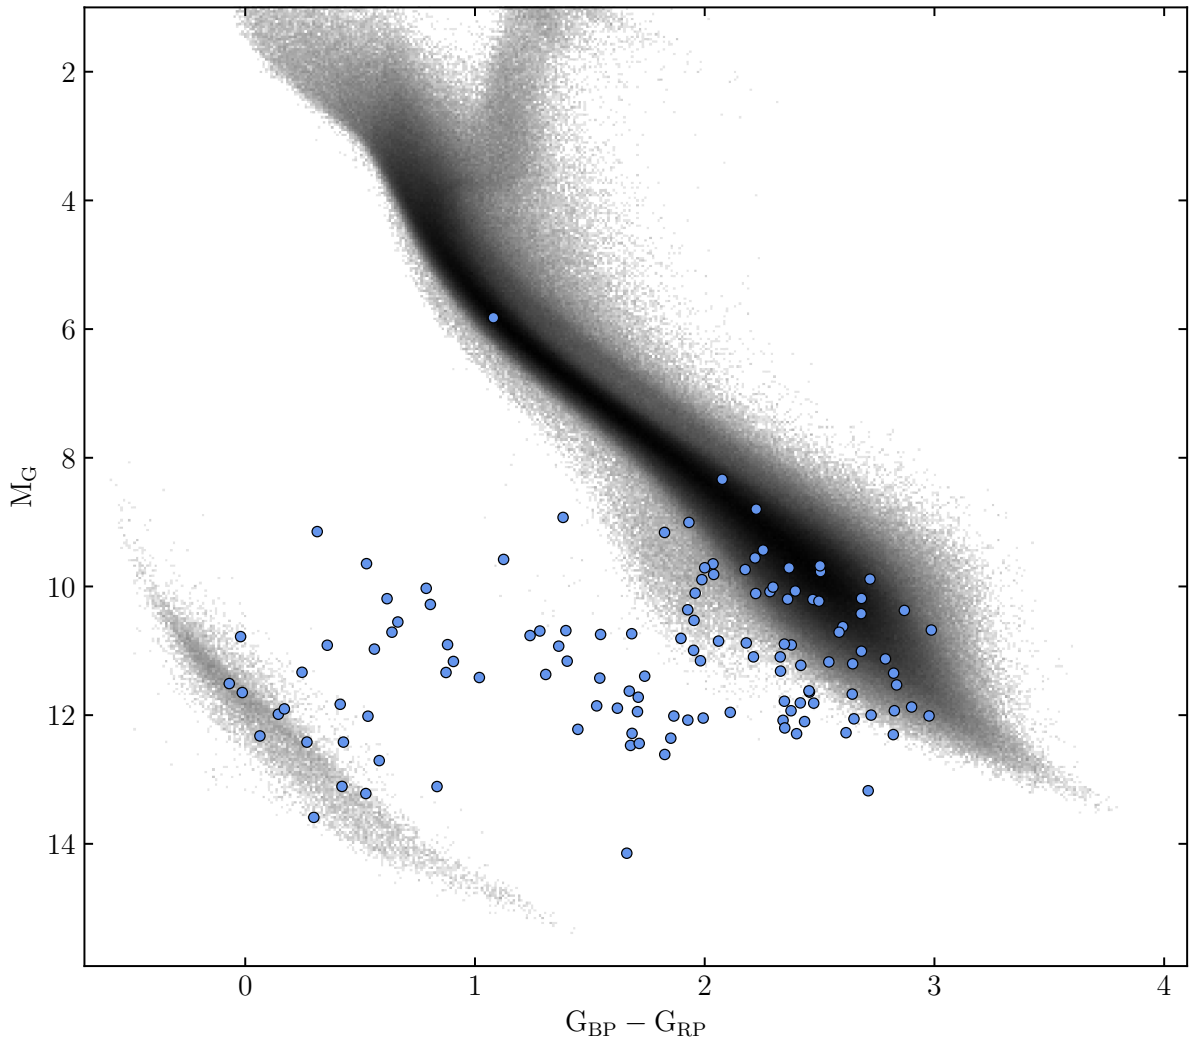

Figure 5.1: Gaia colour-magnitude diagram of the ZTF 200 parsec sample of eclipsing WD + M systems (blue points). A random selection of sources within 500 parsecs is shown in black.

WD + M systems, placing constraints on the dominant angular momentum loss mechanism in close binaries, the key to understanding the evolution of CVs. The use of eclipse photometry will also allow the magnetic fraction of WDs in PCEBs to be measured more accurately than ever before, with the effects of magnetism (namely strong cyclotron emission in the optical and infrared) causing unusual features and discontinuities in the light curves (Figure 4.12), making them relatively easy to detect. Although a volume-limited sample of WD + M binaries will still, inevitably, suffer from some observational biases – most notably against systems with particularly shallow eclipses and long orbital periods – these biases will be relatively simple to account for, as discussed in van Roestel (in preparation).

5.1.2 Analysis of magnetic WDs in detached PCEBs

As mentioned in Chapter 4, magnetic WDs found in eclipsing WD + M binaries offer perhaps the best test of theorised mechanisms of magnetic field generation in WDs. At the end of Chapter 4, I had discovered six of these systems and found a further three candidates.

Since carrying out the work for Chapter 4, I have obtained phase-resolved spectroscopy for many of these systems with X-Shooter (Vernet et al., 2011) on the VLT as well as some spectroscopy from LRIS on the Keck I telescope. This spectroscopic follow-up has confirmed ZTF J0126+1210 as a magnetic PCEB from the presence of Zeeman split Hβ lines (Figure 5.2), meaning I have now discovered 7 confirmed magnetic PCEBs that are eclipsing.

Beyond confirming systems as harbouring magnetic WDs, these spectroscopic data will enable measurements of the magnetic field strength of the WDs from the wavelength,

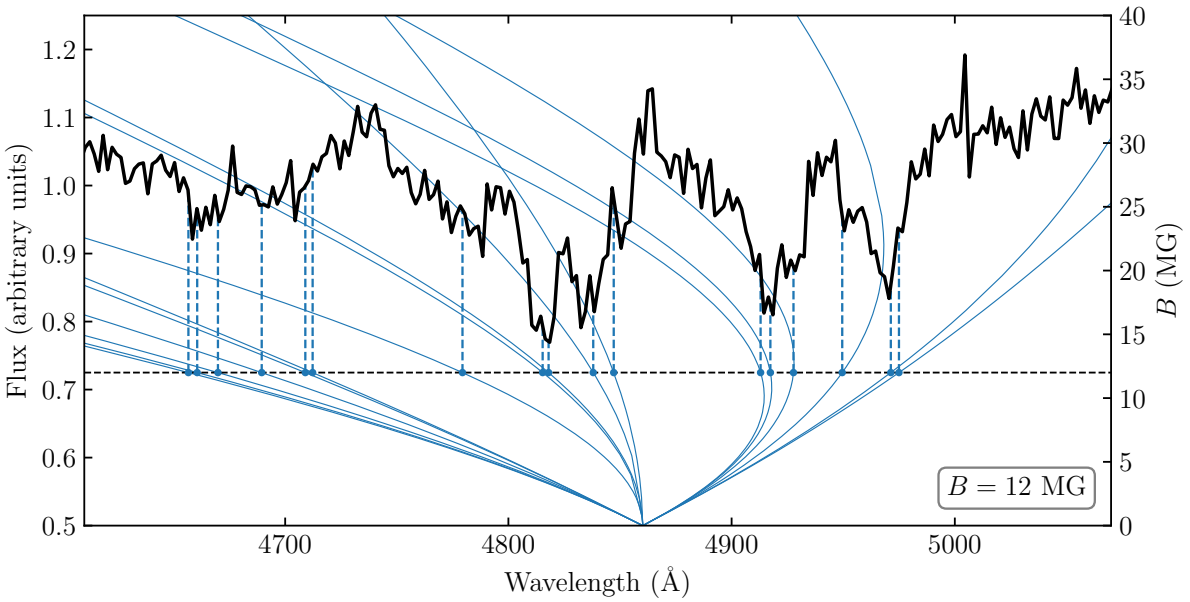

Figure 5.2: Co-added and binned X-Shooter UVB arm spectra of ZTF J0126+1210 around the Hβ absorption line. There is clear Zeeman splitting of the line as a result of its magnetic field. The locations of the Zeeman split components as a function of the magnetic field strength, B, are shown by the solid blue lines, with the splitting for a 12 MG field strength shown by the vertical dashed lines.

λ_n, and separation of the cyclotron lines:

$$\lambda_n = \left(\frac{10700}{n}\right)\left(\frac{10^8}{B}\right) \text{Å}, \tag{5.1}$$

where n is the cyclotron harmonic and B is the magnetic field strength in Gauss (Wickramasinghe & Ferrario, 2000). At what orbital phase these cyclotron lines are strongest also provides a constraint on the longitude of the magnetic pole on the surface of the WD (Figure 5.3). As well as helping characterise some of the parameters of the WD magnetic field, this spectroscopic data will allow for a precise measurement of the radial velocity amplitude of the M dwarf secondary, K_2, as well as the rotational broadening of the M dwarf:

$$V_{rot} \sin i = K_2(1+q)\frac{R_2}{a}. \tag{5.2}$$

Together, these provide information on the mass ratio of the binary and therefore help

constrain the masses of both components. Combining this spectroscopic data with the high-cadence multi-colour eclipse photometry that we have obtained with ULTRACAM and HiPERCAM (Figure 5.4) will enable a thorough analysis of strongly magnetic WDs in close binaries for the first time.

5.1.3 Asteroseismological analysis of ZZ Cetis in eclipsing PCEBs

As mentioned in Chapter 4, pulsations enable an analysis of the internal structure of the WD. In PCEBs, this internal structure will provide information on whether the common-envelope phase affects the remnant core of the giant star that becomes the WD, and if so, how. Since the writing of Chapter 4, we have discovered a third ZZ Ceti WD in an eclipsing PCEB. Phase-resolved spectroscopic follow-up (which we have now obtained for ZTF J1407+2115), together with the eclipse photometry already taken will allow precise model-independent constraints to be placed on the WDs in these systems. With long and continuous photometric observations of these three targets we should be able to detect and identify enough pulsation modes to perform high-quality asteroseismological analyses of WDs in PCEBs for the first time.

5.1.4 Core-composition of High-mass WDs

When a CO-core WD gains sufficient mass – either through a merger or by accretion from a binary companion – such that it approaches the Chandrasekhar limit it will "explode" as a type Ia supernova (Hillebrandt & Niemeyer, 2000). The requirement for a CO-core WD is particularly important though, with ONe-core WDs thought to collapse directly to neutron stars instead (Miyaji et al., 1980; Jones et al., 2013). Unfortunately, the mass at which WDs transition from CO-core compositions to ONe-core compositions is poorly constrained, making it difficult to know whether a particular binary (real or simulated) can be considered as a type Ia supernova progenitor. The precise measurements possible

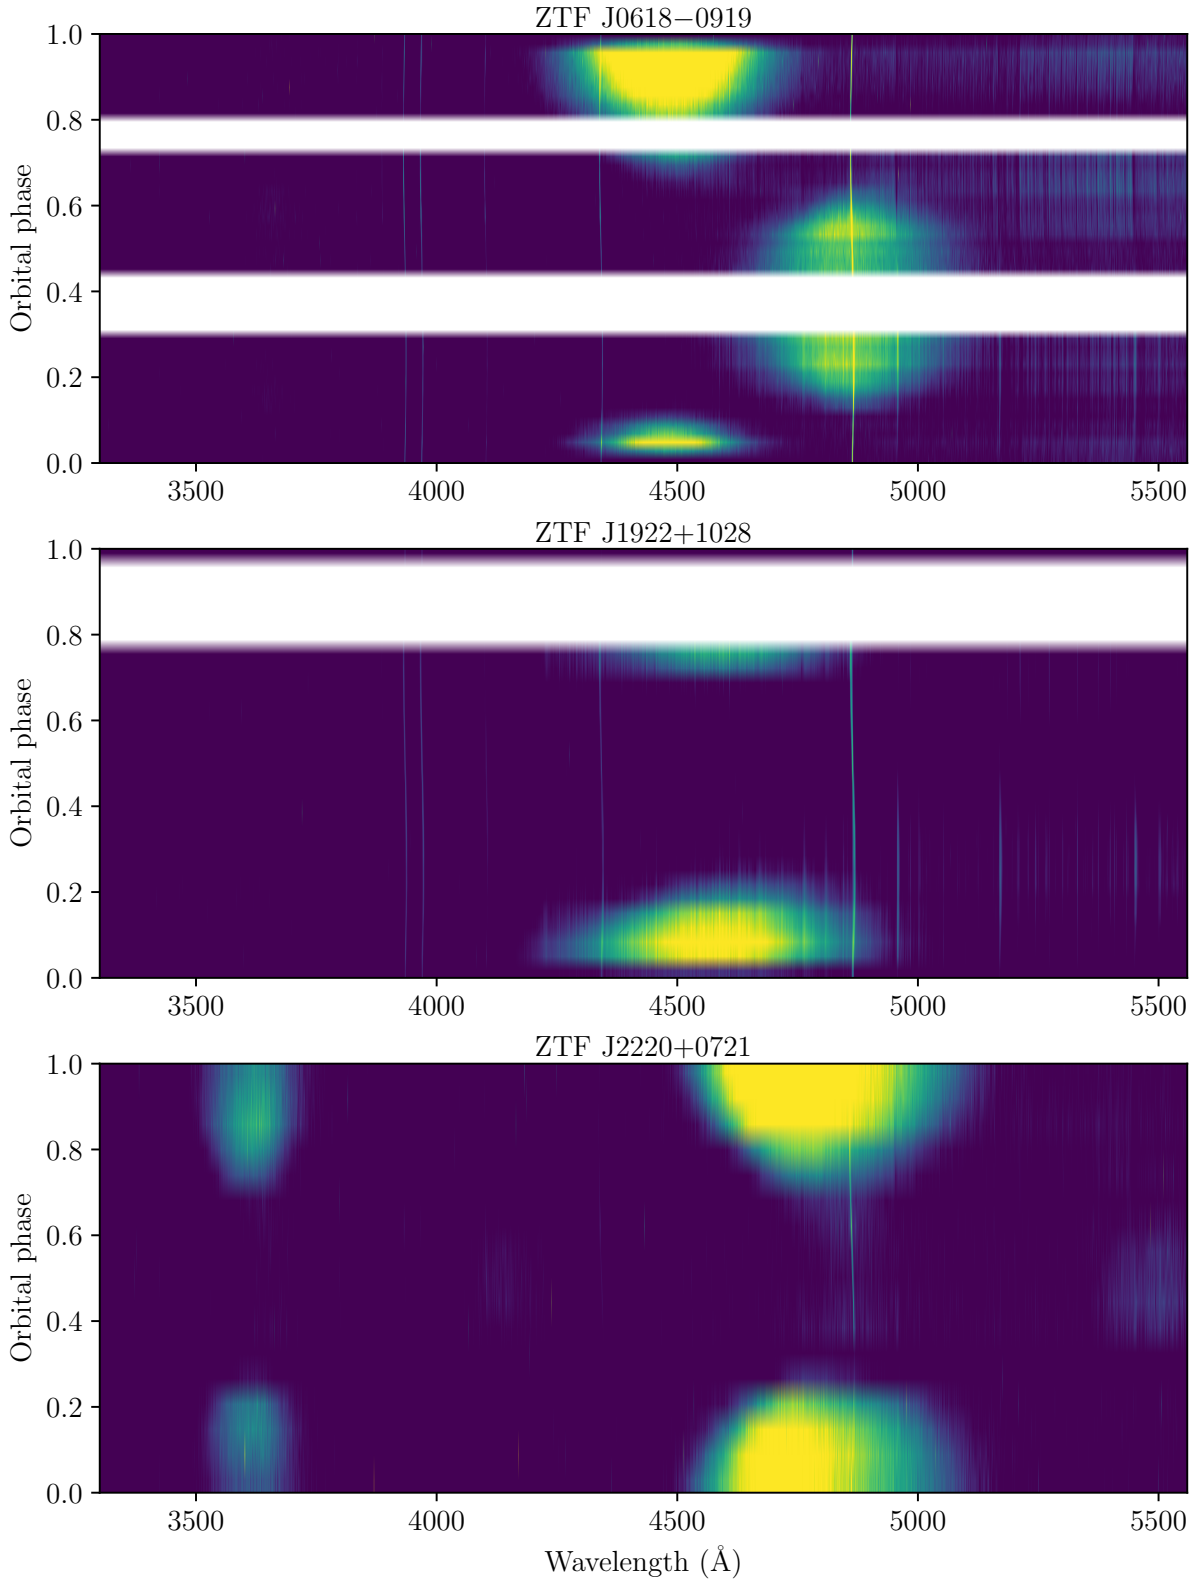

Figure 5.3: Trailed X-Shooter UVB arm spectra (with the M dwarf contribution subtracted) of three of the magnetic eclipsing WD + M systems, plotted as a function of orbital phase. White regions show where no data was taken. Strong cyclotron emission features can be seen in all targets as the broad yellow bands. These are restricted to the orbital phases where a magnetic pole is visible from our line-of-sight as the WD rotates.

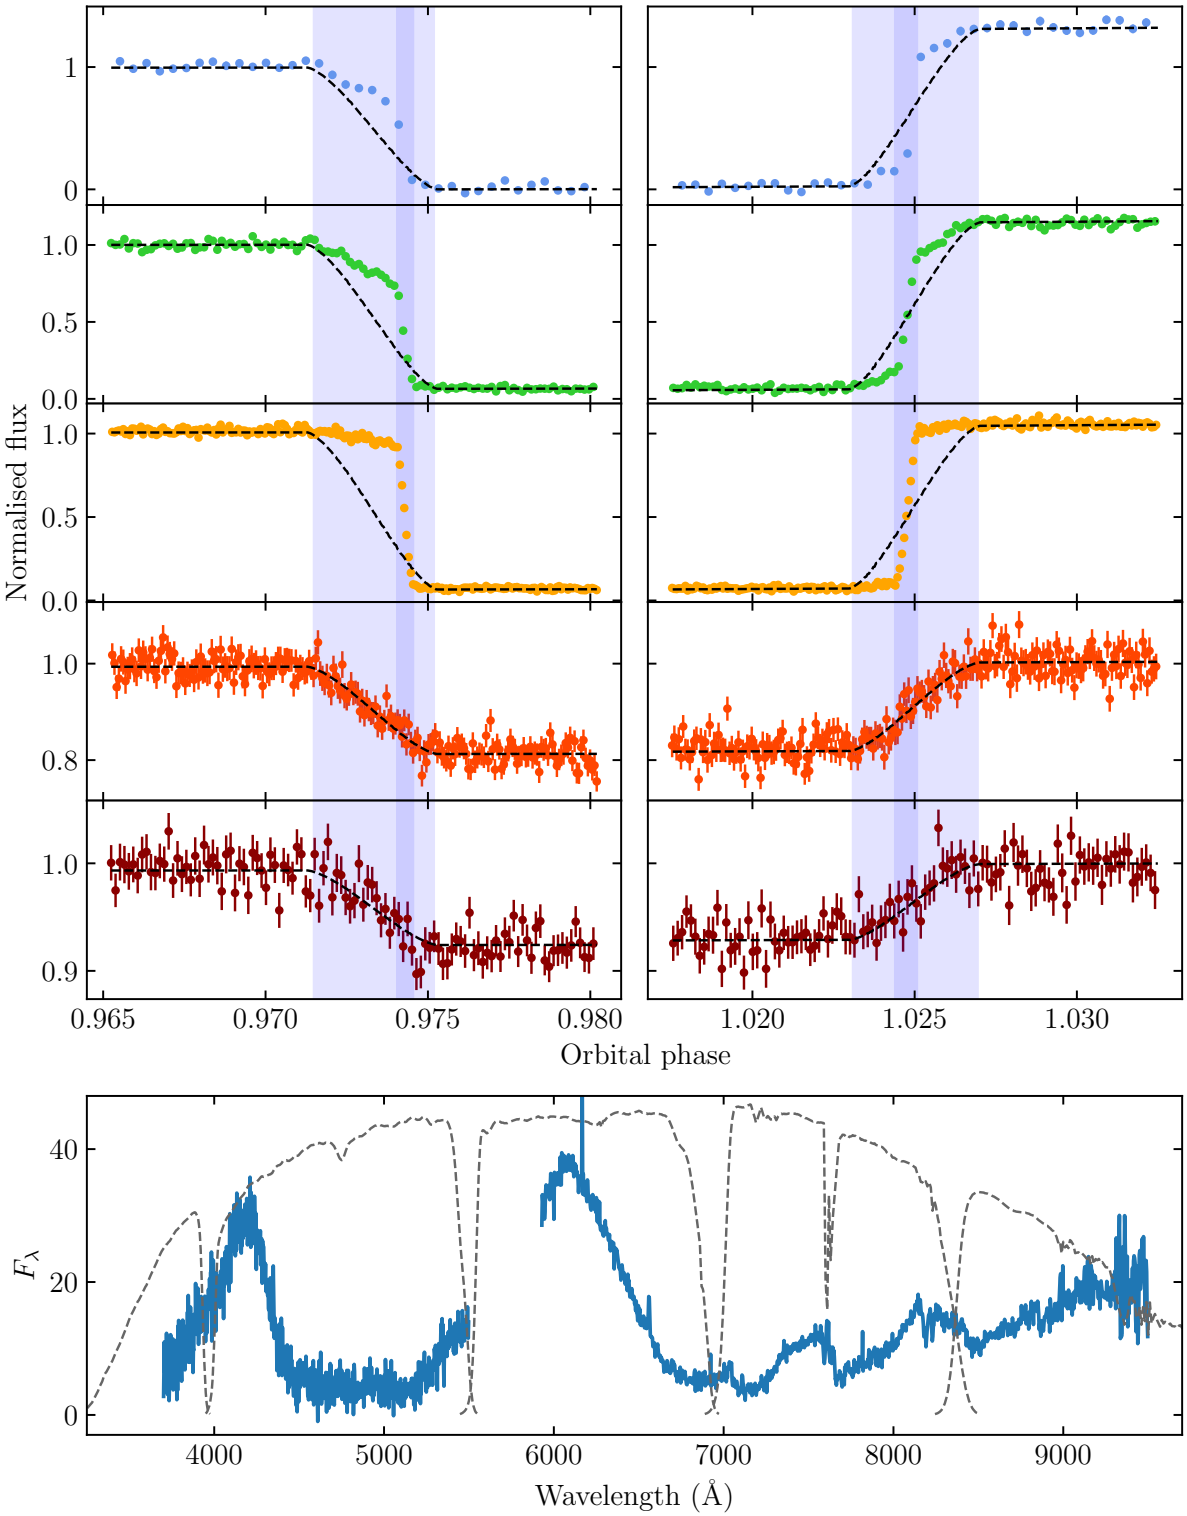

Figure 5.4: *Top*: HiPERCAM $u_s g_s r_s i_s z_s$ eclipse photometry of ZTF J2353+4153 zoomed in around the ingress and egress of the WD. The ingress/egress of the WD photosphere is marked by the light blue region while that of the small magnetic pole is marked by the darker blue region. An eclipse model (fit to the i_s-band photometry) is shown by the dashed line. This model is scaled to the in-eclipse and out-of-eclipse fluxes in the other filters for comparison. *Bottom*: Hale-DBSP spectrum of ZTF J2353+4153 (blue) with the HiPERCAM Super-SDSS bandpasses (grey dashed).

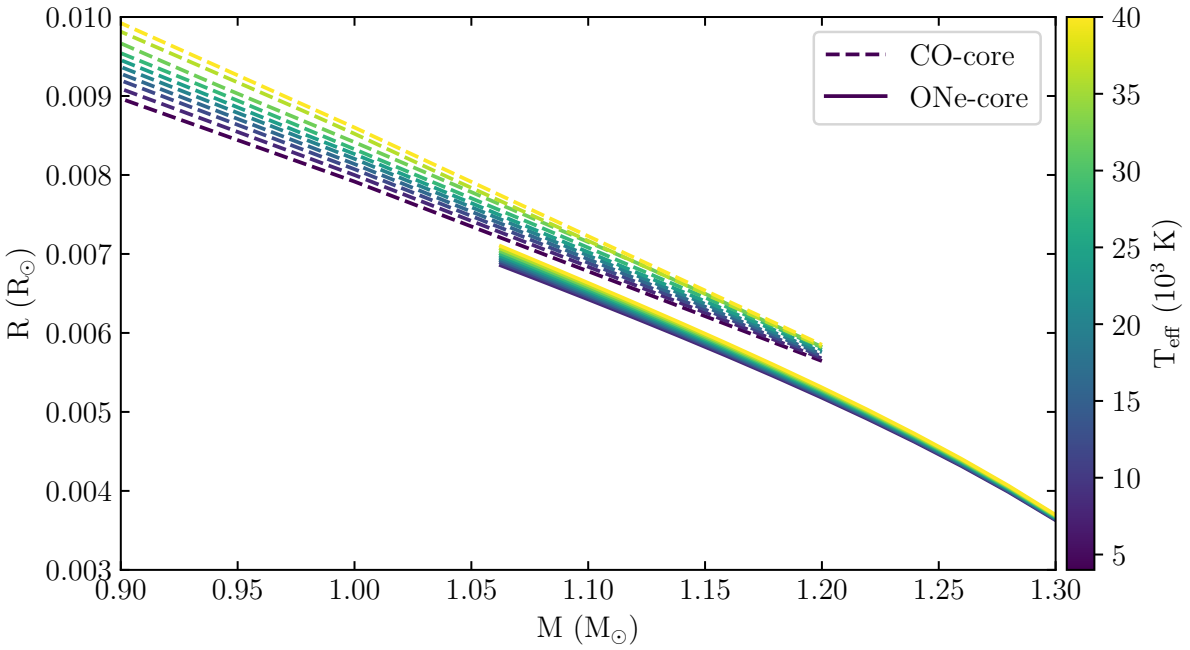

Figure 5.5: Mass-radius relation for WDs with CO- and ONe-core compositions (Bédard et al., 2020; Blouin et al., 2018; Tremblay et al., 2011; Althaus et al., 2005).

in eclipsing binaries may provide a potential method of measuring the mass at which this transition occurs. Theoretical models suggest that, for a given mass, CO-core and ONe-core WDs should have different radii (Figure 5.5). This radius discrepancy is large enough that, with high quality observations, it should be possible to discern a point, above which, observed WDs are more consistent with ONe models than with CO models (for example, this method has previously been used to measure the mass at which WDs transition from He-core models to CO-core models; Parsons et al. 2017a). A dedicated search for high-mass ($M \gtrsim 0.8$ M$_\odot$) WDs in eclipsing PCEBs would be the first step towards this goal, making use of high-speed photometry to follow-up eclipsing PCEBs discovered in ZTF by van Roestel (in preparation) that have been initially characterised as harbouring WDs with masses over 0.8 M$_\odot$. Any systems confirmed to possess high-mass WDs could then be singled out for phase-resolved spectroscopic follow-up to obtain the radial velocity semi-amplitudes for both components and allow for high-precision model-independent measurements of the mass and radius. These can then be compared with the theoretical

tracks shown in Figure 5.5 to give an estimate of the core-composition.

5.2 Summary

In this thesis, I have explained the development and testing of an MCMC code that enables the characterisation of WD + M binaries from multi-band eclipse photometry together with a parallax measurement. I have demonstrated that this method can achieve a significantly higher precision than other methods currently used to parametrise this population of close binaries, and have shown it to be accurate when compared to model-independent measurements. Importantly, I have shown that the clean separation of the two binary components – afforded by the total eclipse of the WD – means that stellar and binary parameters derived via eclipse photometry are far more robust in systems where one component dominates the total flux than previous methods.

Following on from this, I have carried out a programme of high-speed multi-band photometric follow-up of eclipsing WD + M binaries discovered by the ZTF survey using the frame-transfer cameras, ULTRACAM and HiPERCAM. This has resulted in the discovery of at least seven new eclipsing WD + M systems containing magnetic WDs, with two further candidates yet to be confirmed. From these data, I have also characterised 34 of the new eclipsing PCEBs, increasing the published sample of eclipsing WD + M systems by around 50 per cent and finding four with secondary masses and temperatures indicative of substellar companions, doubling the number of eclipsing WD + BD binaries known. In addition, I have discovered and characterised the first two known ZZ Ceti WDs in eclipsing WD + M PCEBs, demonstrating that it is possible to fit their light curves with the application of a GP to model the pulsations.

Finally, I have presented some of the interesting lines of research that naturally follow on from the work carried out in this thesis.

Bibliography

Abbott B. P., et al., 2017, Phys. Rev. Lett., 119, 161101

Allard F., Homeier D., Freytag B., Sharp C. M., 2012, in Reylé C., Charbonnel C., Schultheis M., eds, EAS Publications Series Vol. 57, EAS Publications Series. pp 3–43 (arXiv:1206.1021), doi:10.1051/eas/1257001

Altavilla G., et al., 2015, Astronomische Nachrichten, 336, 515

Altavilla G., et al., 2021, MNRAS, 501, 2848

Althaus L. G., García-Berro E., Isern J., Córsico A. H., 2005, A&A, 441, 689

Ambikasaran S., Foreman-Mackey D., Greengard L., Hogg D. W., O'Neil M., 2015, IEEE Transactions on Pattern Analysis and Machine Intelligence, 38, 252

Bagnulo S., Landstreet J. D., 2021, MNRAS, 507, 5902

Baraffe I., Homeier D., Allard F., Chabrier G., 2015, A&A, 577, A42

Bédard A., Bergeron P., Brassard P., Fontaine G., 2020, ApJ, 901, 93

Bellm E. C., et al., 2019, PASP, 131, 018002

Best W. M. J., Liu M. C., Magnier E. A., Dupuy T. J., 2021, AJ, 161, 42

Beuermann K., et al., 2013, A&A, 558, A96

Bloemen S., et al., 2016, in Hall H. J., Gilmozzi R., Marshall H. K., eds, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series Vol. 9906, Ground-based and Airborne Telescopes VI. p. 990664, doi:10.1117/12.2232522

Blouin S., Dufour P., Allard N. F., 2018, ApJ, 863, 184

Boffin H. M. J., Jones D., 2019, The Importance of Binaries in the Formation and Evolution of Planetary Nebulae, doi:10.1007/978-3-030-25059-1.

Bohlin R. C., Colina L., Finley D. S., 1995, AJ, 110, 1316

Buchner J., et al., 2014, A&A, 564, A125

Camacho J., Torres S., García-Berro E., Zorotovic M., Schreiber M. R., Rebassa-Mansergas A., Nebot Gómez-Morán A., Gänsicke B. T., 2014, A&A, 566, A86

Carroll B. W., Ostlie D. A., 2017, An introduction to modern astrophysics, Second Edition

Casewell S. L., et al., 2018, MNRAS, 476, 1405

Casewell S. L., et al., 2020a, MNRAS, 497, 3571

Casewell S. L., Debes J., Braker I. P., Cushing M. C., Mace G., Marley M. S., Kirkpatrick J. D., 2020b, MNRAS, 499, 5318

Choi J., Dotter A., Conroy C., Cantiello M., Paxton B., Johnson B. D., 2016, ApJ, 823, 102

Claret A., 2000, A&A, 363, 1081

Claret A., 2003, A&A, 406, 623

Claret A., Bloemen S., 2011, A&A, 529, A75

Claret A., Cukanovaite E., Burdge K., Tremblay P. E., Parsons S., Marsh T. R., 2020, A&A, 641, A157

Cojocaru R., Rebassa-Mansergas A., Torres S., García-Berro E., 2017, MNRAS, 470, 1442

Copperwheat C. M., Marsh T. R., Dhillon V. S., Littlefair S. P., Hickman R., Gänsicke B. T., Southworth J., 2010, MNRAS, 402, 1824

Darwin G. H., 1879, Proceedings of the Royal Society of London Series I, 29, 168

Davis P. J., Kolb U., Knigge C., 2012, MNRAS, 419, 287

De Marco O., Passy J.-C., Moe M., Herwig F., Mac Low M.-M., Paxton B., 2011, MNRAS, 411, 2277

Denisenko D. V., Larin I., 2018, arXiv e-prints, p. arXiv:1807.04574

Dhillon V. S., et al., 2007, MNRAS, 378, 825

Dhillon V. S., et al., 2021, MNRAS, 507, 350

Dotter A., 2016, ApJS, 222, 8

Drake A. J., et al., 2009, ApJ, 696, 870

Drake A. J., et al., 2010, arXiv e-prints, p. arXiv:1009.3048

Eggleton P. P., 1983, ApJ, 268, 368

Farihi J., Christopher M., 2004, AJ, 128, 1868

Feroz F., Hobson M. P., Bridges M., 2009, MNRAS, 398, 1601

Fontaine G., Brassard P., 2008, PASP, 120, 1043

Fontaine G., Bergeron P., Billères M., Charpinet S., 2003, ApJ, 591, 1184

Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, PASP, 125, 306

Fukugita M., Ichikawa T., Gunn J. E., Doi M., Shimasaku K., Schneider D. P., 1996, AJ, 111, 1748

Gaia Collaboration et al., 2016, A&A, 595, A1

Gaia Collaboration et al., 2021, A&A, 649, A1

Gianninas A., Bergeron P., Ruiz M. T., 2011, ApJ, 743, 138

Goodman J., Weare J., 2010, Communications in Applied Mathematics and Computational Science, 5, 65

Graham M. J., et al., 2019, PASP, 131, 078001

Greiss S., Gänsicke B. T., Hermes J. J., Steeghs D., Koester D., Ramsay G., Barclay T., Townsley D. M., 2014, MNRAS, 438, 3086

Haefner R., 1989, A&A, 213, L15

Hermes J. J., Montgomery M. H., Winget D. E., Brown W. R., Kilic M., Kenyon S. J., 2012, ApJL, 750, L28

Hermes J. J., et al., 2013a, MNRAS, 436, 3573

Hermes J. J., et al., 2013b, ApJ, 765, 102

Hermes J. J., Kepler S. O., Castanheira B. G., Gianninas A., Winget D. E., Montgomery M. H., Brown W. R., Harrold S. T., 2013c, ApJL, 771, L2

Hermes J. J., et al., 2015, MNRAS, 451, 1701

Hillebrandt W., Niemeyer J. C., 2000, ARAA, 38, 191

Hollands M. A., Tremblay P. E., Gänsicke B. T., Gentile-Fusillo N. P., Toonen S., 2018, MNRAS, 480, 3942

Husser T. O., Wende-von Berg S., Dreizler S., Homeier D., Reiners A., Barman T., Hauschildt P. H., 2013, A&A, 553, A6

Iben Icko J., Livio M., 1993, PASP, 105, 1373

Inight K., Gänsicke B. T., Breedt E., Marsh T. R., Pala A. F., Raddi R., 2021, MNRAS, 504, 2420

Ivanova N., et al., 2013, Astronomy and Astrophysics Review, 21, 59

Ivezić Ž., et al., 2019, ApJ, 873, 111

Izzard R. G., Hall P. D., Tauris T. M., Tout C. A., 2012, IAU Symposium, 283, 95

Jones S., et al., 2013, ApJ, 772, 150

Kesseli A. Y., Muirhead P. S., Mann A. W., Mace G., 2018, AJ, 155, 225

Knigge C., 2006, MNRAS, 373, 484

Koester D., 2010, , 81, 921

Kosakowski A., Kilic M., Brown W. R., Bergeron P., Kupfer T., 2022, MNRAS, 516, 720

Law N. M., et al., 2009, PASP, 121, 1395

Levenberg K., 1944, Quart. Appl. Math., 2, 164

Littlefair S. P., et al., 2014, MNRAS, 445, 2106

Livio M., Soker N., 1988, ApJ, 329, 764

López-Morales M., 2007, ApJ, 660, 732

López-Morales M., Ribas I., 2005, ApJ, 631, 1120

Luo A. L., et al., 2012, Research in Astronomy and Astrophysics, 12, 1243

Mann A. W., et al., 2019, ApJ, 871, 63

Marinoni S., et al., 2016, MNRAS, 462, 3616

Marquardt D. W., 1963, J. Soc. Indust. Appl. Math., 11, 431

Marsh T. R., Dhillon V. S., Duck S. R., 1995, MNRAS, 275, 828

Masci F. J., et al., 2019, PASP, 131, 018003

Miyaji S., Nomoto K., Yokoi K., Sugimoto D., 1980, Publications of the Astronomical Society of Japan, 32, 303

Morrell S., Naylor T., 2019, MNRAS, 489, 2615

Naylor T., 1998, MNRAS, 296, 339

Nebot Gómez-Morán A., et al., 2011, A&A, 536, A43

Nelemans G., Tout C. A., 2005, MNRAS, 356, 753

Nelemans G., Verbunt F., Yungelson L. R., Portegies Zwart S. F., 2000, A&A, 360, 1011

O'Brien M. S., Bond H. E., Sion E. M., 2001, ApJ, 563, 971

Paczynski B., 1976, in Eggleton P., Mitton S., Whelan J., eds, Vol. 73, Structure and Evolution of Close Binary Systems. p. 75

Pala A. F., et al., 2020, MNRAS, 494, 3799

Pancino E., et al., 2012, MNRAS, 426, 1767

Pancino E., et al., 2021, MNRAS, 503, 3660

Panei J. A., Althaus L. G., Chen X., Han Z., 2007, MNRAS, 382, 779

Parsons S. G., Marsh T. R., Copperwheat C. M., Dhillon V. S., Littlefair S. P., Gänsicke B. T., Hickman R., 2010, MNRAS, 402, 2591

Parsons S. G., et al., 2013a, MNRAS, 429, 256

Parsons S. G., Marsh T. R., Gänsicke B. T., Schreiber M. R., Bours M. C. P., Dhillon V. S., Littlefair S. P., 2013b, MNRAS, 436, 241

Parsons S. G., et al., 2015, MNRAS, 449, 2194

Parsons S. G., Rebassa-Mansergas A., Schreiber M. R., Gänsicke B. T., Zorotovic M., Ren J. J., 2016, MNRAS, 463, 2125

Parsons S. G., et al., 2017a, MNRAS, 470, 4473

Parsons S. G., et al., 2017b, MNRAS, 471, 976

Parsons S. G., et al., 2018, MNRAS, 481, 1083

Parsons S. G., et al., 2020, Nature Astronomy, 4, 690

Parsons S. G., Gänsicke B. T., Schreiber M. R., Marsh T. R., Ashley R. P., Breedt E., Littlefair S. P., Meusinger H., 2021, MNRAS, 502, 4305

Pyrzas S., et al., 2015, MNRAS, 447, 691

Raghavan D., et al., 2010, ApJS, 190, 1

Rebassa-Mansergas A., Gänsicke B. T., Rodríguez-Gil P., Schreiber M. R., Koester D., 2007, MNRAS, 382, 1377

Rebassa-Mansergas A., Gänsicke B. T., Schreiber M. R., Koester D., Rodríguez-Gil P., 2010, MNRAS, 402, 620

Rebassa-Mansergas A., Nebot Gómez-Morán A., Schreiber M. R., Gänsicke B. T., Schwope A., Gallardo J., Koester D., 2012, MNRAS, 419, 806

Rebassa-Mansergas A., Ren J. J., Parsons S. G., Gänsicke B. T., Schreiber M. R., García-Berro E., Liu X. W., Koester D., 2016, MNRAS, 458, 3808

Rebassa-Mansergas A., et al., 2017, MNRAS, 472, 4193

Rebassa-Mansergas A., et al., 2021, MNRAS, 506, 5201

Ren J. J., et al., 2014, A&A, 570, A107

Ren J. J., Rebassa-Mansergas A., Parsons S. G., Liu X. W., Luo A. L., Kong X., Zhang H. T., 2018, MNRAS, 477, 4641

Ren J. J., et al., 2020, ApJ, 905, 38

Ritter H., Zhang Z. Y., Kolb U., 2000, A&A, 360, 959

Romero A. D., Lauffer G. R., Istrate A. G., Parsons S. G., 2022a, MNRAS, 510, 858

Romero A. D., et al., 2022b, MNRAS, 511, 1574

Saumon D., Marley M. S., 2008, ApJ, 689, 1327

Scherbak P., Fuller J., 2023, MNRAS, 518, 3966

Schreiber M. R., Zorotovic M., Wijnen T. P. G., 2016, MNRAS, 455, L16

Schreiber M. R., Belloni D., Gänsicke B. T., Parsons S. G., Zorotovic M., 2021, Nature Astronomy, 5, 648

Smith J. A., et al., 2002, AJ, 123, 2121

Spring K., Fellers T., Davidson M., 2023, Introduction to Charge-Coupled Devices (CCDs), https://www.microscopyu.com/digital-imaging/introduction-to-charge-coupled-devices-ccds

Steeghs D., et al., 2022, MNRAS, 511, 2405

Steele P. R., Burleigh M. R., Dobbie P. D., Jameson R. F., Barstow M. A., Satterthwaite R. P., 2011, MNRAS, 416, 2768

Steinfadt J. D. R., Bildsten L., Kaplan D. L., Fulton B. J., Howell S. B., Marsh T. R., Ofek E. O., Shporer A., 2012, PASP, 124, 1

Stobie R. S., et al., 1997, MNRAS, 287, 848

Teledyne Imaging 2022, How a Charge Coupled Device (CCD) Image Sensor Works, https://www.teledyne-e2v.com/download/4eceef0b-685b-4f38-a0f9-486455d46efb/

Toonen S., Nelemans G., 2013, A&A, 557, A87

Tremblay P. E., Bergeron P., Gianninas A., 2011, ApJ, 730, 128

Tylenda R., et al., 2011, A&A, 528, A114

Vanderburg A., et al., 2020, Nature, 585, 363

Vernet J., et al., 2011, A&A, 536, A105

Webbink R. F., 1984, ApJ, 277, 355

Wickramasinghe D. T., Ferrario L., 2000, PASP, 112, 873

Wild J. F., et al., 2022, MNRAS, 509, 5086

Willems B., Kolb U., 2004, A&A, 419, 1057

Winget D. E., Kepler S. O., 2008, ARAA, 46, 157

Zenati Y., Toonen S., Perets H. B., 2019, MNRAS, 482, 1135

Zorotovic M., Schreiber M., 2022, MNRAS, 513, 3587

Zorotovic M., Schreiber M. R., Gänsicke B. T., Nebot Gómez-Morán A., 2010, A&A, 520, A86

Zorotovic M., Schreiber M. R., Gänsicke B. T., 2011, A&A, 536, A42

Zorotovic M., et al., 2016, MNRAS, 457, 3867

de Kool M., van den Heuvel E. P. J., Pylyser E., 1987, A&A, 183, 47

van Roestel J., et al., 2021, ApJL, 919, L26

van den Heuvel E. P. J., 1976, in Eggleton P., Mitton S., Whelan J., eds, Vol. 73, Structure and Evolution of Close Binary Systems. p. 35

Appendix A

ULTRACAM and HiPERCAM standard star catalogue

As mentioned in the Chapter 2, while the HiPERCAM photometric system approximates the SDSS system (Figure 2.8), there are some departures, most significantly in the u_s band where the difference is on the order of a few tenths of a magnitude. To assess any colour terms I follow the procedure of Wild et al. (2022), using synthetic photometry of main sequence stars and WDs to provide corrections as a function of SDSS colour.

For main sequence stars I use PHOENIX spectral models (Allard et al., 2012) at effective temperatures and surface gravities defined by a MIST (Dotter, 2016; Choi et al., 2016) isochrone with an age of $10^{8.5}$ yr, covering masses ranging from $0.1 - 3$ $M_\odot$ and surface gravities from $3.7 - 4.7$ dex. For WDs we use Koester (2010) models with a $\log(g) = 8.5$ dex over the full range of available model temperatures. I then use Equation 2.7 to generate synthetic photometry for the SDSS and HiPERCAM Super-SDSS systems. Corrections as a function of colour are shown in Figure A.1 with best fit colour terms shown in Table A.1 and Table A.2 for WDs and main sequence stars respectively. Corrections should be possible to an accuracy of a couple of per cent in the g_s, r_s, i_s, and z_s bands using main sequence stars but it is clear that no easy correction can be made for the u_s band where there is no consistent correlation with colour. As such, any correction to the u_s-band magnitudes using main sequence stars in combination with these colour terms should be avoided if possible. Using WDs make this a lot easier with tight relations between the colour and corrections that only weakly depend on the surface gravity of the WD (as demonstrated by the similarity of $\log(g) = 8.0$ and $\log(g) = 8.5$ relations).

In order to make flux calibration simpler and more robust, we define a set of ULTRA-

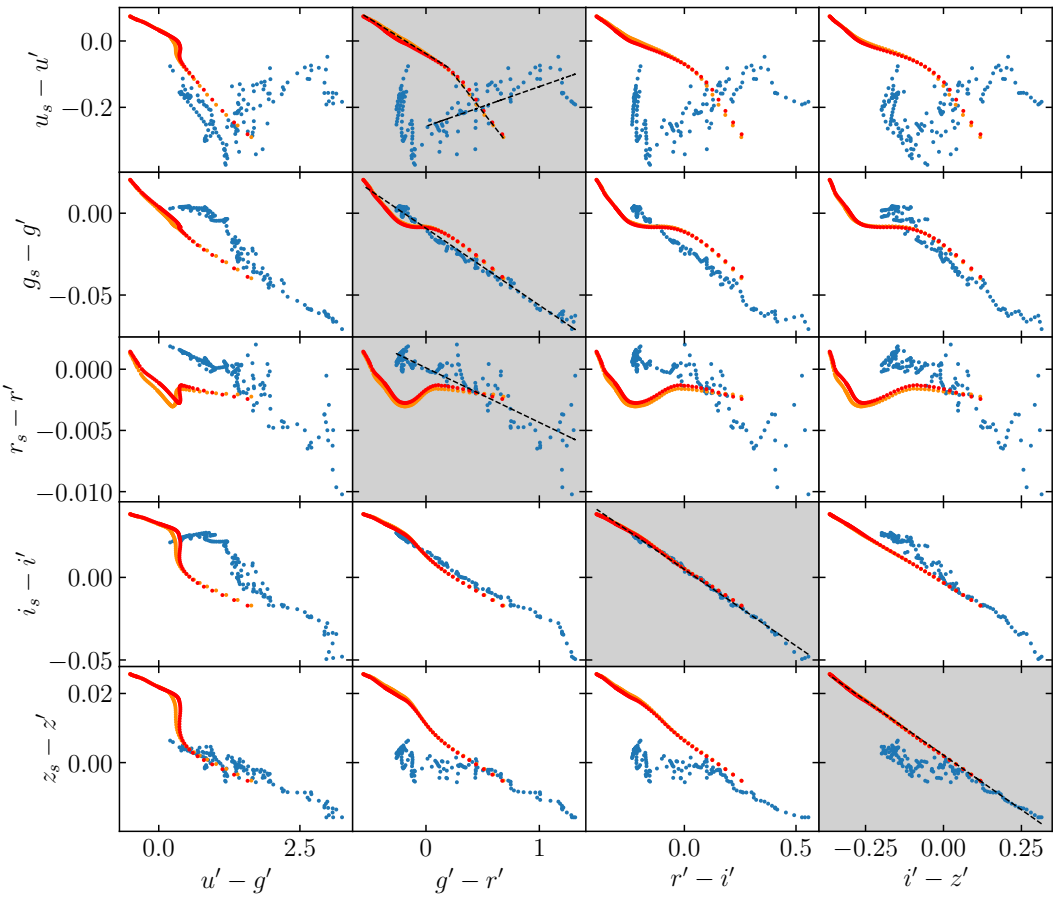

Figure A.1: Magnitude offsets between the HiPERCAM Super-SDSS ($u_s g_s r_s i_s z_s$) photometric system (Dhillon et al., 2021) and the SDSS primed ($u'g'r'i'z'$) photometric system (Fukugita et al., 1996) as a function of SDSS colour for main sequence stars (Allard et al., 2012) (blue) with age = $10^{8.5}$ yr and for WDs (Koester, 2010) with a $\log(g) = 8.0$ (red) and $\log(g) = 8.5$ (orange). Shaded plots indicate relations to which colour terms are fit and these best fit corrections (listed in Table A.1 and Table A.2) are shown by a black dashed line.

Table A.1: SDSS to HiPERCAM Super-SDSS colour terms for WDs. Validity shows the range of colours spanned by the models that the colour terms were fit to. These take the form of a straight line, e.g. $u_s - u' = -0.211(g' - r') - 0.038$

Correction	Gradient	Variable	y-intercept	Validity
$u_s - u'$	-0.211	$g' - r'$	-0.038	$-0.55 < g' - r' \leq 0.20$
	-0.438	$g' - r'$	0.006	$0.20 < g' - r' \leq 0.70$
$g_s - g'$	-0.047	$g' - r'$	-0.009	$-0.55 < g' - r' \leq 0.70$
$r_s - r'$	–	–	–	–
$i_s - i'$	-0.093	$r' - i'$	0.005	$-0.40 < r' - i' \leq 0.25$
$z_s - z'$	-0.047	$i' - z'$	-0.009	$-0.35 < i' - z' \leq 0.10$

Table A.2: As Table A.1 but for main sequence models.

Correction	Gradient	Variable	y-intercept	Validity
$u_s - u'$	0.120	$g' - r'$	-0.257	$0.00 < g' - r' \leq 1.30$
$g_s - g'$	-0.047	$g' - r'$	-0.009	$-0.25 < g' - r' \leq 1.30$
$r_s - r'$	-0.004	$g' - r'$	0.000	$-0.25 < g' - r' \leq 1.30$
$i_s - i'$	-0.093	$r' - i'$	0.005	$-0.25 < g' - r' \leq 0.55$
$z_s - z'$	0.000	$i' - z'$	0.000	$-0.20 < i' - z' \leq 0.10$
	-0.047	$i' - z'$	-0.009	$0.10 < i' - z' \leq 0.30$

CAM and HiPERCAM standard stars (Table A.3 and Table A.4 respectively). These standard stars are the *Gaia* spectro-photometric standard stars collated and evaluated by Pancino et al. (2012); Altavilla et al. (2015); Marinoni et al. (2016); Altavilla et al. (2021); Pancino et al. (2021) who provide high quality spectra of these standards covering the full range of the ULTRACAM and HiPERCAM photometric system. We again use Equation 2.7 to produce synthetic AB magnitudes for these stars in both the ULTRACAM and HiPERCAM systems. Pancino et al. (2021) mentions that the scatter in the spectro-photometric standard stars when compared with literature is of order 1 per cent with discrepant behaviour of a similar order in the red for faint blue stars. Additionally, due to the use of spectral models to extend the flux tables below 400 nm and above 800 nm, we estimate an uncertainty of 2 per cent in u_s and z_s, and 1 per cent in g_s, r_s, and i_s.

Table A.3: *Gaia* spectro-photometric standard stars (Pancino et al., 2012; Altavilla et al., 2015; Marinoni et al., 2016; Altavilla et al., 2021; Pancino et al., 2021) with AB magnitudes computed for the ULTRACAM Super-SDSS ($u_s g_s r_s i_s z_s$) photometric system using the flux tables of Pancino et al. (2021). The 'Type' column indicates the status of a flux standard as either 'Pillar', 'Primary', or 'Secondary' as described in Pancino et al. (2012) ('0', '1', and '2' respectively in the table). The pillars denoted here are the same three stars on which the CALSPEC system is based (Bohlin et al., 1995) and the primary stars are all bright, well-known spectro-photometric standards that are already tied to – or are easy to tie to – the CALSPEC flux scale. Secondary standards are then calibrated from these primary stars. The 'Stability' column shows which standards have been confirmed as photometrically constant by the variability monitoring campaign (Marinoni et al., 2016). Standards that are not yet confirmed as photometrically constant are still considered likely to be constant (Marinoni et al., 2016, see section 3.4). We therefore choose not to discard them.

Name	RA	DEC	u_s	g_s	r_s	i_s	z_s	Type	SpType	Stability
WD 0004+330	00:07:32.26	33:17:27.60	13.125	13.523	14.042	14.424	14.781	2	DA1	
WD 0018-267	00:21:30.73	-26:26:11.46	15.302	14.094	13.61	13.433	13.385	2	DA9	
WD 0038+555	00:41:21.99	55:50:08.40	14.019	13.98	14.113	14.28	14.471	2	DQ5	Confirmed
LTT377	00:41:30.47	-33:37:32.03	13.889	11.296	9.948	9.283	9.062	2	K9	Confirmed
WD 0046+051	00:49:09.90	05:23:19.01	13.574	12.567	12.298	12.268	12.407	2	DZ7	Confirmed
WD 0047-524	00:50:03.68	-52:08:15.60	14.171	14.054	14.439	14.738	15.069	2	DA2	
WD 0050-332	00:53:17.44	-32:59:56.60	12.742	13.086	13.607	13.932	14.222	2	DA1	
WD 0109-264	01:12:11.65	-26:13:27.69	12.691	12.875	13.372	13.722	14.058	2	DA1	
WD 0123-262	01:25:24.45	-26:00:43.90	15.58	15.149	15.006	15.013	15.145	2	DC	Confirmed
G245-31	01:38:39.39	69:38:01.50	15.953	14.86	14.228	13.985	13.865	2	K	Confirmed
GJ70	01:43:20.18	04:19:17.97	14.173	11.834	10.416	9.297	8.761	2	M2	Confirmed
LTT1020	01:54:50.27	-27:28:35.74	12.614	11.745	11.353	11.227	11.216	1	G	Confirmed
LP885-23	02:07:06.33	-30:24:22.90	17.62	14.914	13.515	12.152	11.516	2	M0	Confirmed
EG21	03:10:31.02	-68:36:03.39	11.467	11.254	11.578	11.851	12.179	1	DA3	Confirmed
HG7-15	03:48:11.86	07:08:46.47	14.134	11.528	10.25	9.716	9.409	2	M1	Confirmed
LTT1788	03:48:22.67	-39:08:37.20	14.098	13.334	13.018	12.909	12.908	1	F	Confirmed
GD50	03:48:50.20	-00:58:31.20	13.367	13.788	14.308	14.681	15.081	1	DA2	
HZ2	04:12:43.55	11:51:49.00	13.706	13.703	14.092	14.407	14.714	1	DA3	

Name	RA	DEC	u_s	g_s	r_s	i_s	z_s	Type	SpType	Stability
WD 0455-282	04:57:13.90	-28:07:54.00	13.237	13.64	14.199	14.578	14.93	2	DA1	
WD 0501-289	05:03:55.51	-28:54:34.57	13.057	13.572	14.141	14.557	14.936	2	DO	Confirmed
G191-B2B	05:05:30.61	52:49:51.95	11.031	11.483	12.029	12.423	12.813	0	DA0	Confirmed
GD71	05:52:27.63	15:53:13.37	12.451	12.775	13.277	13.651	14.034	0	DA1	Confirmed
LTT2415	05:56:24.74	-27:51:32.35	13.156	12.354	12.11	12.033	12.053	1	G	Confirmed
HD270477	05:59:33.36	-67:01:13.72	11.528	10.506	10.314	10.306	10.363	2	F8	
HD271747	05:59:58.62	-66:06:08.91	12.904	11.699	11.294	11.195	11.2	2	G0	
HD271759	06:00:41.34	-66:03:14.03	12.011	10.863	10.881	10.99	11.1	2	A2	Confirmed
HD271783	06:02:11.36	-66:34:59.13	13.31	12.157	11.775	11.698	11.761	2	F5	
WD 0604-203	06:06:13.39	-20:21:07.20	11.788	11.825	12.285	12.605	12.937	2	DA	Confirmed
WD 0621-376	06:23:12.63	-37:41:28.01	11.341	11.783	12.343	12.738	13.111	2	DA1	Confirmed
HILT600	06:45:13.37	02:08:14.70	10.82	10.458	10.455	10.535	10.651	1	B1	Confirmed
WD 0644+375	06:47:37.99	37:30:57.07	11.814	11.876	12.275	12.603	12.944	2	DA2	
WD 0646-253	06:48:56.09	-25:23:47.00	13.199	13.411	13.881	14.244	14.532	2	DA2	Confirmed
G193-26	07:03:26.29	54:52:06.00	14.019	13.232	12.802	12.637	12.575	2	G	Confirmed
WD 0721-276	07:23:20.10	-27:47:21.60	13.949	14.288	14.788	15.145	15.511	2	DA1	Confirmed
LTT3218	08:41:32.56	-32:56:34.90	12.274	11.876	11.908	12.023	12.203	1	DA	Confirmed
G114-25	08:59:03.37	-06:23:46.19	13.05	12.179	11.73	11.552	11.524	2	F7	Confirmed
WD 0859-039	09:02:17.30	-04:06:55.45	12.841	12.979	13.401	13.736	14.096	2	DA2	Confirmed
GD108	10:00:47.37	-07:33:30.50	13.198	13.328	13.772	14.099	14.431	1	B	Confirmed
LTT3864	10:32:13.60	-35:37:41.80	13.229	12.366	12.027	11.915	11.898	1	F	Confirmed
WD 1031-114	10:33:42.76	-11:41:38.35	12.619	12.801	13.236	13.566	13.91	2	DA2	Confirmed
WD 1034+001	10:37:03.81	-00:08:19.30	12.416	12.916	13.49	13.901	14.288	2	DOZ1	Confirmed
Feige34	10:39:36.74	43:06:09.25	10.42	10.871	11.411	11.777	12.152	1	DO	Confirmed
WD 1105-048	11:07:59.95	-05:09:25.90	13.193	12.953	13.26	13.541	13.85	2	DA3	Confirmed
SDSS09310	11:38:02.62	57:29:23.89	16.003	15.093	14.989	15.011	15.082	2	A0/F3	Confirmed
LTT4364	11:45:42.92	-64:50:29.46	11.778	11.518	11.479	11.556	11.709	1	DQ6	Confirmed

Name	RA	DEC	u_s	g_s	r_s	i_s	z_s	Type	SpType	Stability
Feige56	12:06:47.23	11:40:12.64	11.179	10.867	11.207	11.494	11.741	1	B5p	Confirmed
HD106355	12:14:10.53	-17:14:20.19	12.858	10.66	9.716	9.358	9.119	2	G8IV	Confirmed
Feige66	12:37:23.52	25:03:59.87	9.901	10.228	10.75	11.134	11.505	1	O	Confirmed
SA104-428	12:41:41.28	-00:26:26.20	15.087	13.105	12.324	12.063	11.911	2	G8	Confirmed
SA104-490	12:44:33.46	-00:25:51.70	13.902	12.788	12.441	12.367	12.378	2	G3	Confirmed
GD153	12:57:02.33	22:01:52.52	12.692	13.072	13.594	13.984	14.361	0	DA1	Confirmed
G14-24	13:02:01.58	-02:05:21.42	14.275	13.16	12.57	12.339	12.226	2	K0	Confirmed
GJ2097	13:07:04.31	20:48:38.54	15.946	13.237	11.822	10.797	10.289	2	M1	Confirmed
HZ44	13:23:35.26	36:07:59.51	10.971	11.382	11.904	12.307	12.685	1	O	Confirmed
SA105-663	13:37:30.34	-00:13:17.37	9.929	8.861	8.729	8.708	8.799	1	F2	Confirmed
GJ521	13:39:24.10	46:11:11.37	13.52	10.948	9.645	8.752	8.301	2	M2	Confirmed
HD121968	13:58:51.17	-02:54:52.32	9.95	10.034	10.437	10.752	11.052	1	B1	Confirmed
WD 1408+323	14:10:26.95	32:08:36.10	13.912	13.828	14.185	14.491	14.827	2	DA3	
SDSS09626	14:29:51.06	39:28:25.43	16.082	15.007	15.047	15.151	15.247	2	A0	Confirmed
G15-10	15:09:46.02	-04:45:06.61	13.232	12.32	11.803	11.61	11.548	2	G2	Confirmed
WD 1509+322	15:11:27.66	32:04:17.80	14.265	14.005	14.292	14.565	14.86	2	DA3	Confirmed
M5-S1490	15:17:38.64	02:02:25.60	16.328	14.575	13.795	13.478	13.308	2		Confirmed
G167-50	15:35:31.55	27:51:02.20	14.829	13.834	13.294	13.092	13.016	2	G	Confirmed
SA107-544	15:36:48.10	-00:15:07.11	10.339	9.157	8.97	8.946	8.986	1	F3	Confirmed
LTT6248	15:38:59.66	-28:35:36.87	12.817	11.999	11.651	11.527	11.523	1	A	Confirmed
G179-54	15:46:08.25	39:14:16.40	14.668	13.771	13.312	13.15	13.094	2	F	Confirmed
G224-83	15:46:14.68	62:26:39.60	13.838	12.993	12.553	12.393	12.339	2	K	Confirmed
G16-20	15:58:18.62	02:03:06.11	12.069	11.051	10.605	10.442	10.38	2	K	Confirmed
P177-D	15:59:13.57	47:36:41.90	14.991	13.734	13.27	13.135	13.113	1	G0	Confirmed
WD 1606+422	16:08:22.20	42:05:43.20	14.173	13.756	13.972	14.2	14.465	2	DA4	Confirmed
WD 1615-154	16:17:55.26	-15:35:51.90	12.896	13.189	13.666	13.999	14.329	2	DA2	
EG274	16:23:33.84	-39:13:46.16	10.699	10.804	11.255	11.58	11.922	1	DA2	

Name	RA	DEC	u_s	g_s	r_s	i_s	z_s	Type	SpType	Stability
G180-58	16:28:16.87	44:40:38.28	12.611	11.612	11.085	10.891	10.808	2	G/K	Confirmed
WD 1626+368	16:28:25.03	36:46:15.40	13.993	13.808	13.862	13.994	14.193	2	DZA6	Confirmed
G170-47	17:32:41.63	23:44:11.64	10.198	9.201	8.717	8.53	8.459	2	G0	Confirmed
1743045	17:43:04.48	66:55:01.60	14.698	13.56	13.535	13.605	13.67	1	A5	Confirmed
KF08T3	17:55:16.23	66:10:11.70	15.601	13.747	12.989	12.74	12.594	1	K0	Confirmed
2MASS J1757+6703	17:57:13.25	67:03:40.90	13.022	11.845	11.819	11.905	11.979	2	A3	Confirmed
KF06T2	17:58:37.99	66:46:52.20	16.93	14.538	13.579	13.237	13.02	1	K1	Confirmed
BD+661071	18:02:10.92	66:12:26.39	11.7	10.694	10.472	10.45	10.489	2	F5	Confirmed
1805292	18:05:29.28	64:27:52.00	13.358	12.236	12.38	12.535	12.654	1	A6	Confirmed
1812095	18:12:09.57	63:29:42.30	12.929	11.733	11.787	11.911	12.01	1	A5	Confirmed
LTT7379	18:36:25.95	-44:18:36.94	11.56	10.482	10.069	9.925	9.955	1	G0	Confirmed
G184-17	18:40:29.27	19:36:06.65	15.661	14.466	13.832	13.609	13.482	2	K	Confirmed
WD 1845+019	18:47:39.08	01:57:35.62	12.489	12.719	13.173	13.395	13.426	2	DA2	
GJ745A	19:07:05.56	20:53:16.97	14.315	11.591	10.145	9.104	8.544	2	M2	Confirmed
GJ745B	19:07:13.20	20:52:37.24	14.311	11.589	10.146	9.092	8.531	2	M2	Confirmed
WD 1914-598	19:18:44.84	-59:46:33.80	14.284	14.213	14.598	14.885	15.235	2	DA	Confirmed
EG131	19:20:34.93	-07:40:00.05	12.236	12.218	12.366	12.544	12.779	1	DBQA5	Confirmed
WD 1919+145	19:21:40.40	14:40:43.00	13.123	12.877	13.175	13.467	13.782	2	DA3	Confirmed
G23-14	19:51:49.61	05:36:45.84	12.242	11.071	10.487	10.27	10.161	2	G5	Confirmed
WD 2004-605	20:09:05.24	-60:25:41.60	12.695	13.077	13.612	13.97	14.363	2	DA1	Confirmed
LTT7987	20:10:56.85	-30:13:06.64	12.389	12.122	12.435	12.696	13.011	1	DA4	Confirmed
WD 2032+248	20:34:21.88	25:03:49.72	11.332	11.34	11.739	12.075	12.421	2	DA2	
WD 2034-532	20:38:16.84	-53:04:25.40	14.083	14.212	14.535	14.781	15.045	2	DB4	Confirmed
G24-25	20:40:16.10	00:33:19.74	11.756	10.815	10.398	10.268	10.212	2	G0	Confirmed
WD 2039-682	20:44:21.47	-68:05:21.30	13.342	13.198	13.511	13.773	14.101	2	DA3	Confirmed
WD 2047+372	20:49:06.69	37:28:13.90	13.221	12.915	13.175	13.428	13.726	2	DA3	Confirmed
WD 2111+498	21:12:44.05	50:06:17.80	12.435	12.783	13.289	13.659	14.03	2	DA1	

Name	RA	DEC	u_s	g_s	r_s	i_s	z_s	Type	SpType	Stability
WD 2105-820	21:13:13.90	-81:49:04.00	14.004	13.659	13.741	13.88	14.084	2	DA5	Confirmed
WD 2117+539	21:18:56.27	54:12:41.25	12.532	12.245	12.541	12.825	13.123	2	DA3	Confirmed
WD 2134+218	21:36:36.30	22:04:33.00	14.442	14.324	14.666	14.964	15.265	2	DA3	
WD 2140+207	21:42:41.00	20:59:58.24	13.445	13.247	13.253	13.354	13.499	2	DQ6	Confirmed
G93-48	21:52:25.38	02:23:19.56	12.741	12.618	12.966	13.258	13.589	1	DA3	
WD 2216-657	22:19:48.35	-65:29:18.11	14.606	14.485	14.596	14.74	14.94	2	DZ5	Confirmed
LTT9239	22:52:41.03	-20:35:32.89	13.292	12.335	11.895	11.726	11.734	1	F	Confirmed
WD 2251-634	22:55:10.00	-63:10:27.00	14.041	13.963	14.343	14.612	14.938	2	DA	Confirmed
WD 2309+105	23:12:21.62	10:47:04.25	12.356	12.803	13.342	13.731	14.122	2	DA1	
Feige110	23:19:58.40	-05:09:56.21	11.14	11.539	12.077	12.443	12.821	1	O	Confirmed
WD 2329+407	23:31:35.65	41:01:30.70	13.948	13.739	14.056	14.345	14.657	2	DA3	
WD 2331-475	23:34:02.20	-47:14:26.50	12.73	13.14	13.68	14.039	14.345	2	DA1	
WD 2352+401	23:54:56.25	40:27:30.10	15.26	14.992	14.917	14.977	15.137	2	DQ6	Confirmed

Table A.4: The same as in Table A.3 but for the HiPERCAM Super-SDSS system.

Name	RA	DEC	u_s	g_s	r_s	i_s	z_s	Type	SpType	Stability
WD 0004+330	00:07:32.26	33:17:27.60	13.158	13.531	14.023	14.426	14.742	2	DA1	
WD 0018-267	00:21:30.73	-26:26:11.46	15.252	14.083	13.621	13.431	13.391	2	DA9	
WD 0038+555	00:41:21.99	55:50:08.40	14.014	13.98	14.106	14.28	14.442	2	DQ5	Confirmed
LTT377	00:41:30.47	-33:37:32.03	13.783	11.279	9.973	9.275	9.06	2	K9	Confirmed
WD 0046+051	00:49:09.90	05:23:19.01	13.578	12.556	12.301	12.268	12.395	2	DZ7	Confirmed
WD 0047-524	00:50:03.68	-52:08:15.60	14.172	14.059	14.422	14.741	15.036	2	DA2	
WD 0050-332	00:53:17.44	-32:59:56.60	12.767	13.095	13.589	13.934	14.191	2	DA1	
WD 0109-264	01:12:11.65	-26:13:27.69	12.695	12.883	13.355	13.724	14.025	2	DA1	
WD 0123-262	01:25:24.45	-26:00:43.90	15.531	15.144	15.008	15.013	15.13	2	DC	Confirmed
G245-31	01:38:39.39	69:38:01.50	15.886	14.849	14.243	13.983	13.865	2	K	Confirmed
GJ70	01:43:20.18	04:19:17.97	13.931	11.815	10.445	9.286	8.799	2	M2	Confirmed
LTT1020	01:54:50.27	-27:28:35.74	12.542	11.736	11.361	11.226	11.215	1	G	Confirmed
LP885-23	02:07:06.33	-30:24:22.90	17.474	14.894	13.549	12.138	11.558	2	M0	Confirmed
EG21	03:10:31.02	-68:36:03.39	11.468	11.258	11.563	11.855	12.151	1	DA3	Confirmed
HG7-15	03:48:11.86	07:08:46.47	14.059	11.513	10.274	9.71	9.426	2	M1	Confirmed
LTT1788	03:48:22.67	-39:08:37.20	14.013	13.327	13.025	12.907	12.902	1	F	Confirmed
GD50	03:48:50.20	-00:58:31.20	13.401	13.797	14.29	14.683	15.04	1	DA2	
HZ2	04:12:43.55	11:51:49.00	13.719	13.709	14.076	14.41	14.688	1	DA3	
WD 0455-282	04:57:13.90	-28:07:54.00	13.262	13.649	14.181	14.581	14.887	2	DA1	
WD 0501-289	05:03:55.51	-28:54:34.57	13.086	13.583	14.122	14.559	14.899	2	DO	Confirmed
G191-B2B	05:05:30.61	52:49:51.95	11.069	11.492	12.01	12.427	12.778	0	DA0	Confirmed
GD71	05:52:27.63	15:53:13.37	12.48	12.783	13.258	13.654	14.0	0	DA1	Confirmed
LTT2415	05:56:24.74	-27:51:32.35	13.045	12.348	12.115	12.032	12.041	1	G	Confirmed
HD270477	05:59:33.36	-67:01:13.72	11.392	10.501	10.316	10.306	10.354	2	F8	
HD271747	05:59:58.62	-66:06:08.91	12.824	11.689	11.303	11.194	11.196	2	G0	
HD271759	06:00:41.34	-66:03:14.03	11.803	10.862	10.877	10.991	11.09	2	A2	Confirmed
HD271783	06:02:11.36	-66:34:59.13	13.209	12.148	11.783	11.697	11.755	2	F5	
WD 0604-203	06:06:13.39	-20:21:07.20	11.773	11.831	12.269	12.608	12.903	2	DA	Confirmed

Name	RA	DEC	u_s	g_s	r_s	i_s	z_s	Type	SpType	Stability
WD 0621-376	06:23:12.63	-37:41:28.01	11.372	11.792	12.325	12.741	13.069	2	DA1	Confirmed
HILT600	06:45:13.37	02:08:14.70	10.751	10.458	10.452	10.535	10.631	1	B1	Confirmed
WD 0644+375	06:47:37.99	37:30:57.07	11.832	11.882	12.258	12.606	12.906	2	DA2	
WD 0646-253	06:48:56.09	-25:23:47.00	13.22	13.419	13.863	14.247	14.507	2	DA2	Confirmed
G193-26	07:03:26.29	54:52:06.00	13.955	13.223	12.812	12.635	12.567	2	G	Confirmed
WD 0721-276	07:23:20.10	-27:47:21.60	13.973	14.297	14.771	15.149	15.47	2	DA1	Confirmed
LTT3218	08:41:32.56	-32:56:34.90	12.248	11.875	11.903	12.025	12.186	1	DA	Confirmed
G114-25	08:59:03.37	-06:23:46.19	12.976	12.169	11.741	11.551	11.519	2	F7	Confirmed
WD 0859-039	09:02:17.30	-04:06:55.45	12.859	12.986	13.385	13.739	14.055	2	DA2	Confirmed
GD108	10:00:47.37	-07:33:30.50	13.195	13.337	13.756	14.101	14.393	1	B	Confirmed
LTT3864	10:32:13.60	-35:37:41.80	13.14	12.358	12.034	11.914	11.902	1	F	Confirmed
WD 1031-114	10:33:42.76	-11:41:38.35	12.638	12.808	13.22	13.569	13.871	2	DA2	Confirmed
WD 1034+001	10:37:03.81	-00:08:19.30	12.451	12.925	13.47	13.904	14.246	2	DOZ1	Confirmed
Feige34	10:39:36.74	43:06:09.25	10.451	10.88	11.393	11.781	12.108	1	DO	Confirmed
WD 1105-048	11:07:59.95	-05:09:25.90	13.197	12.957	13.245	13.544	13.814	2	DA3	Confirmed
SDSS09310	11:38:02.62	57:29:23.89	15.855	15.089	14.989	15.01	15.068	2	A0/F3	Confirmed
LTT4364	11:45:42.92	-64:50:29.46	11.745	11.516	11.477	11.557	11.688	1	DQ6	Confirmed
Feige56	12:06:47.23	11:40:12.64	11.068	10.872	11.193	11.496	11.713	1	B5p	Confirmed
HD106355	12:14:10.53	-17:14:20.19	12.822	10.641	9.738	9.355	9.146	2	G8IV	Confirmed
Feige66	12:37:23.52	25:03:59.87	9.924	10.236	10.731	11.138	11.464	1	O	Confirmed
SA104-428	12:41:41.28	-00:26:26.20	15.035	13.088	12.342	12.061	11.914	2	G8	Confirmed
SA104-490	12:44:33.46	-00:25:51.70	13.799	12.779	12.448	12.366	12.367	2	G3	Confirmed
GD153	12:57:02.33	22:01:52.52	12.723	13.081	13.575	13.986	14.326	0	DA1	Confirmed
G14-24	13:02:01.58	-02:05:21.42	14.208	13.148	12.584	12.337	12.224	2	K0	Confirmed
GJ2097	13:07:04.31	20:48:38.54	15.809	13.219	11.85	10.787	10.313	2	M1	Confirmed
HZ44	13:23:35.26	36:07:59.51	10.999	11.389	11.885	12.31	12.653	1	O	Confirmed
SA105-663	13:37:30.34	-00:13:17.37	9.755	8.856	8.73	8.706	8.795	1	F2	Confirmed
GJ521	13:39:24.10	46:11:11.37	13.354	10.933	9.673	8.743	8.32	2	M2	Confirmed
HD121968	13:58:51.17	-02:54:52.32	9.924	10.041	10.423	10.754	11.023	1	B1	Confirmed

Name	RA	DEC	u_s	g_s	r_s	i_s	z_s	Type	SpType	Stability
WD 1408+323	14:10:26.95	32:08:36.10	13.924	13.833	14.169	14.493	14.79	2	DA3	
SDSS09626	14:29:51.06	39:28:25.43	15.865	15.006	15.042	15.15	15.233	2	A0	Confirmed
G15-10	15:09:46.02	-04:45:06.61	13.17	12.31	11.816	11.608	11.548	2	G2	Confirmed
WD 1509+322	15:11:27.66	32:04:17.80	14.267	14.009	14.276	14.567	14.825	2	DA3	Confirmed
M5-S1490	15:17:38.64	02:02:25.60	16.234	14.558	13.813	13.474	13.313	2		Confirmed
G167-50	15:35:31.55	27:51:02.20	14.78	13.824	13.305	13.09	13.014	2	G	Confirmed
SA107-544	15:36:48.10	-00:15:07.11	10.163	9.152	8.973	8.945	8.987	1	F3	Confirmed
LTT6248	15:38:59.66	-28:35:36.87	12.72	11.993	11.659	11.526	11.512	1	A	Confirmed
G179-54	15:46:08.25	39:14:16.40	14.612	13.761	13.322	13.148	13.088	2	F	Confirmed
G224-83	15:46:14.68	62:26:39.60	13.777	12.984	12.563	12.391	12.333	2	K	Confirmed
G16-20	15:58:18.62	02:03:06.11	11.966	11.042	10.615	10.441	10.373	2	K	Confirmed
P177-D	15:59:13.57	47:36:41.90	14.936	13.723	13.279	13.133	13.116	1	G0	Confirmed
WD 1606+422	16:08:22.20	42:05:43.20	14.163	13.758	13.959	14.201	14.435	2	DA4	Confirmed
WD 1615-154	16:17:55.26	-15:35:51.90	12.919	13.198	13.649	14.002	14.285	2	DA2	
EG274	16:23:33.84	-39:13:46.16	10.712	10.81	11.237	11.583	11.883	1	DA2	
G180-58	16:28:16.87	44:40:38.28	12.554	11.601	11.096	10.889	10.804	2	G/K	Confirmed
WD 1626+368	16:28:25.03	36:46:15.40	14.028	13.805	13.856	13.995	14.168	2	DZA6	Confirmed
G170-47	17:32:41.63	23:44:11.64	10.09	9.191	8.728	8.529	8.457	2	G0	Confirmed
1743045	17:43:04.48	66:55:01.60	14.495	13.558	13.531	13.605	13.665	1	A5	Confirmed
KF08T3	17:55:16.23	66:10:11.70	15.568	13.731	13.006	12.738	12.605	1	K0	Confirmed
2MASS J1757+6703	17:57:13.25	67:03:40.90	12.811	11.842	11.816	11.905	11.97	2	A3	Confirmed
KF06T2	17:58:37.99	66:46:52.20	16.856	14.518	13.602	13.233	13.038	1	K1	Confirmed
BD+661071	18:02:10.92	66:12:26.39	11.579	10.688	10.475	10.449	10.479	2	F5	Confirmed
1805292	18:05:29.28	64:27:52.00	13.127	12.238	12.372	12.536	12.649	1	A6	Confirmed
1812095	18:12:09.57	63:29:42.30	12.715	11.731	11.781	11.911	12.005	1	A5	Confirmed
LTT7379	18:36:25.95	-44:18:36.94	11.488	10.472	10.077	9.924	9.943	1	G0	Confirmed
G184-17	18:40:29.27	19:36:06.65	15.621	14.453	13.845	13.607	13.49	2	K	Confirmed
WD 1845+019	18:47:39.08	01:57:35.62	12.513	12.726	13.156	13.396	13.418	2	DA2	
GJ745A	19:07:05.56	20:53:16.97	14.176	11.573	10.171	9.093	8.575	2	M2	Confirmed

Name	RA	DEC	u_s	g_s	r_s	i_s	z_s	Type	SpType	Stability
GJ745B	19:07:13.20	20:52:37.24	14.177	11.57	10.173	9.08	8.562	2	M2	Confirmed
WD 1914-598	19:18:44.84	-59:46:33.80	14.292	14.219	14.582	14.888	15.203	2	DA	Confirmed
EG131	19:20:34.93	-07:40:00.05	12.228	12.219	12.359	12.544	12.75	1	DBQA5	Confirmed
WD 1919+145	19:21:40.40	14:40:43.00	13.128	12.881	13.159	13.469	13.753	2	DA3	Confirmed
G23-14	19:51:49.61	05:36:45.84	12.159	11.059	10.5	10.268	10.164	2	G5	Confirmed
WD 2004-605	20:09:05.24	-60:25:41.60	12.721	13.086	13.594	13.972	14.326	2	DA1	Confirmed
LTT7987	20:10:56.85	-30:13:06.64	12.39	12.126	12.42	12.699	12.979	1	DA4	Confirmed
WD 2032+248	20:34:21.88	25:03:49.72	11.348	11.346	11.721	12.078	12.388	2	DA2	
WD 2034-532	20:38:16.84	-53:04:25.40	14.089	14.214	14.522	14.783	15.018	2	DB4	Confirmed
G24-25	20:40:16.10	00:33:19.74	11.677	10.806	10.406	10.266	10.216	2	G0	Confirmed
WD 2039-682	20:44:21.47	-68:05:21.30	13.354	13.202	13.496	13.776	14.071	2	DA3	Confirmed
WD 2047+372	20:49:06.69	37:28:13.90	13.221	12.918	13.16	13.43	13.693	2	DA3	Confirmed
WD 2111+498	21:12:44.05	50:06:17.80	12.463	12.792	13.27	13.662	13.994	2	DA1	
WD 2105-820	21:13:13.90	-81:49:04.00	13.999	13.66	13.734	13.883	14.061	2	DA5	Confirmed
WD 2117+539	21:18:56.27	54:12:41.25	12.53	12.249	12.526	12.828	13.097	2	DA3	Confirmed
WD 2134+218	21:36:36.30	22:04:33.00	14.454	14.329	14.65	14.966	15.232	2	DA3	
WD 2140+207	21:42:41.00	20:59:58.24	13.425	13.245	13.25	13.354	13.48	2	DQ6	Confirmed
G93-48	21:52:25.38	02:23:19.56	12.749	12.622	12.95	13.259	13.552	1	DA3	
WD 2216-657	22:19:48.35	-65:29:18.11	14.633	14.484	14.589	14.742	14.92	2	DZ5	Confirmed
LTT9239	22:52:41.03	-20:35:32.89	13.228	12.325	11.905	11.725	11.729	1	F	Confirmed
WD 2251-634	22:55:10.00	-63:10:27.00	14.046	13.968	14.328	14.614	14.901	2	DA	Confirmed
WD 2309+105	23:12:21.62	10:47:04.25	12.395	12.812	13.323	13.733	14.079	2	DA1	
Feige110	23:19:58.40	-05:09:56.21	11.168	11.547	12.059	12.446	12.787	1	O	Confirmed
WD 2329+407	23:31:35.65	41:01:30.70	13.953	13.744	14.04	14.347	14.623	2	DA3	
WD 2331-475	23:34:02.20	-47:14:26.50	12.76	13.149	13.662	14.042	14.317	2	DA1	
WD 2352+401	23:54:56.25	40:27:30.10	15.229	14.99	14.917	14.977	15.11	2	DQ6	Confirmed

Appendix B

Light curves

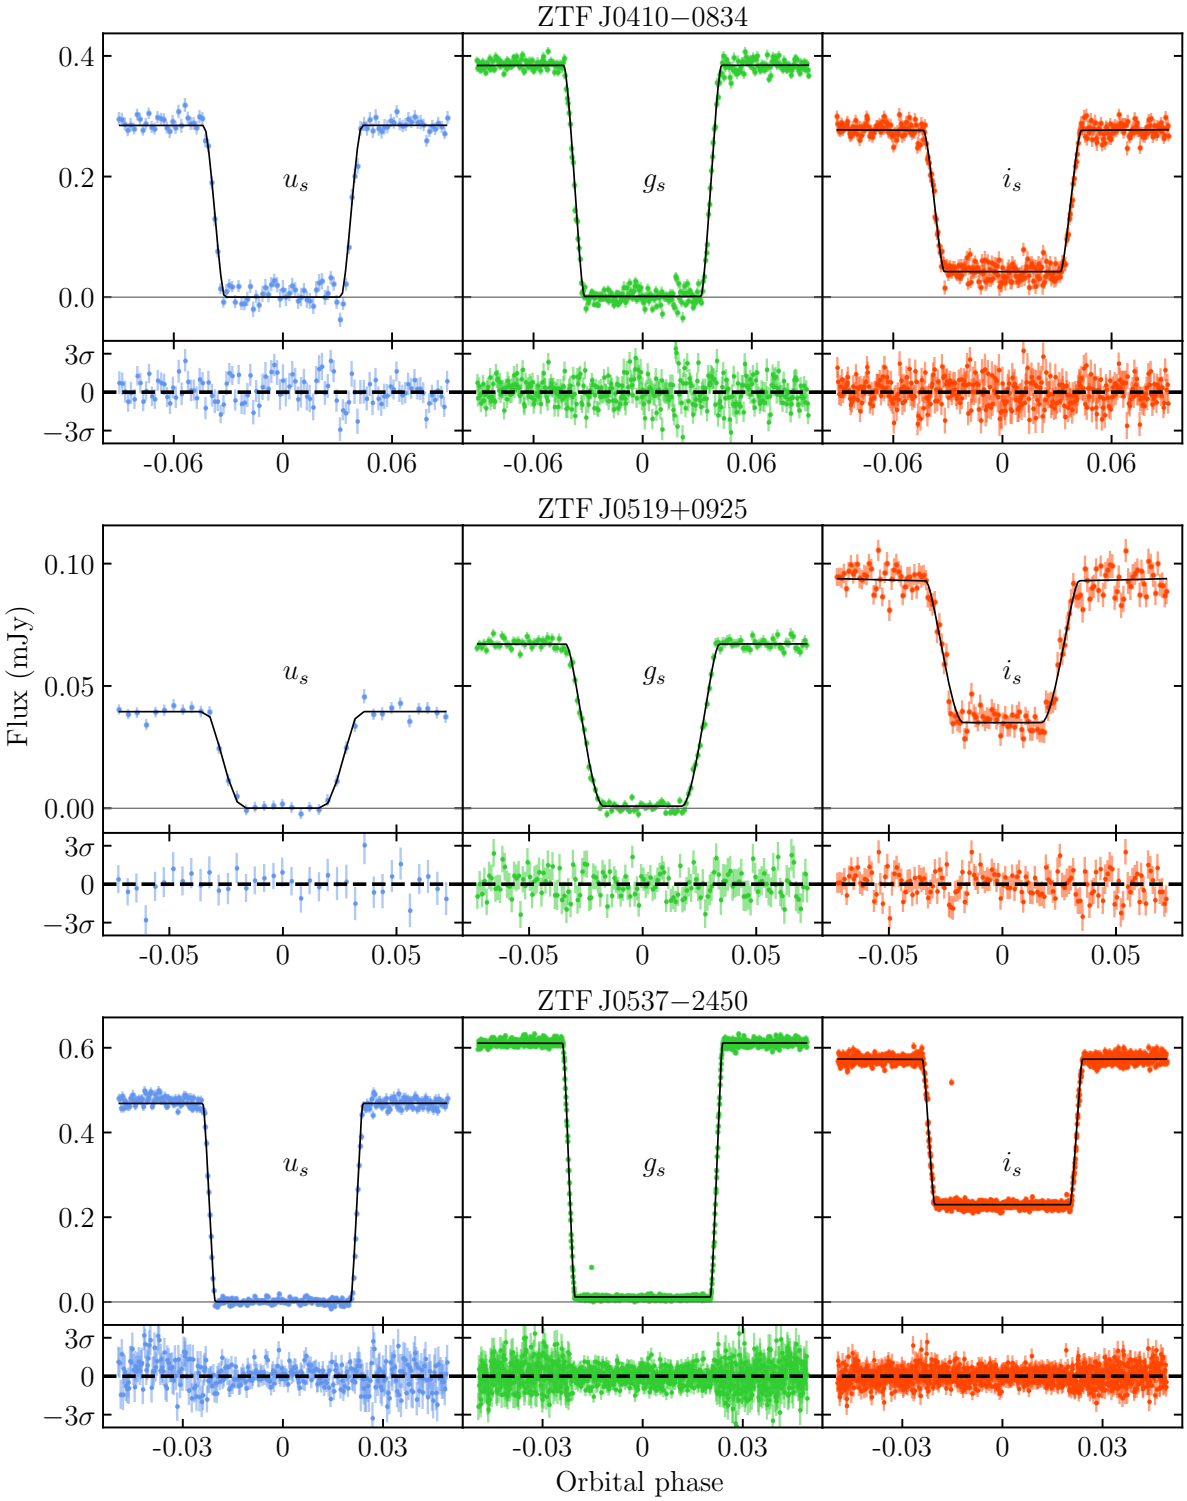

Figure B.1: ULTRACAM $u_s g_s i_s$ eclipse photometry with the best-fit light curve model over-plotted in black. A flux of zero is shown by the thin grey line. Residuals from this best-fit model are displayed in the lower panels.

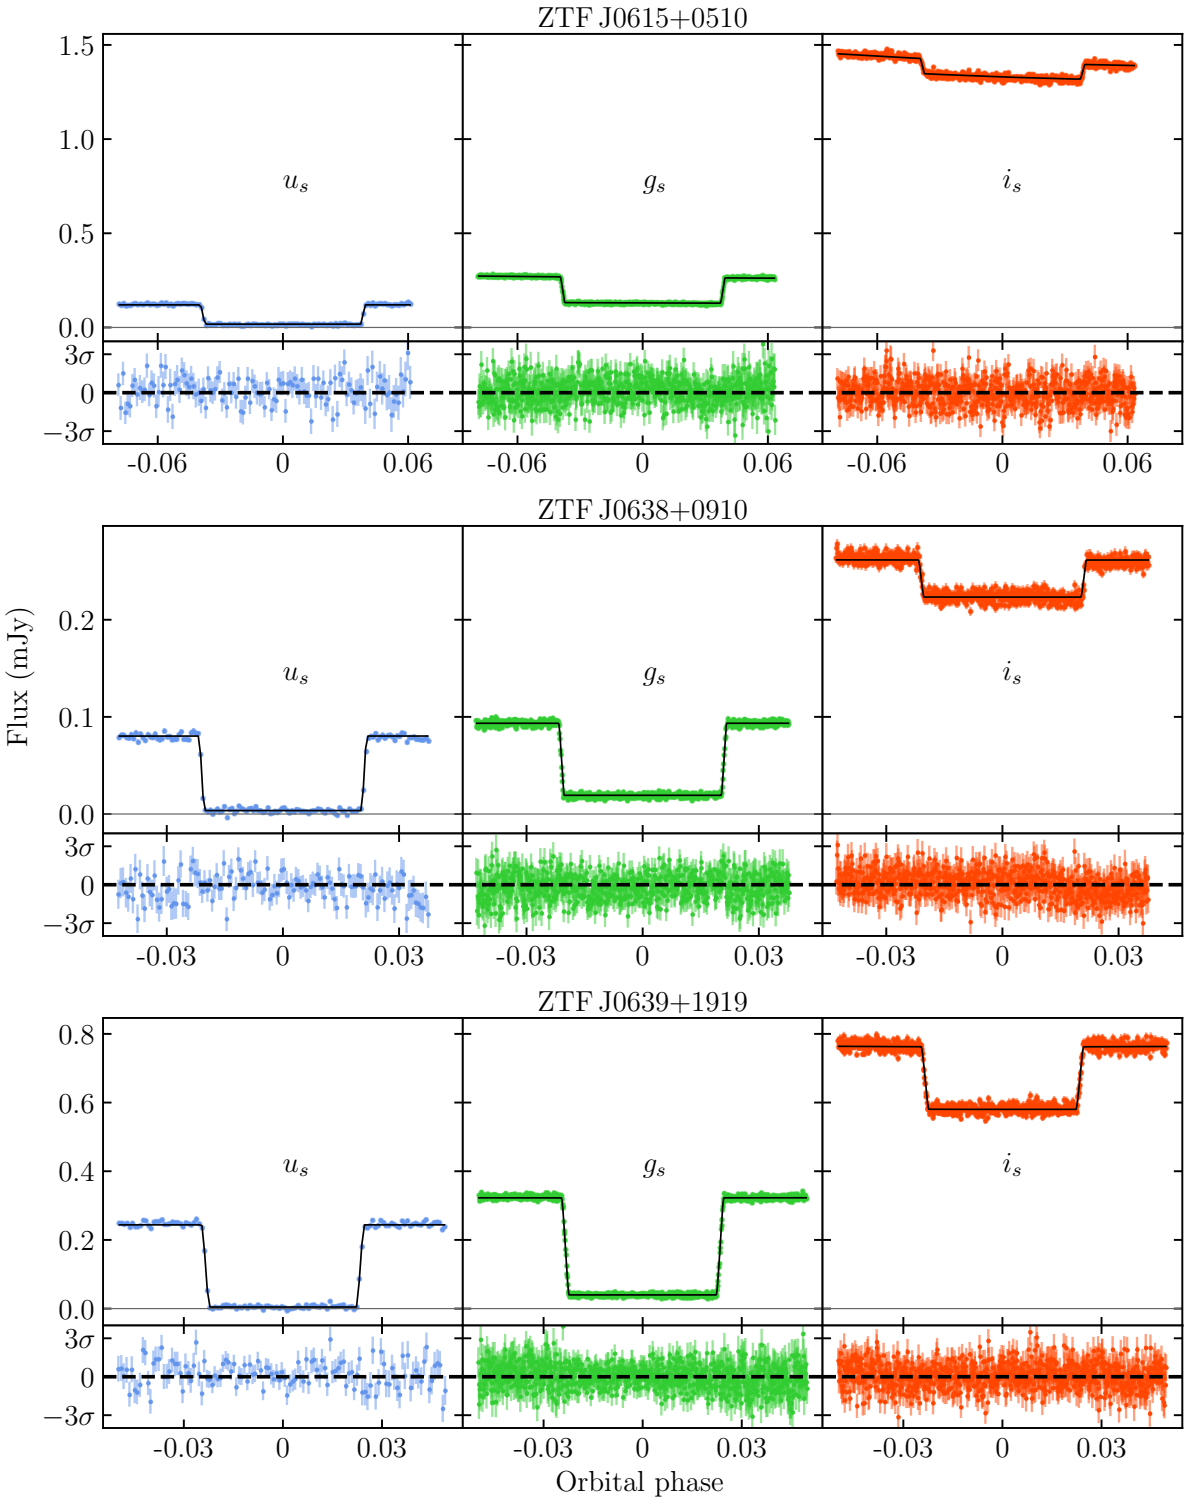

Figure B.2: As in Figure B.1.

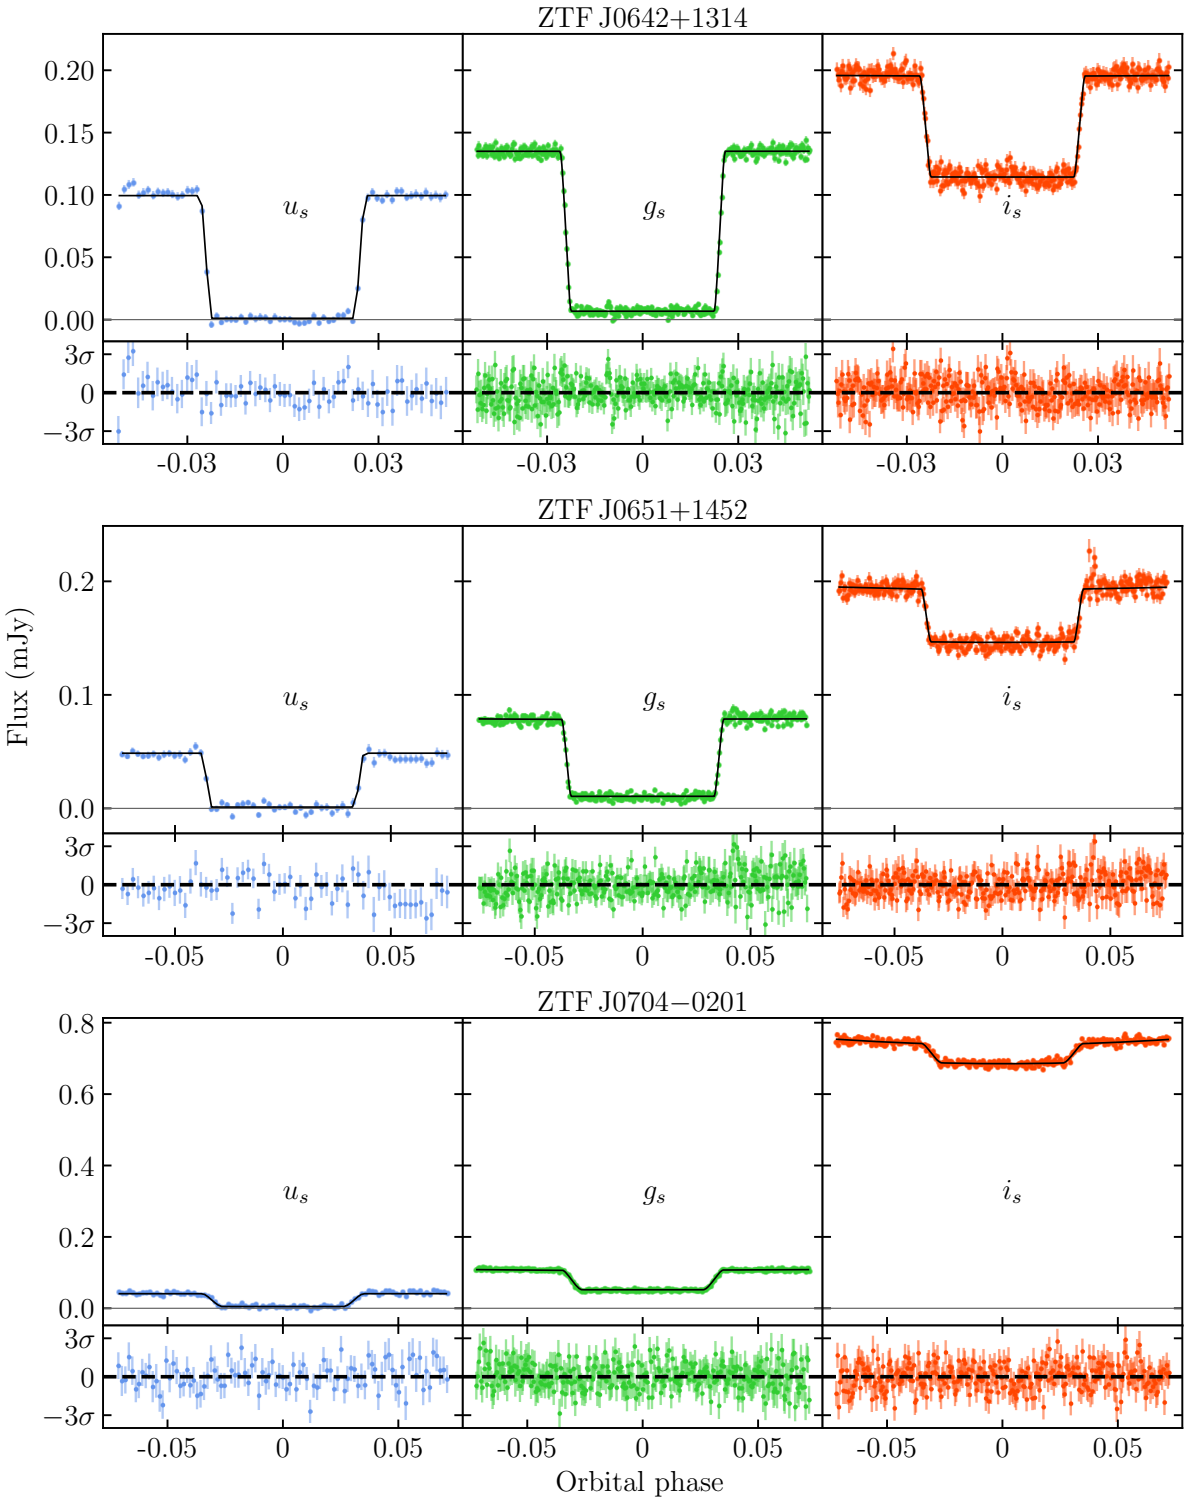

Figure B.3: As in Figure B.1.

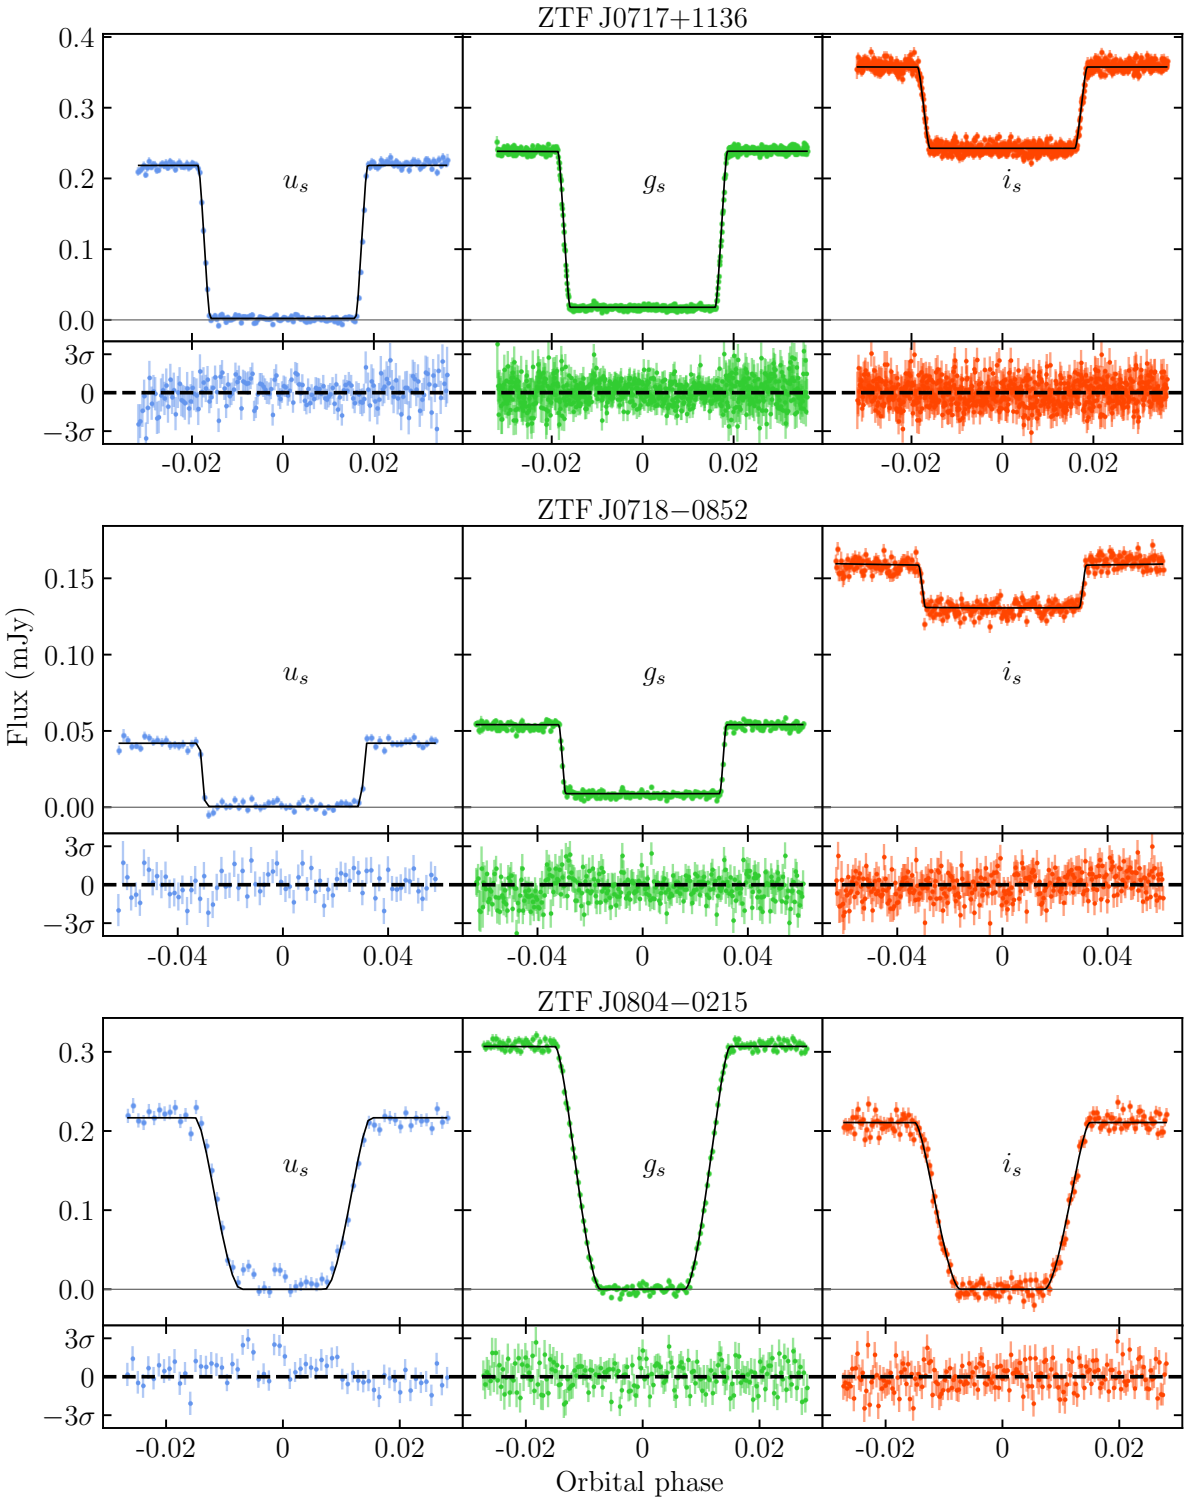

Figure B.4: As in Figure B.1.

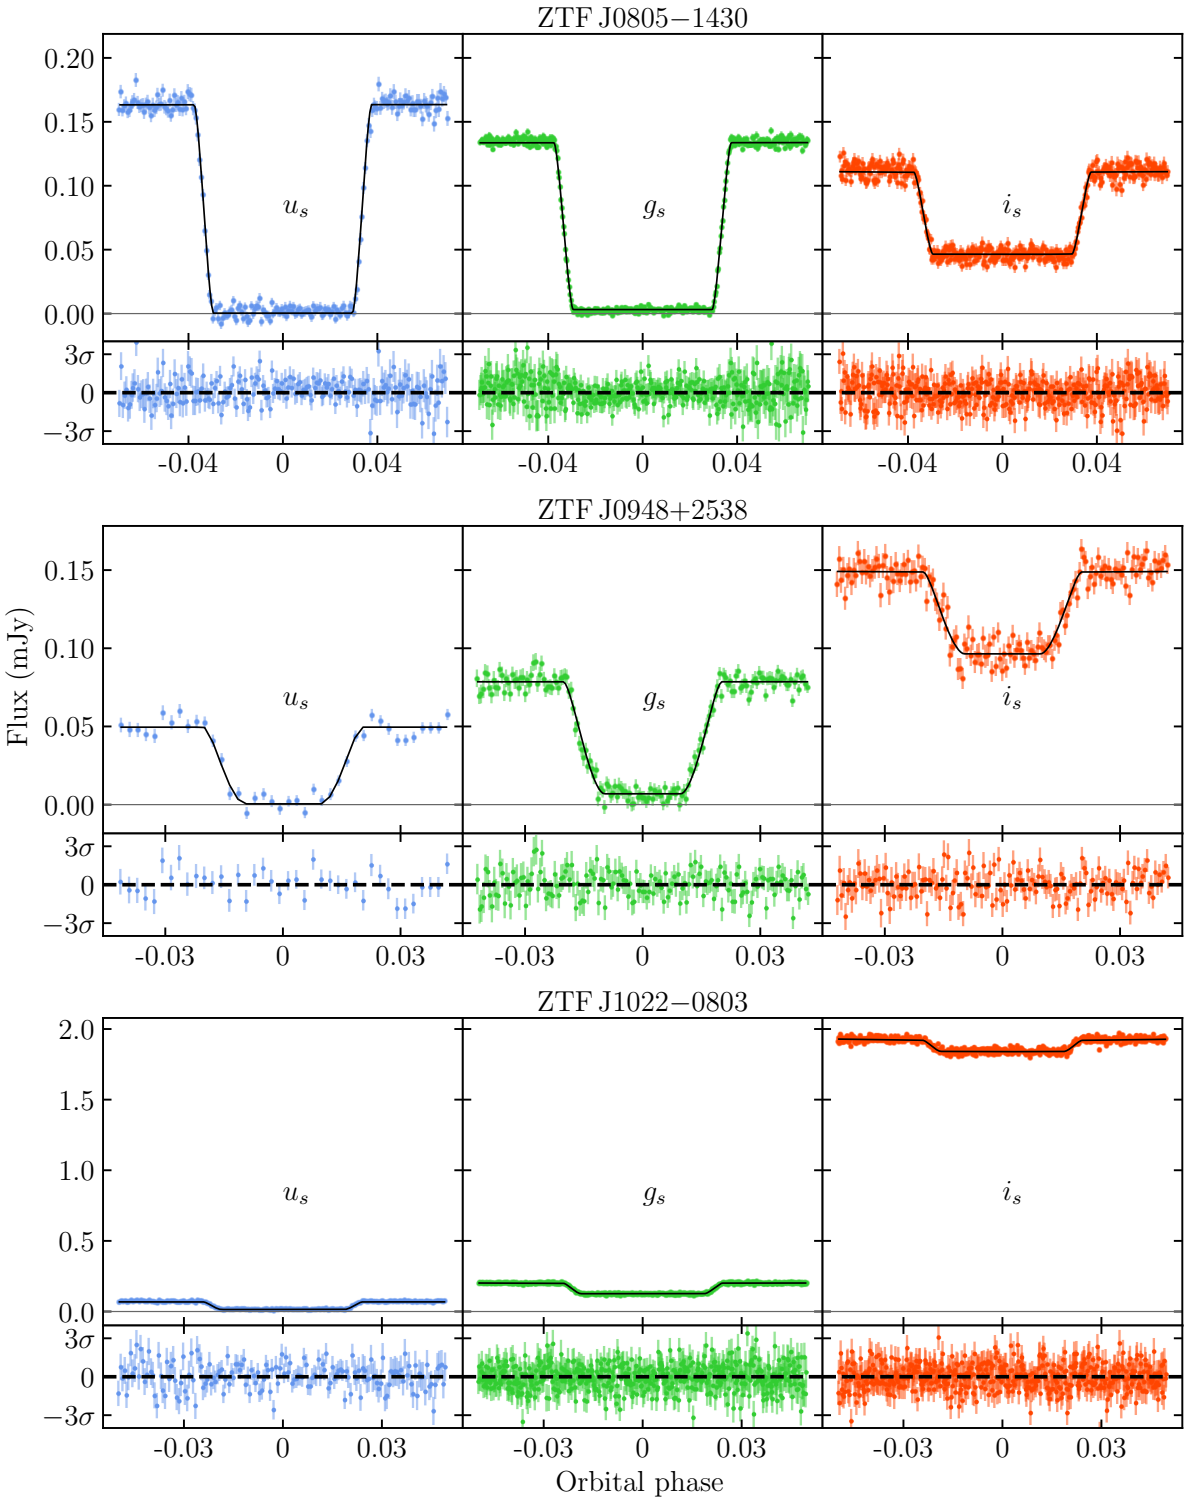

Figure B.5: As in Figure B.1.

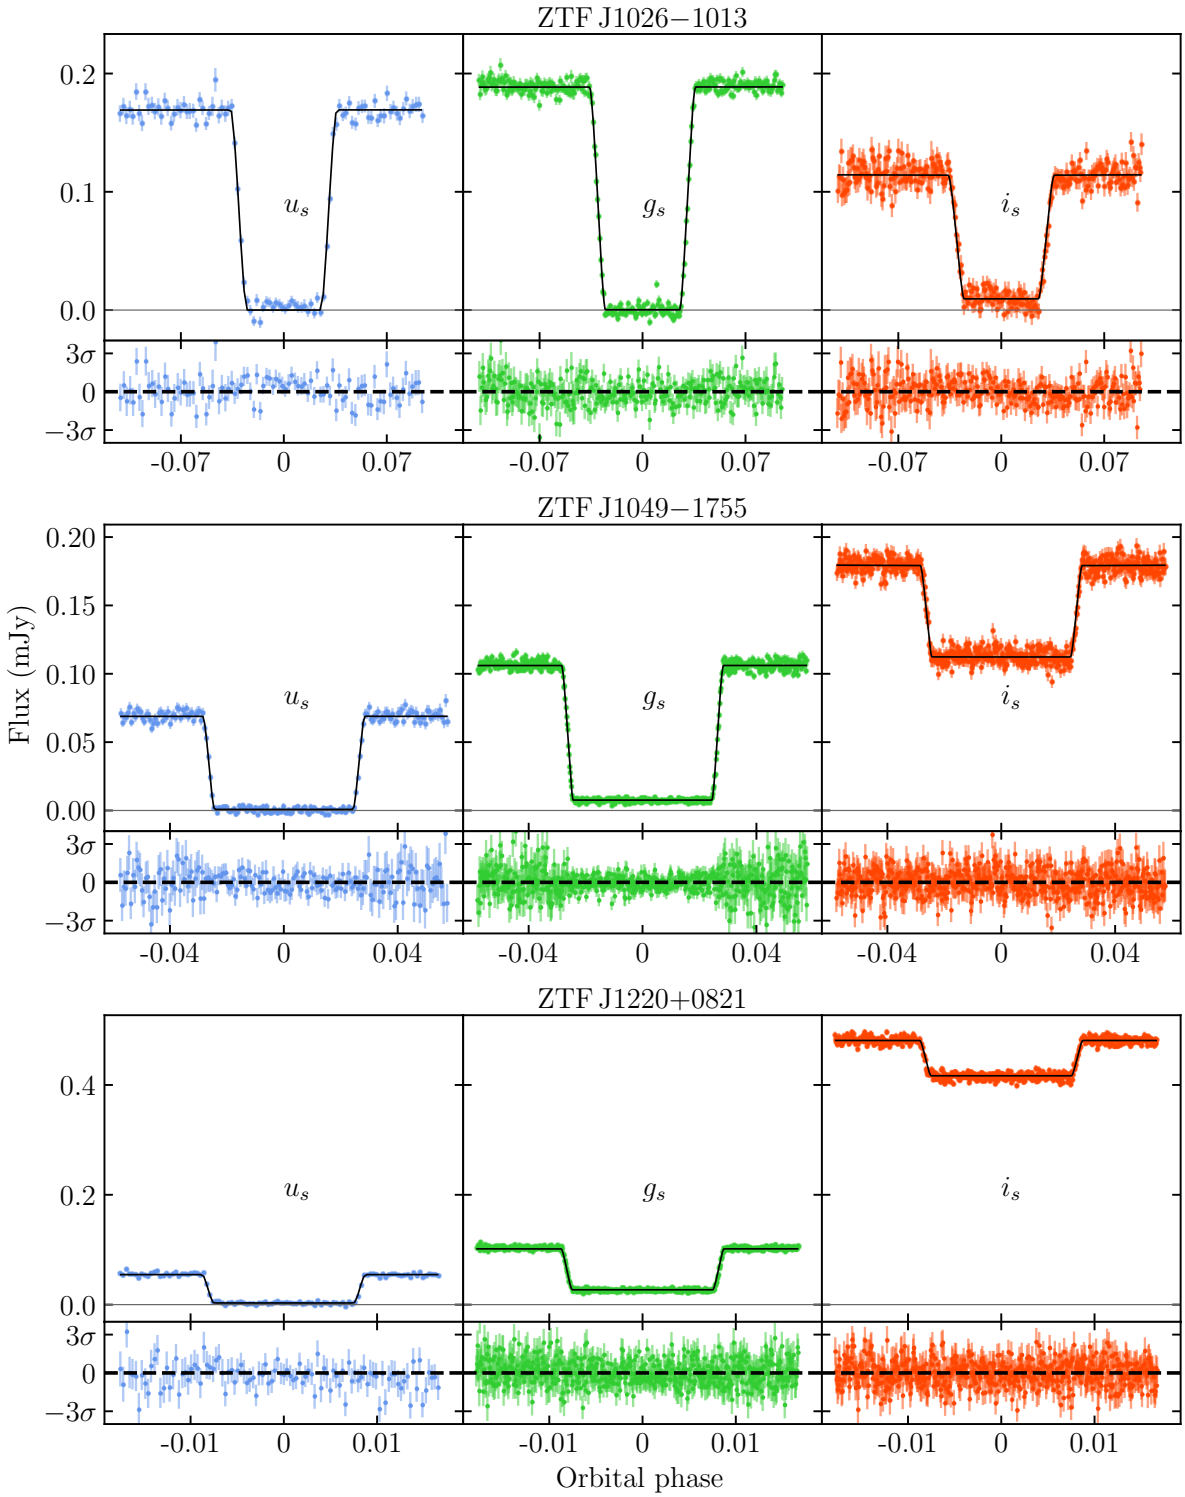

Figure B.6: As in Figure B.1.

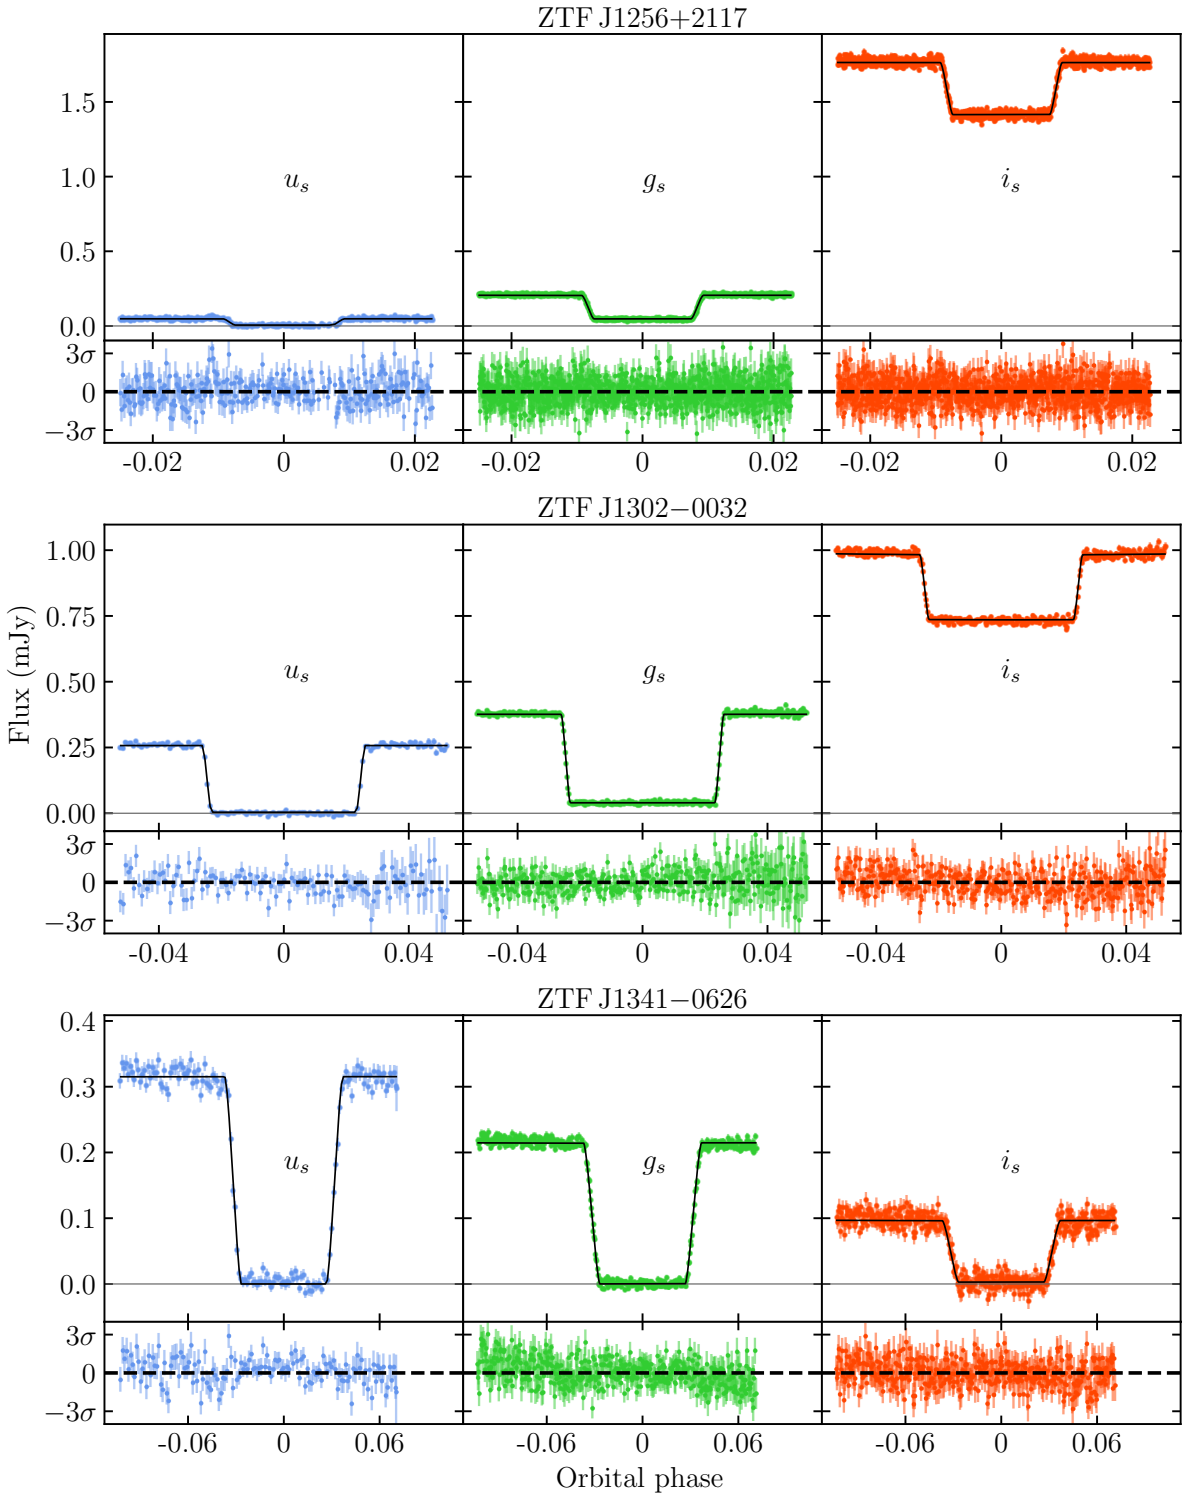

Figure B.7: As in Figure B.1.

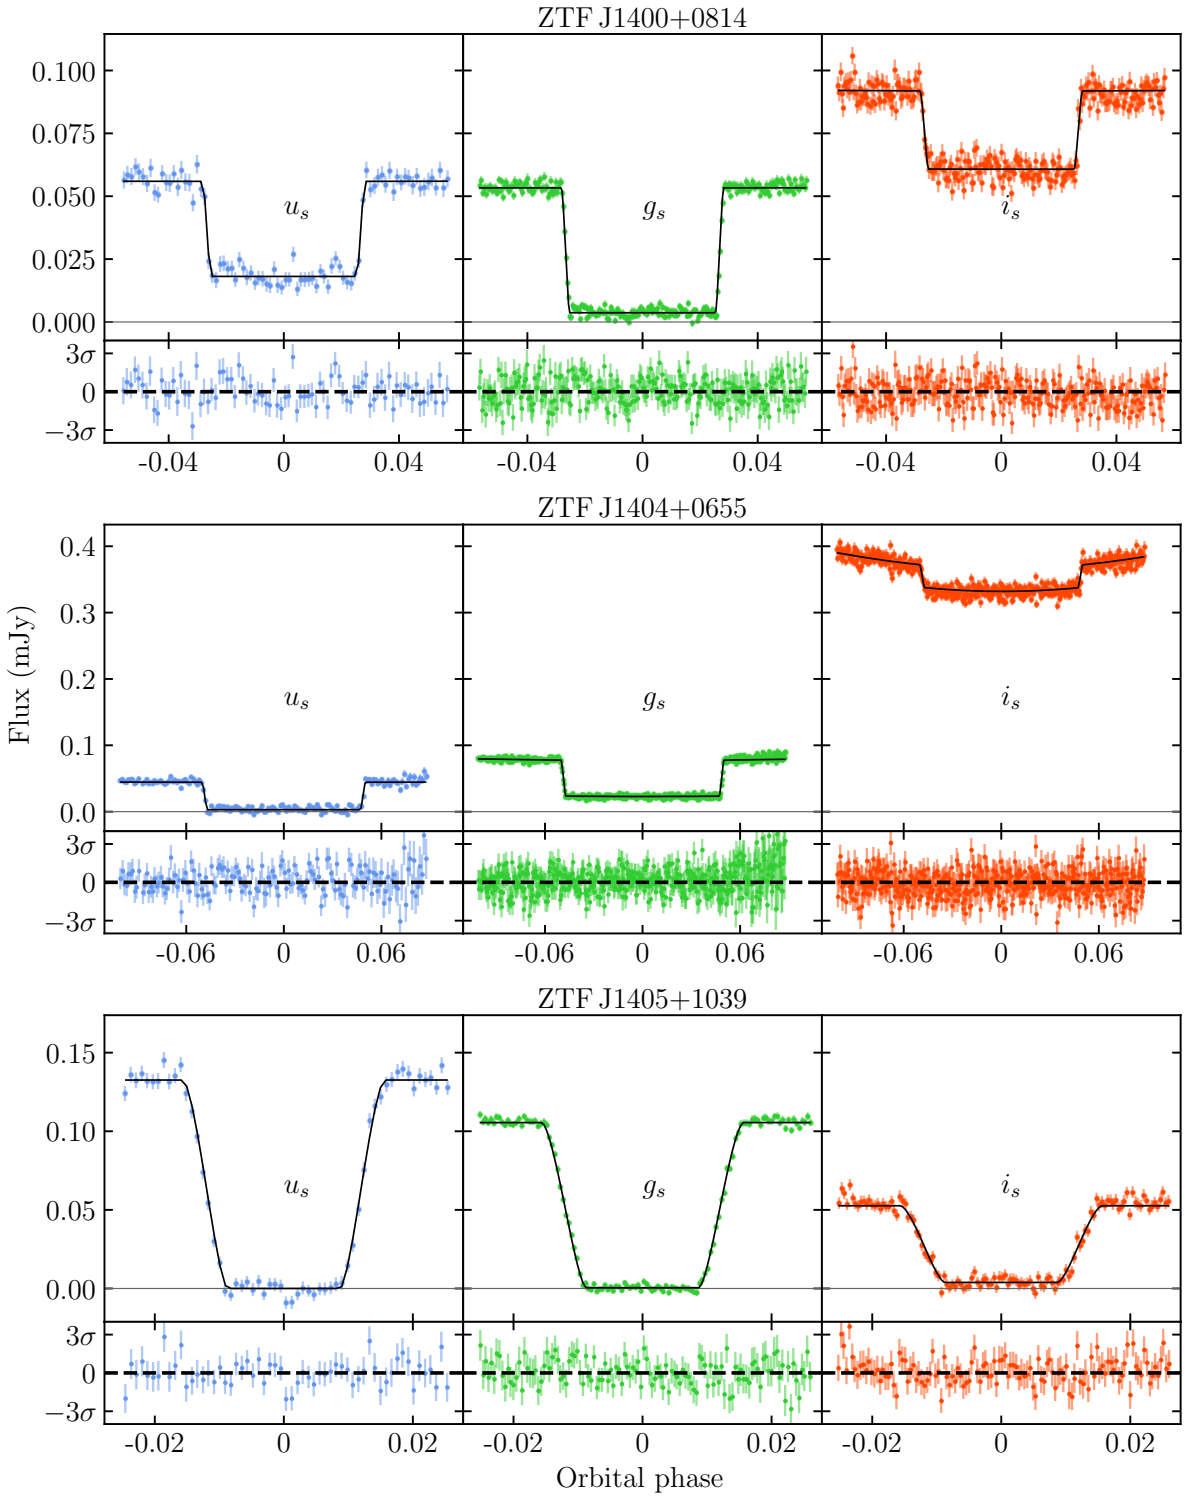

Figure B.8: As in Figure B.1.

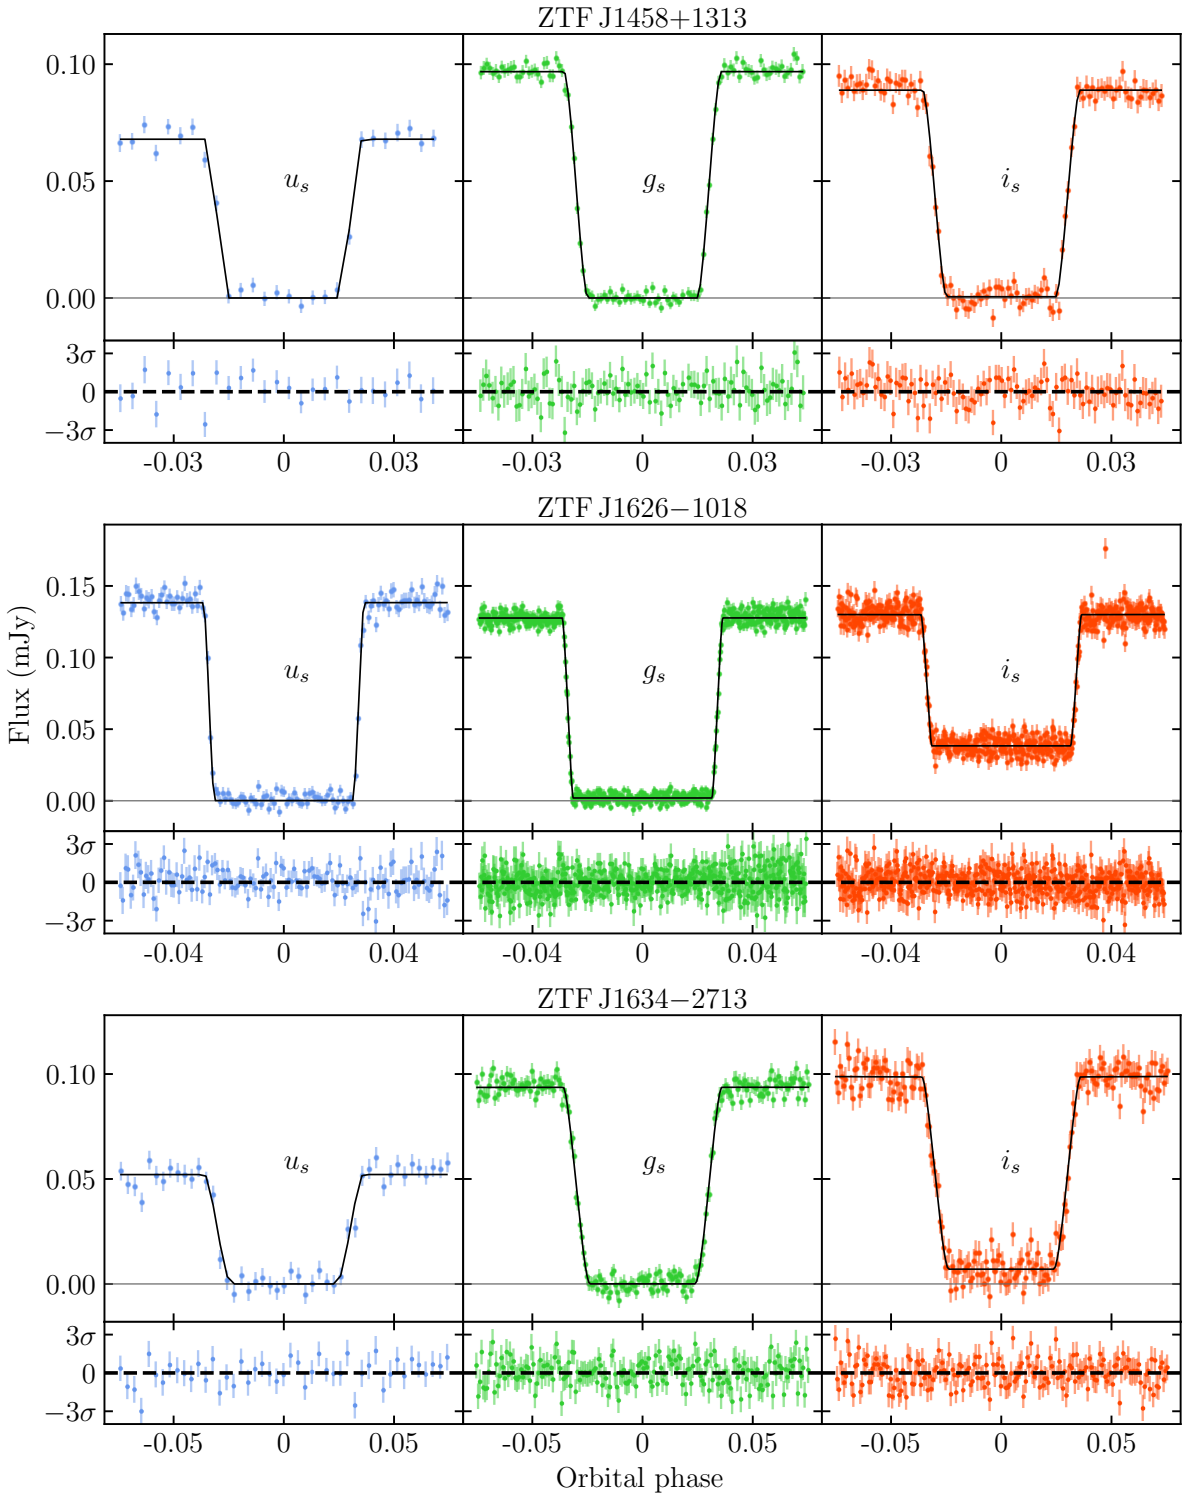

Figure B.9: As in Figure B.1.

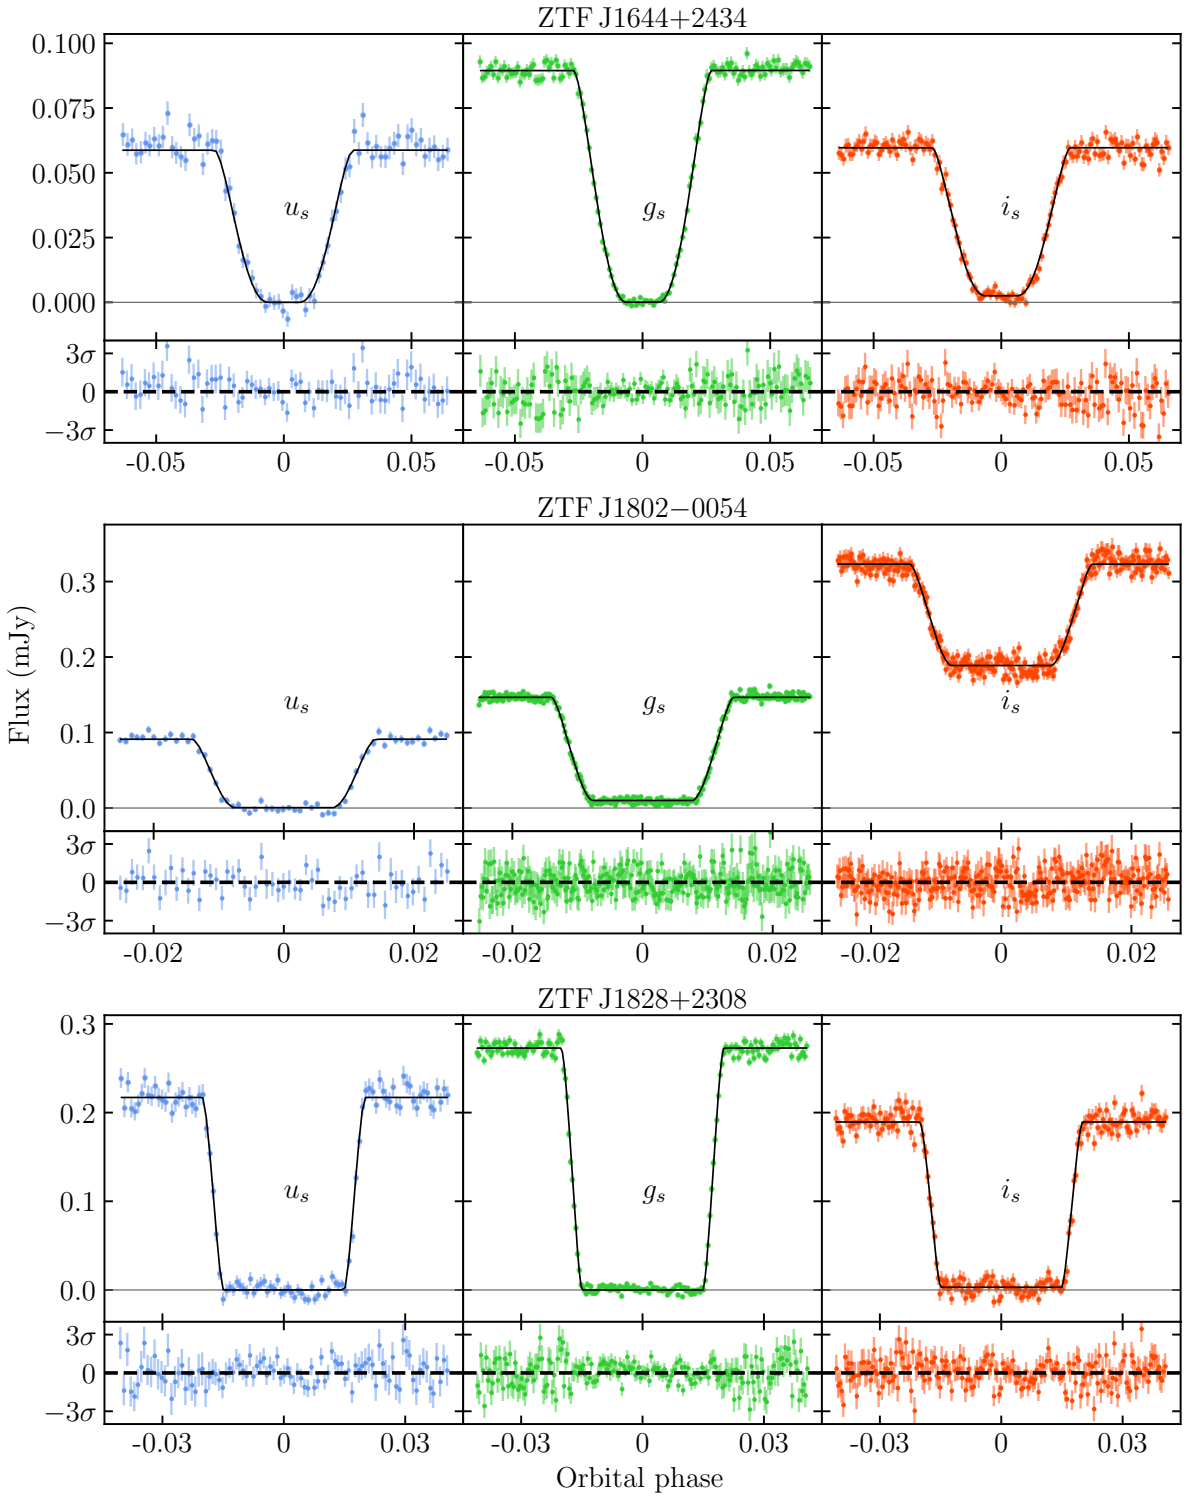

Figure B.10: As in Figure B.1.

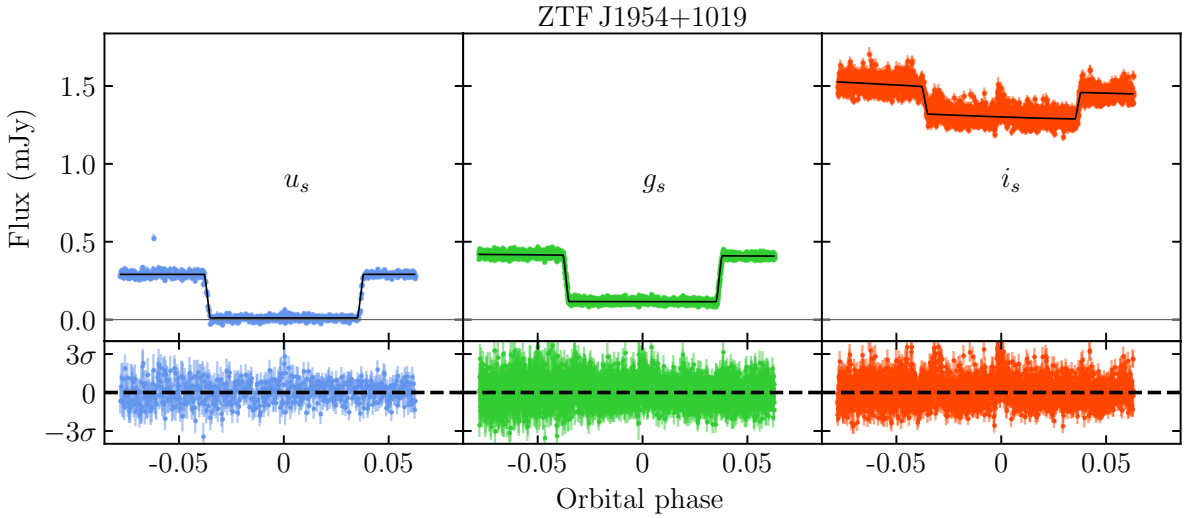

Figure B.11: As in Figure B.1.

Appendix C

SEDs

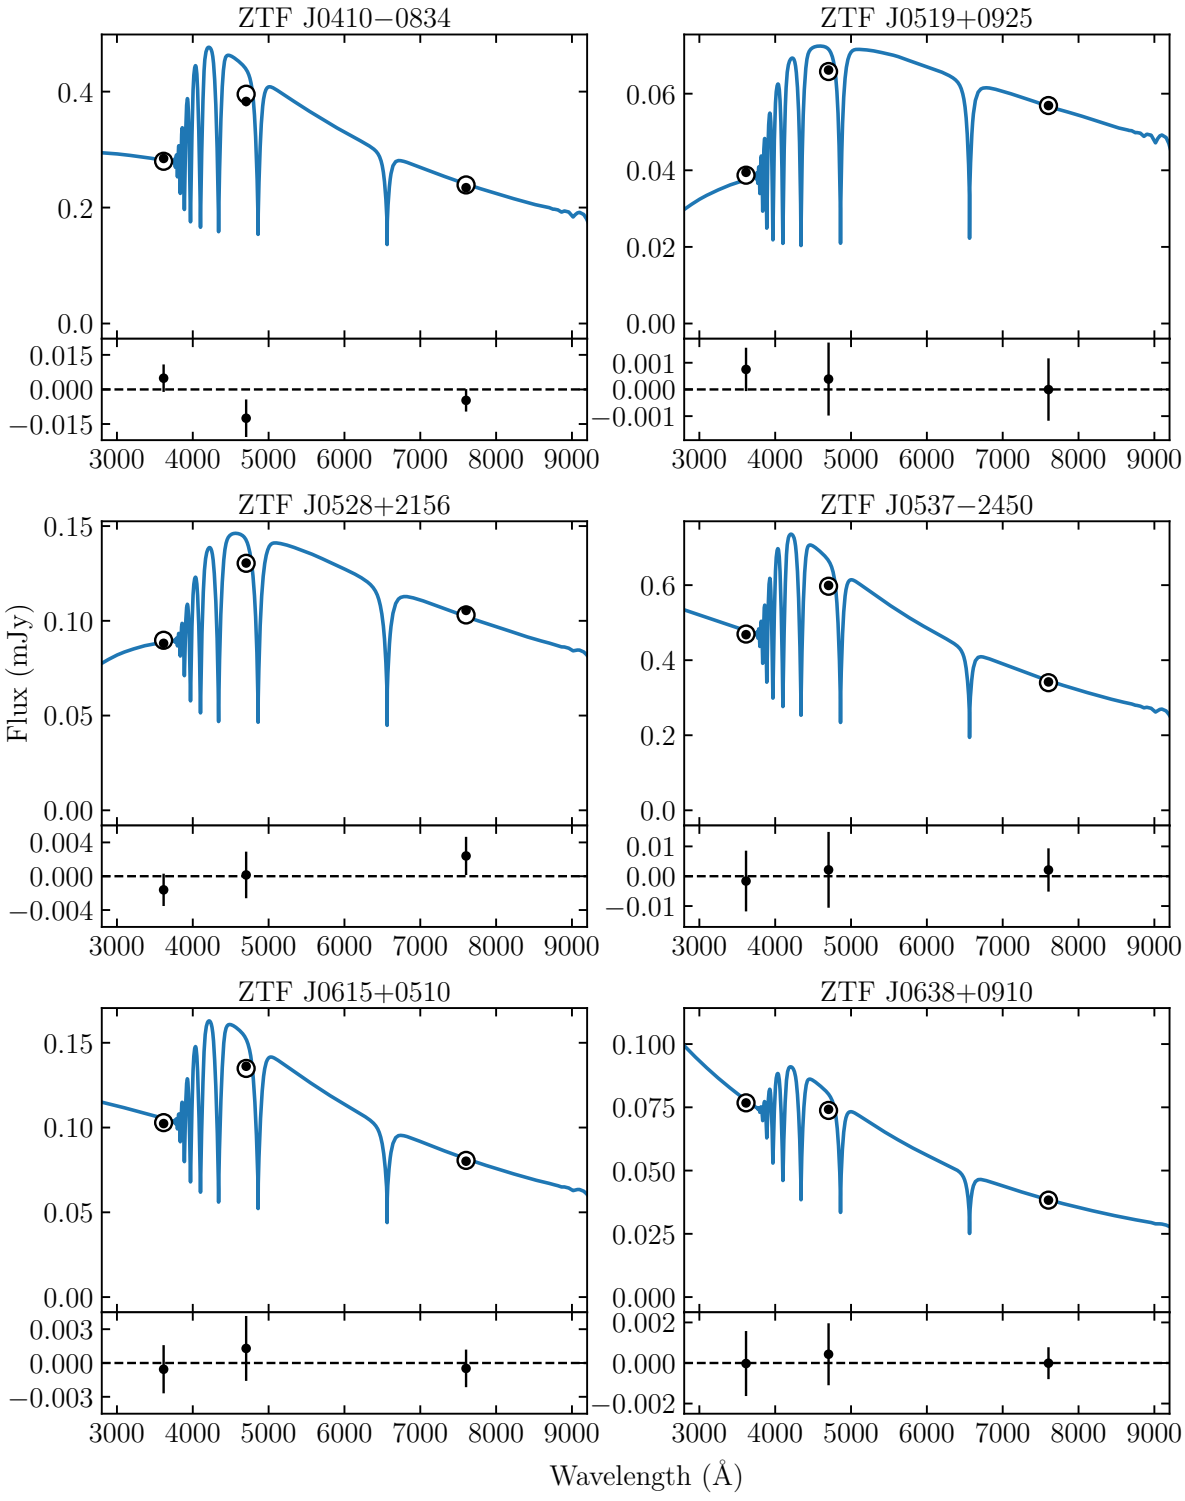

Figure C.1: The SED of the WD (black errorbars) measured from the eclipse photometry shown alongside the synthetic SED for the best-fit WD model (open circles) with the residuals displayed in the panels below. The spectrum of the WD is plotted in blue (Koester, 2010).

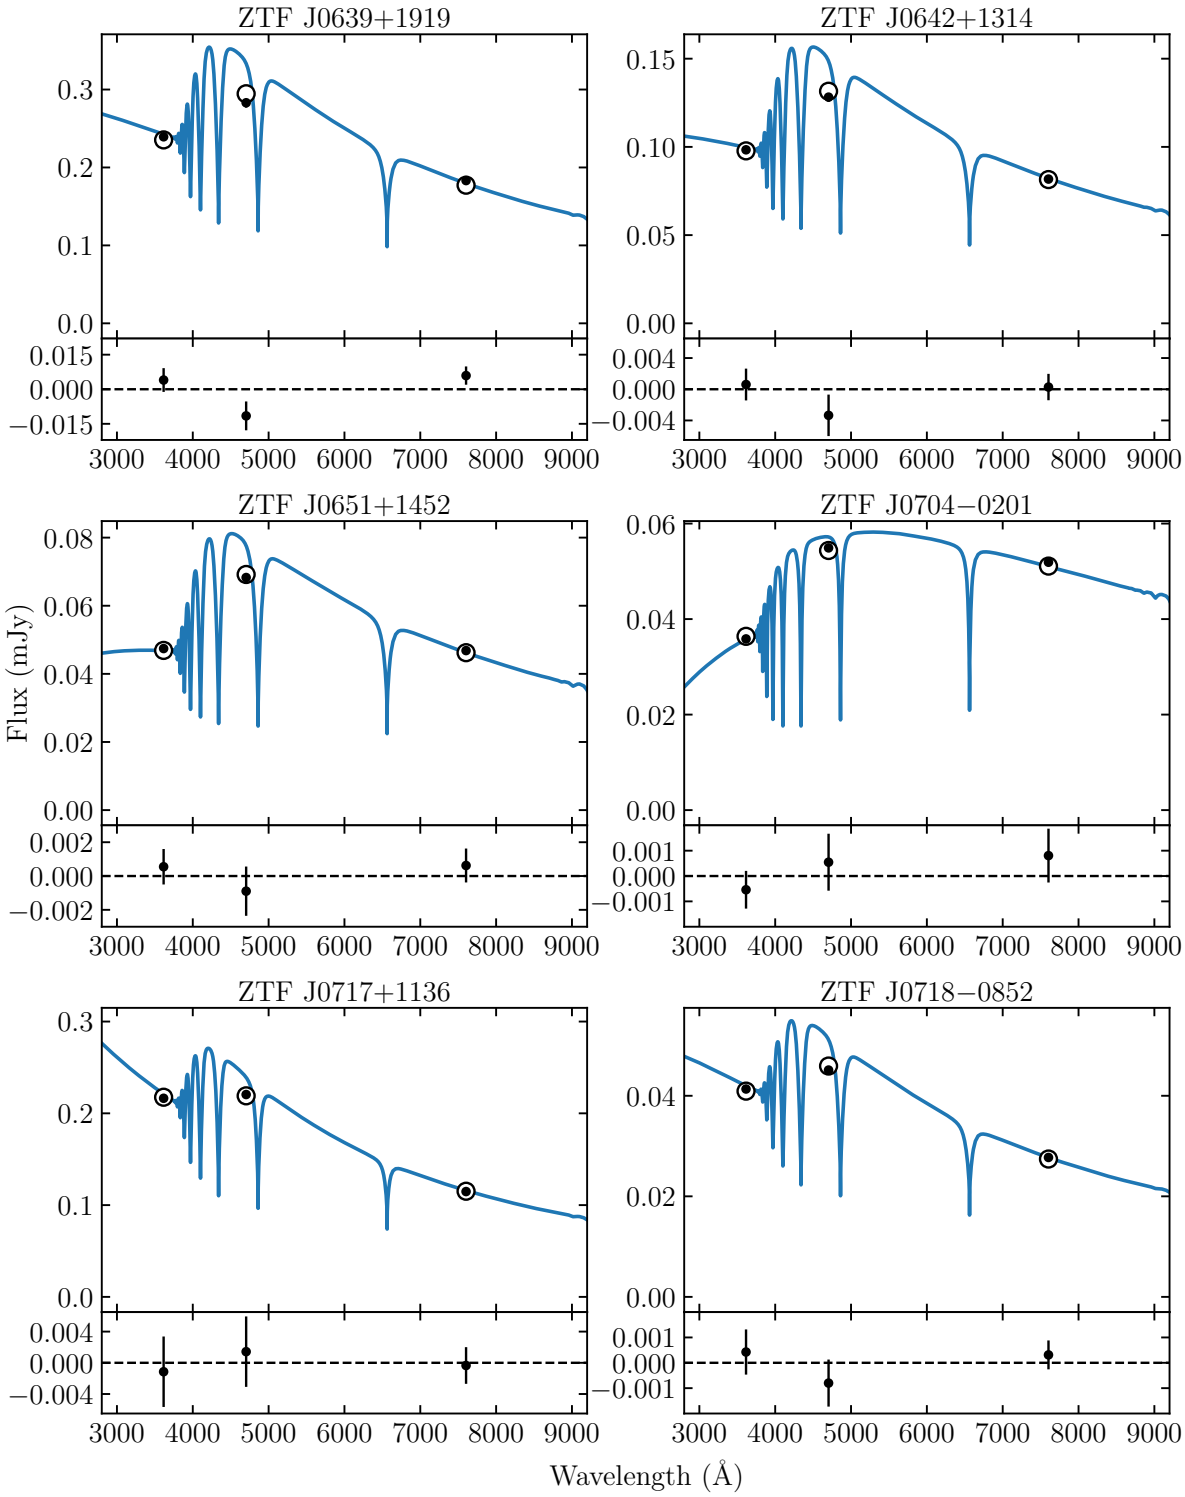

Figure C.2: As in Figure C.1.

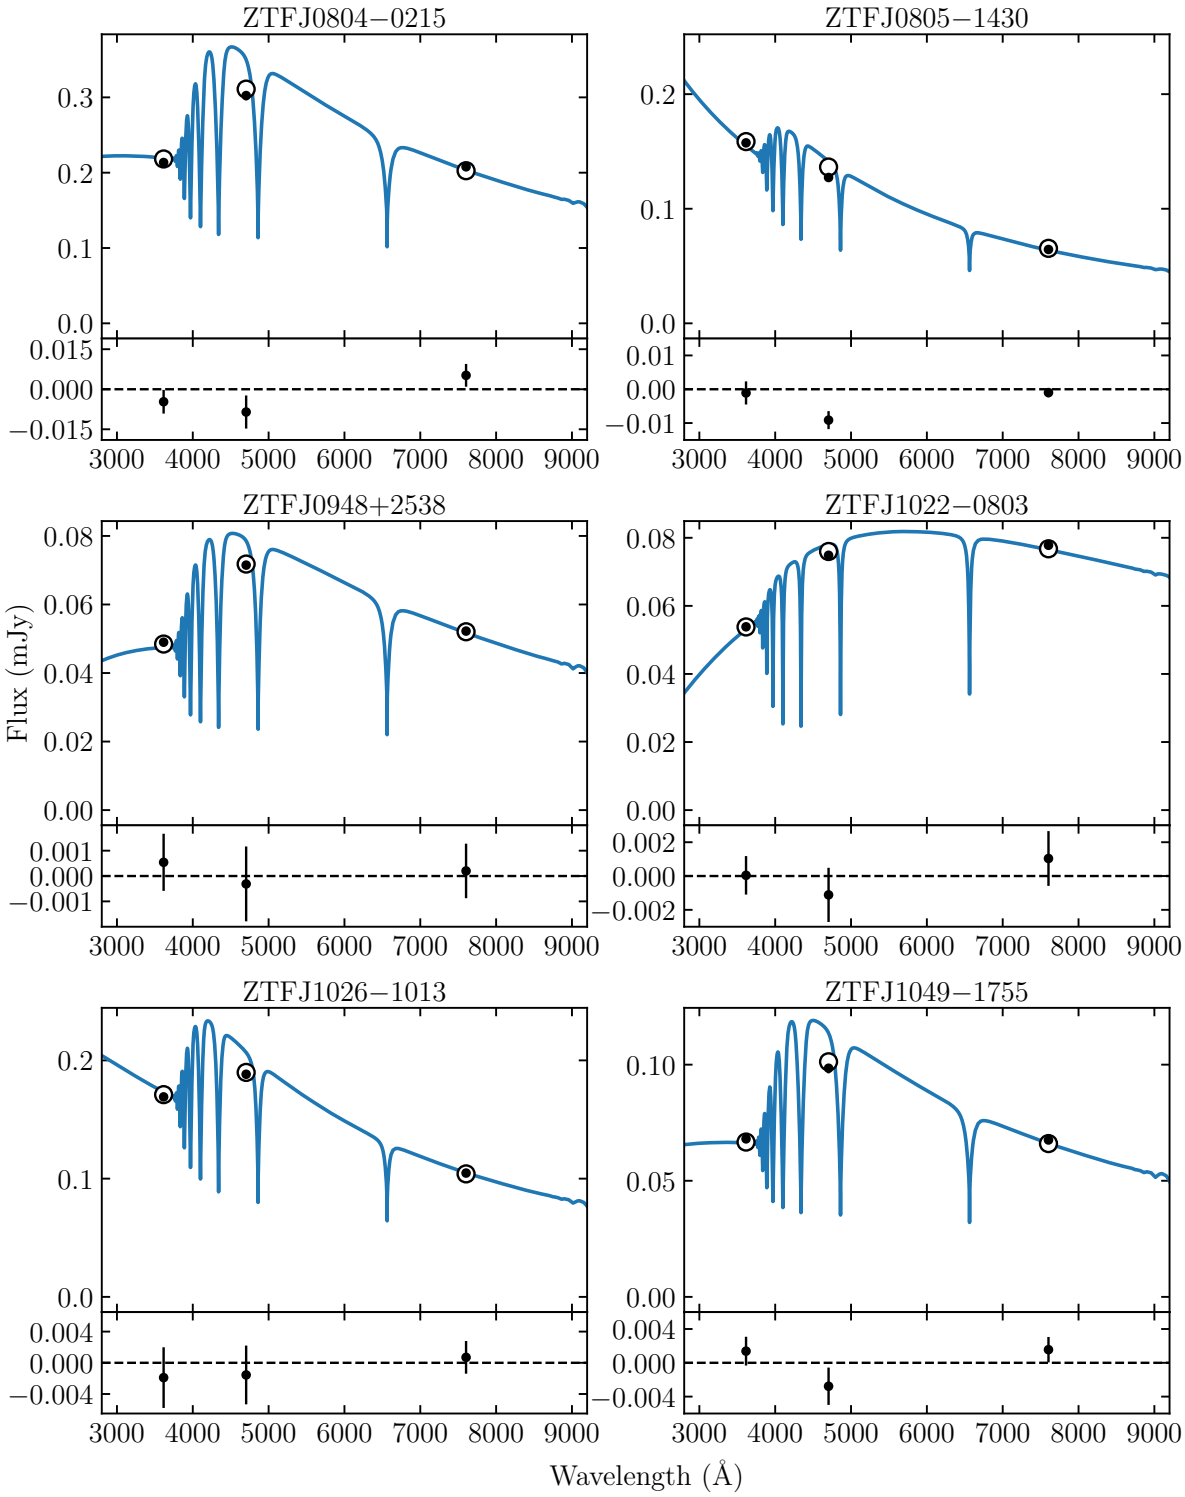

Figure C.3: As in Figure C.1.

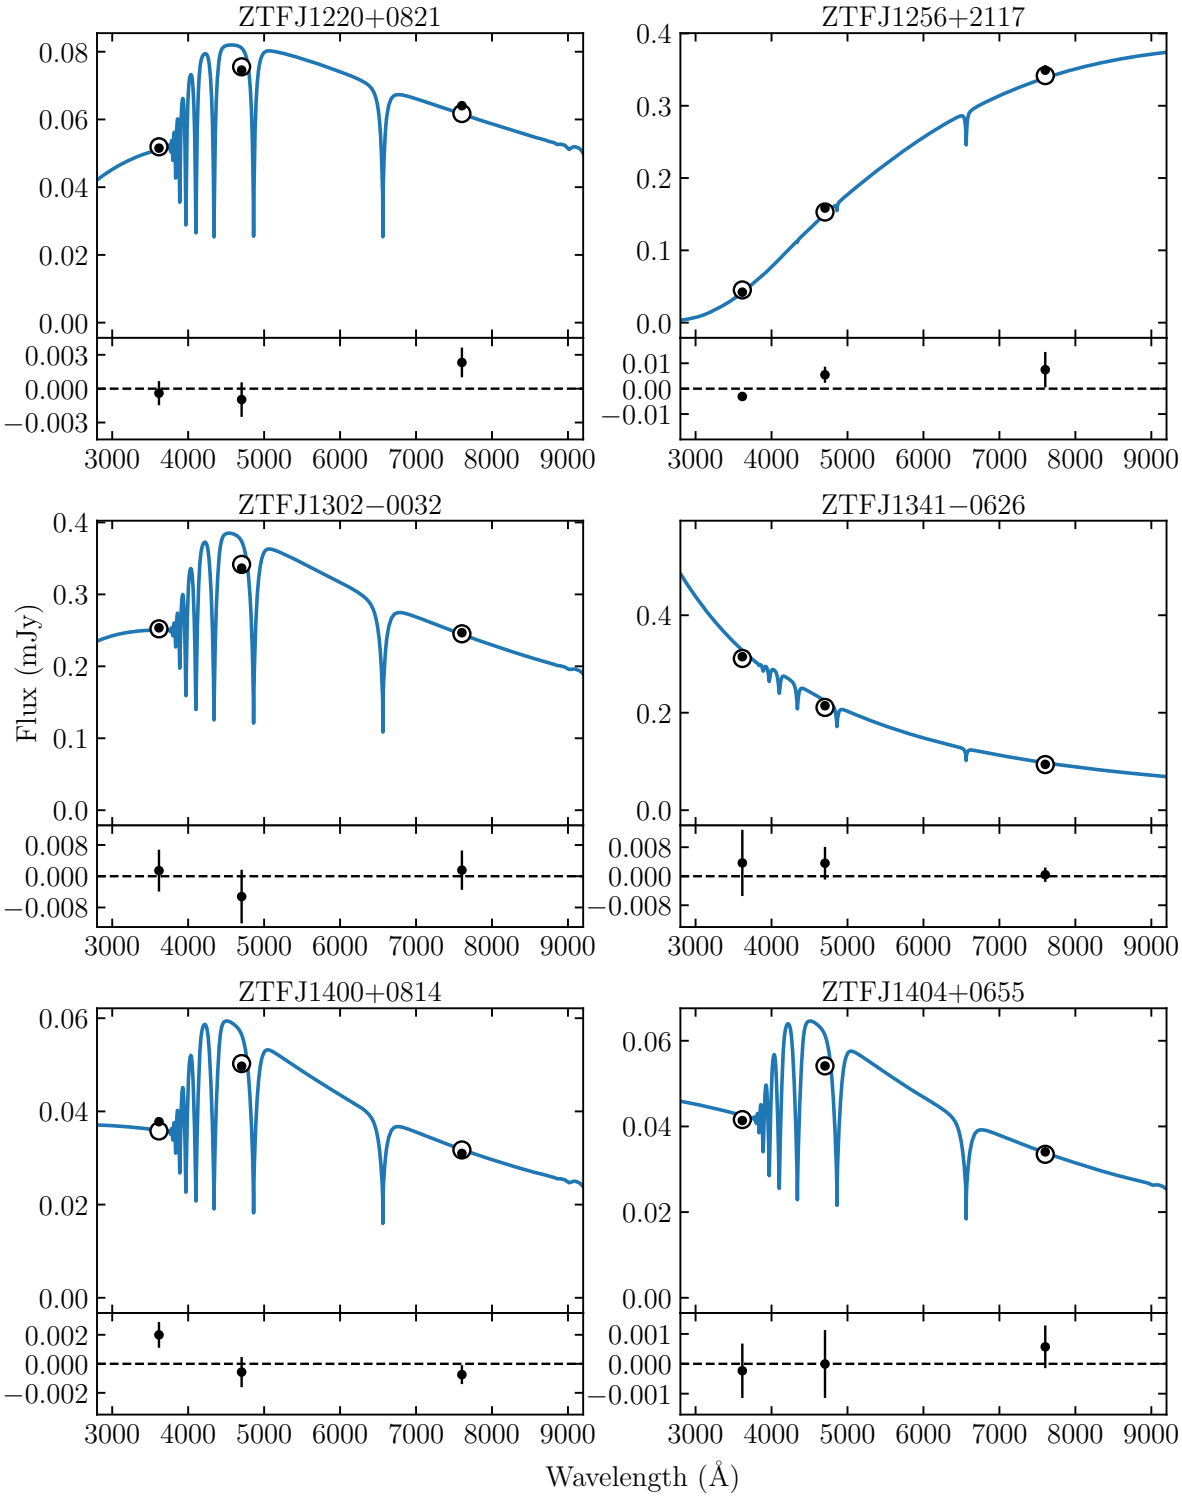

Figure C.4: As in Figure C.1.

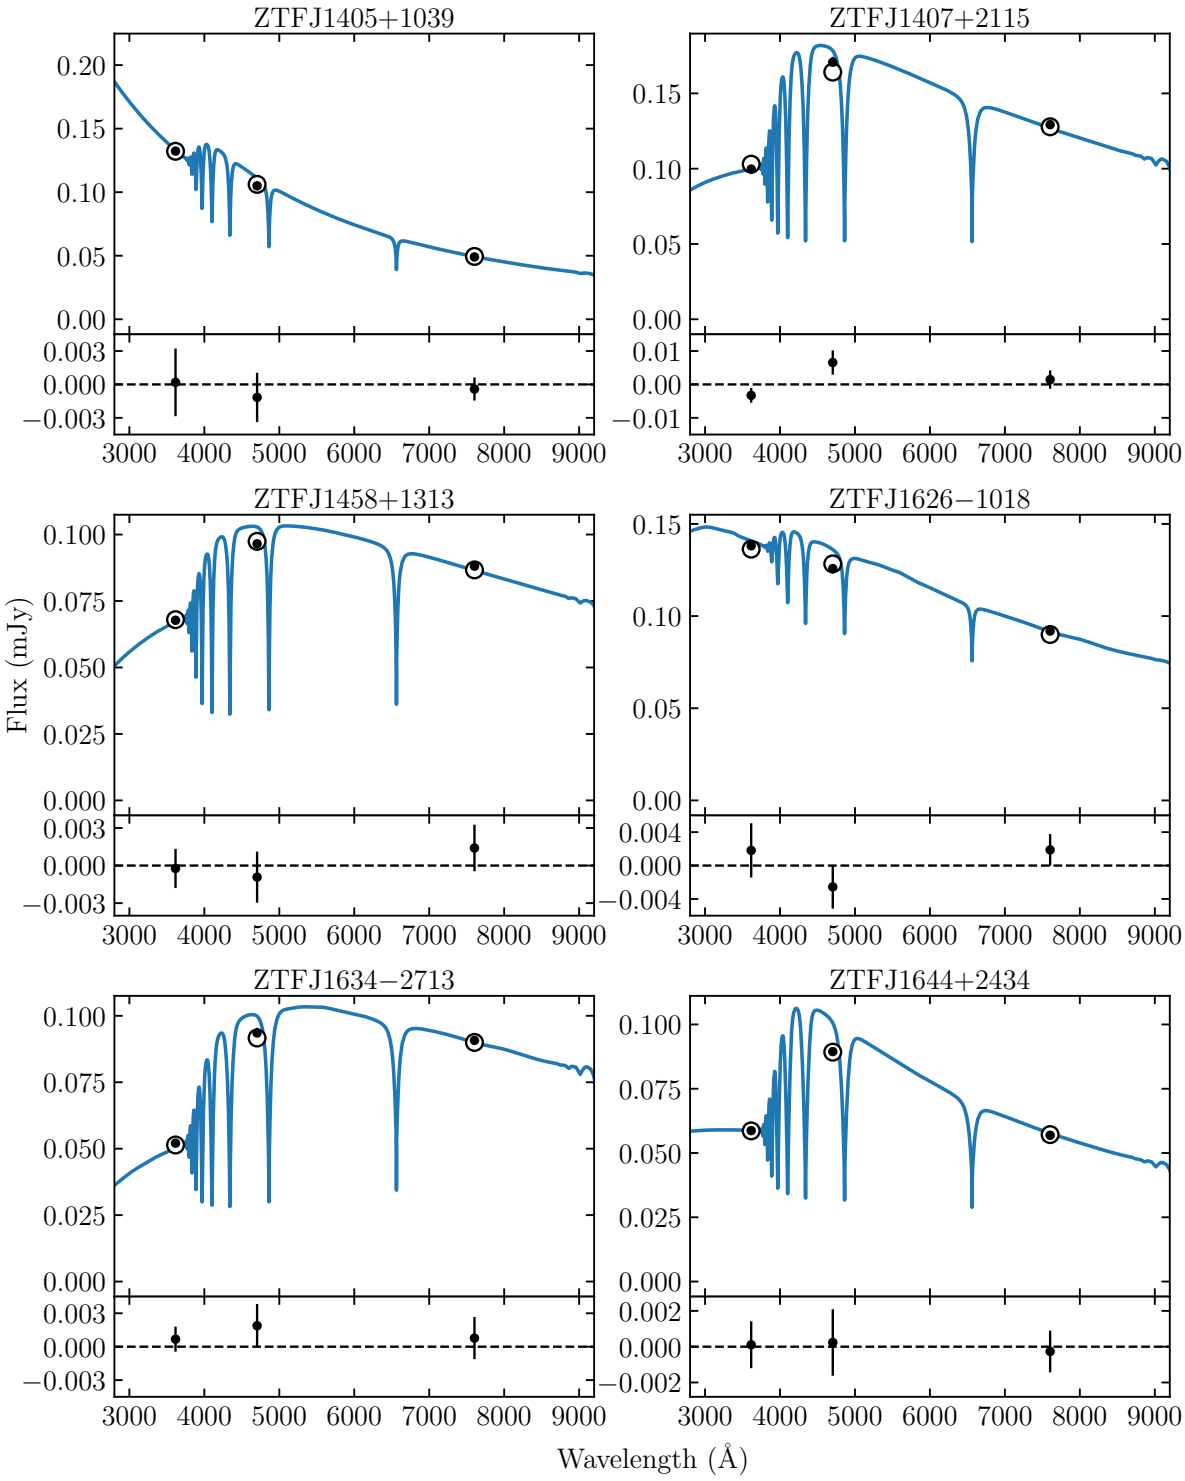

Figure C.5: As in Figure C.1.

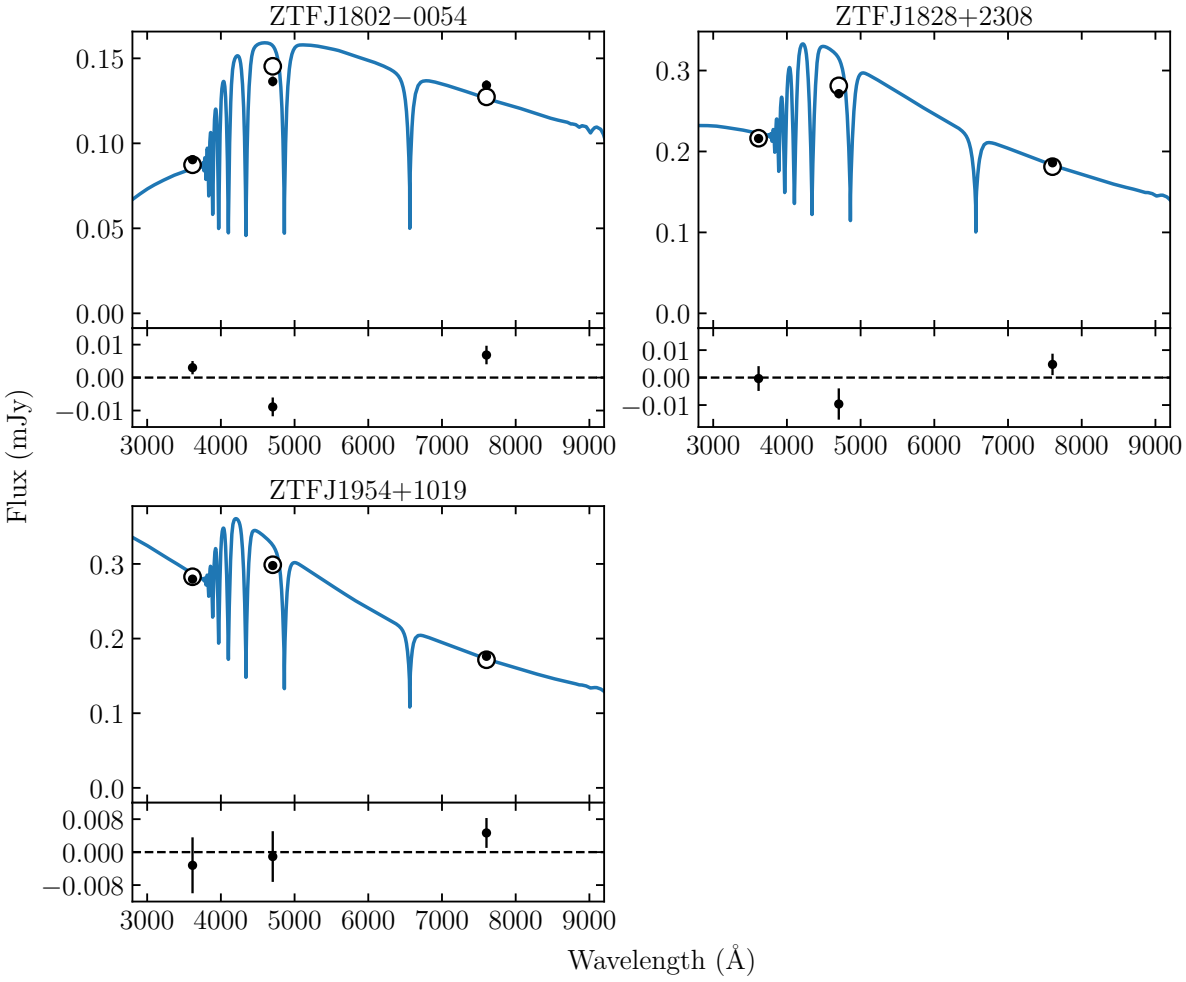

Figure C.6: As in Figure C.1.

Appendix D

Cornerplots

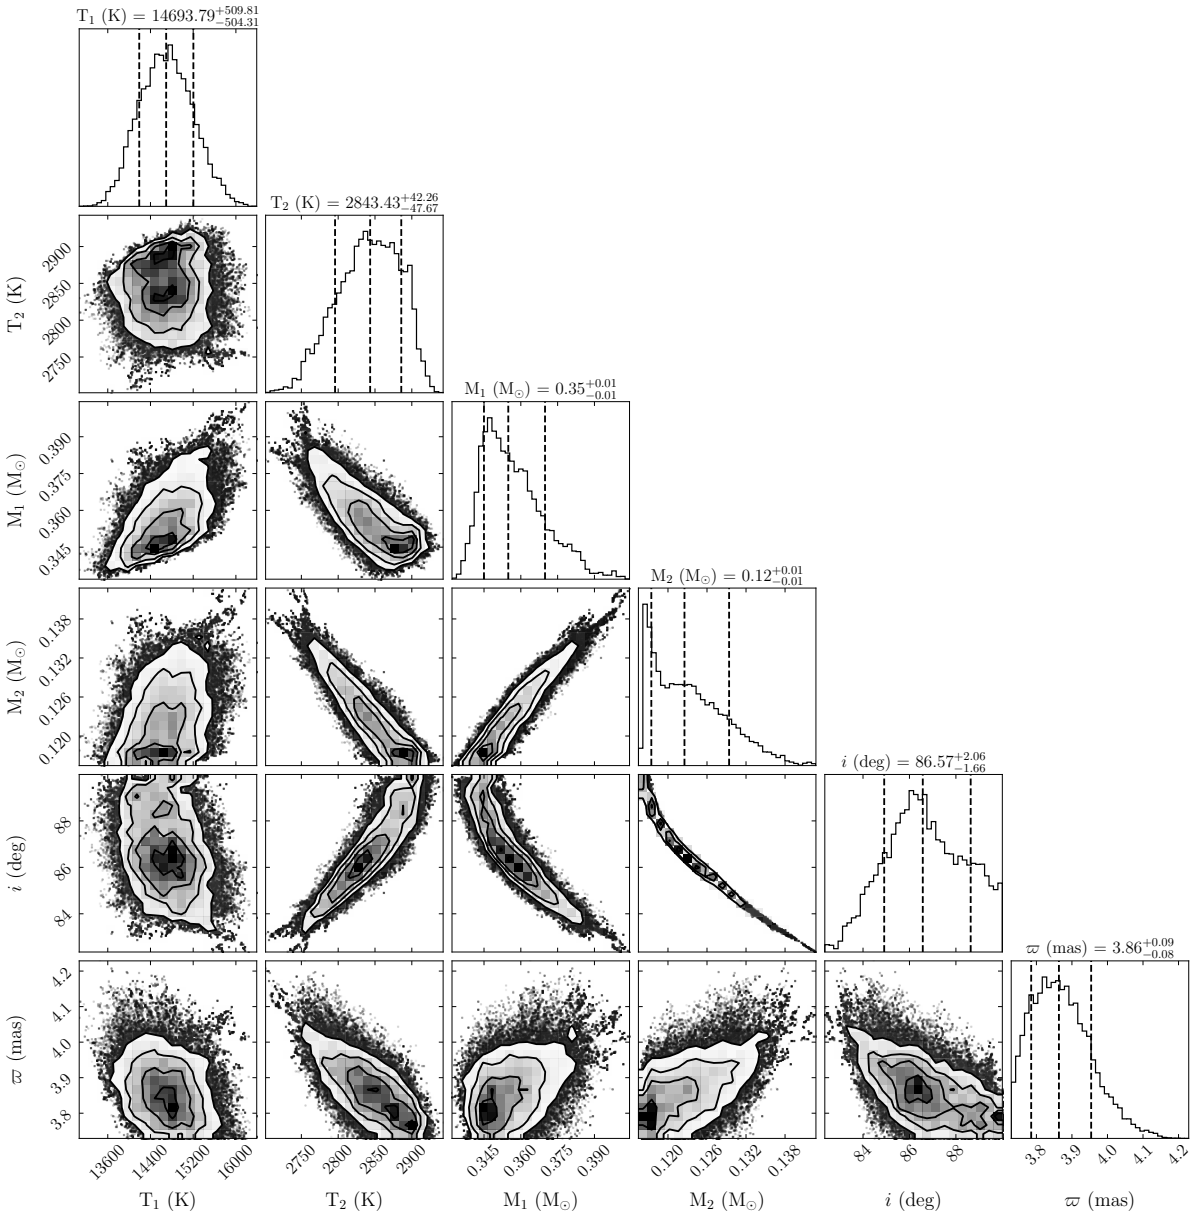

Figure D.1: 2D distributions of the MCMC samples for the final 3000 steps of the light curve fit to ZTF J0410−0834. Marginalised distributions for these parameters are shown above with the 16th, 50th, and 84th percentiles marked by the dashed lines. For clarity, only a subset of the fitted parameters are shown.

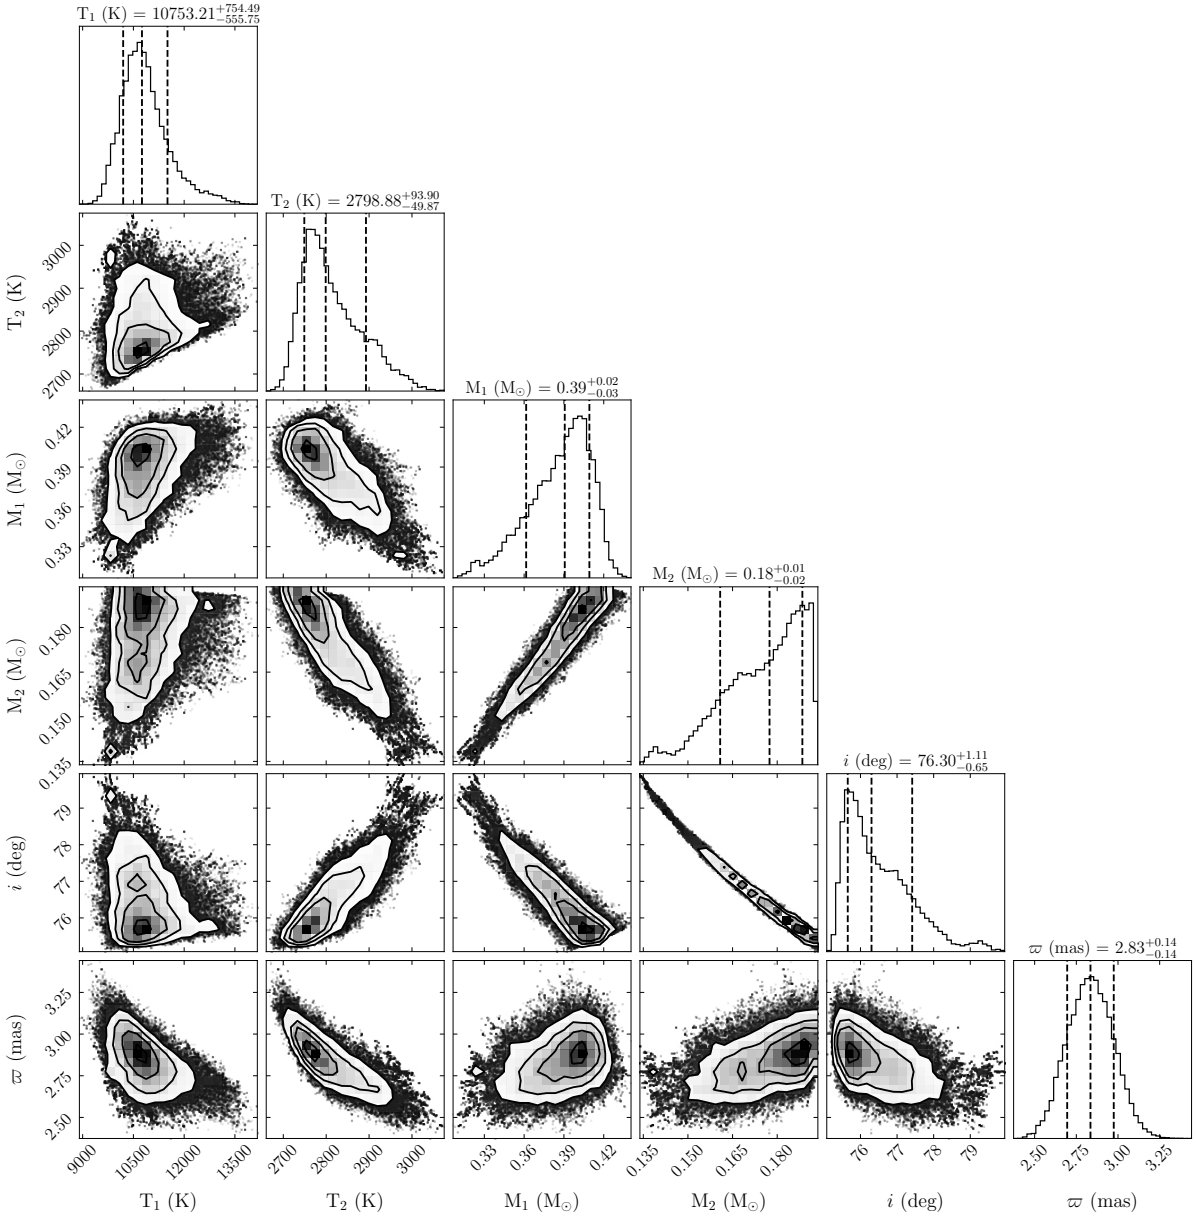

Figure D.2: As in Figure D.1 but for ZTF J0519+0925.

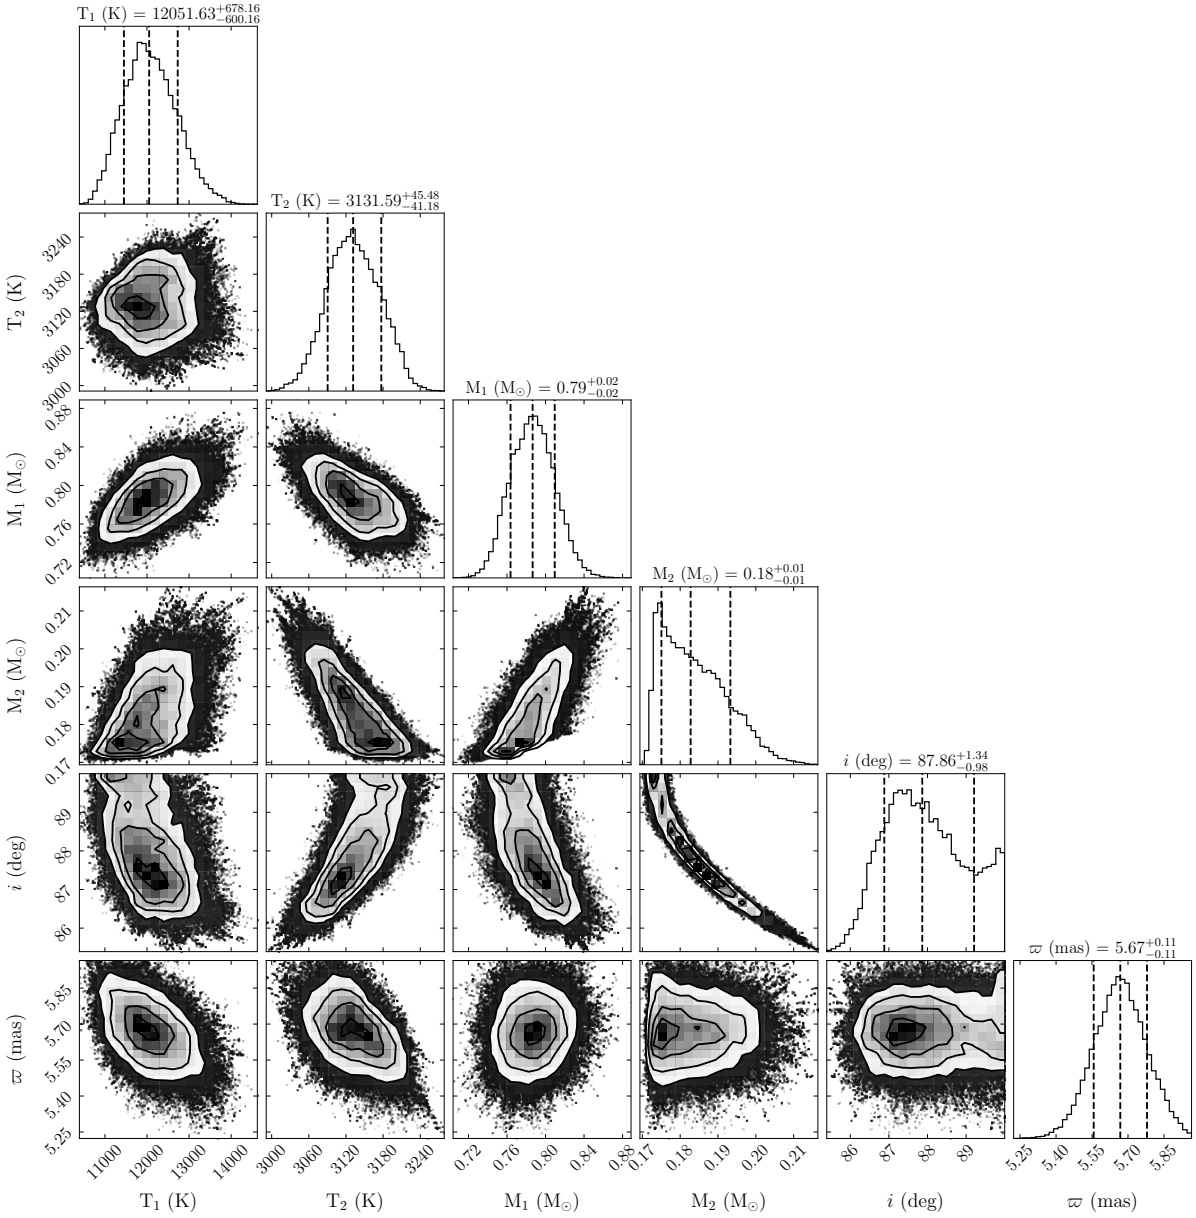

Figure D.3: As in Figure D.1 but for ZTF J0528+2156.

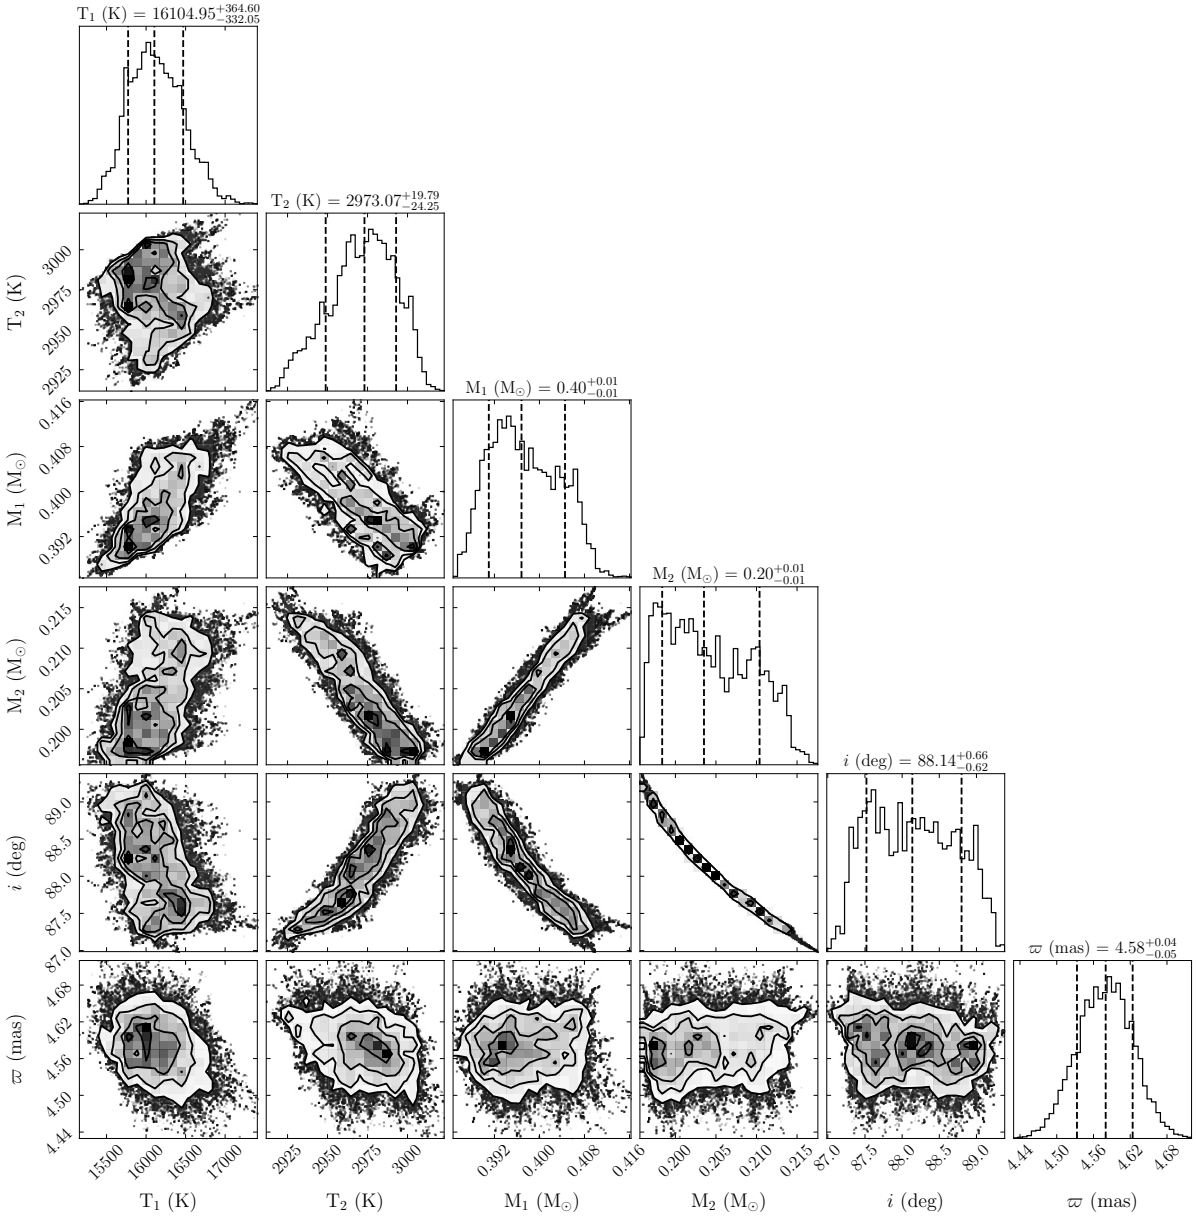

Figure D.4: As in Figure D.1 but for ZTF J0537−2450.

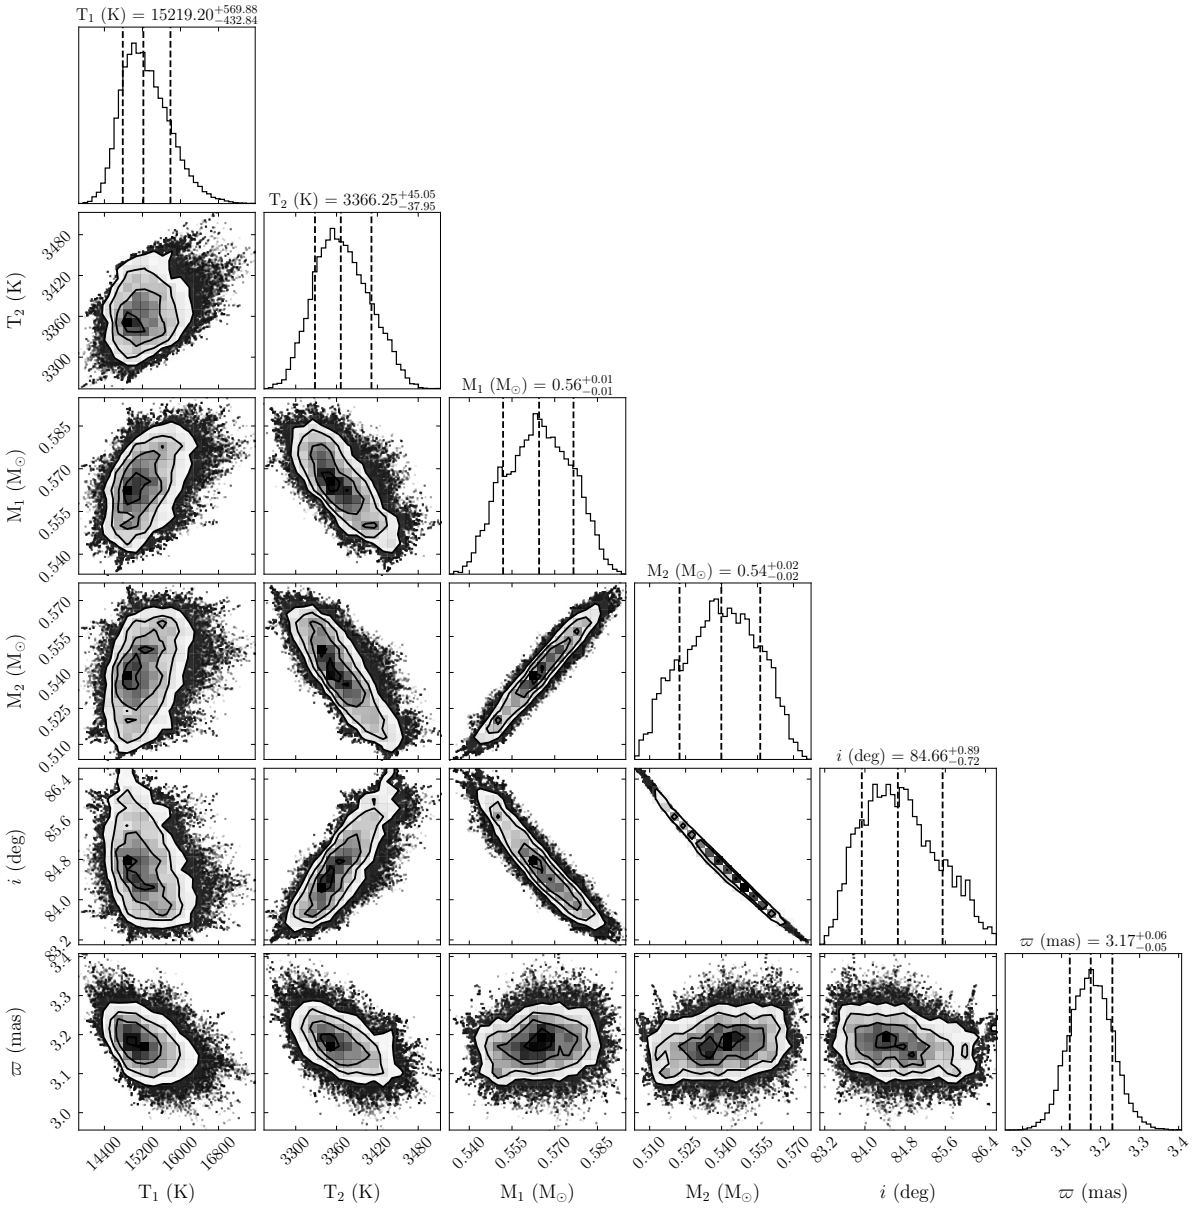

Figure D.5: As in Figure D.1 but for ZTF J0615+0510.

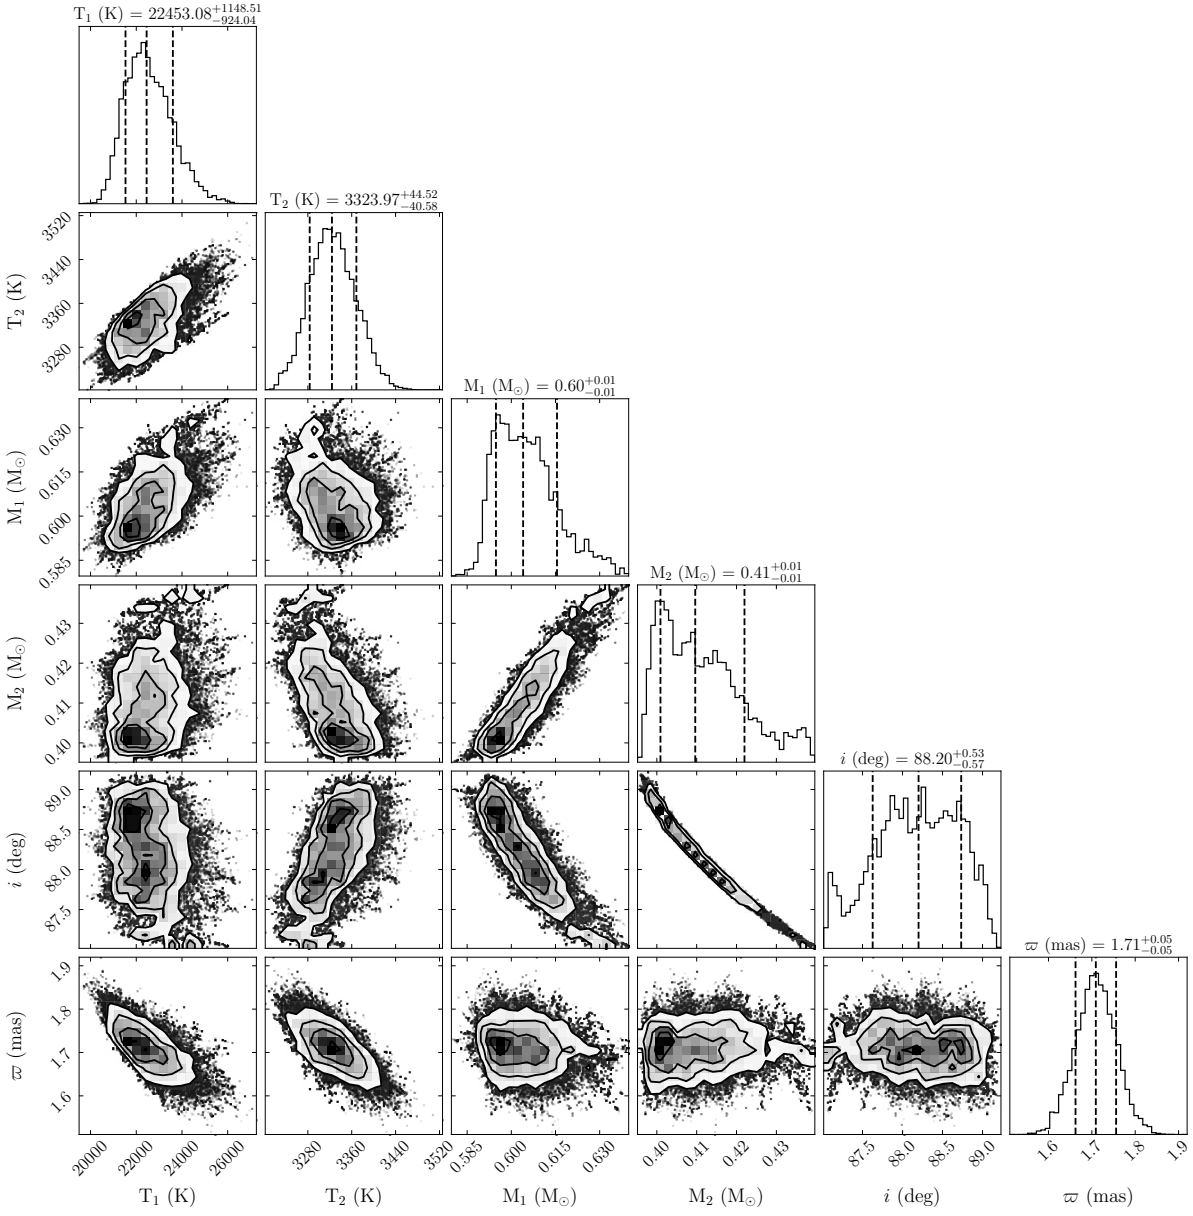

Figure D.6: As in Figure D.1 but for ZTF J0638+0910.

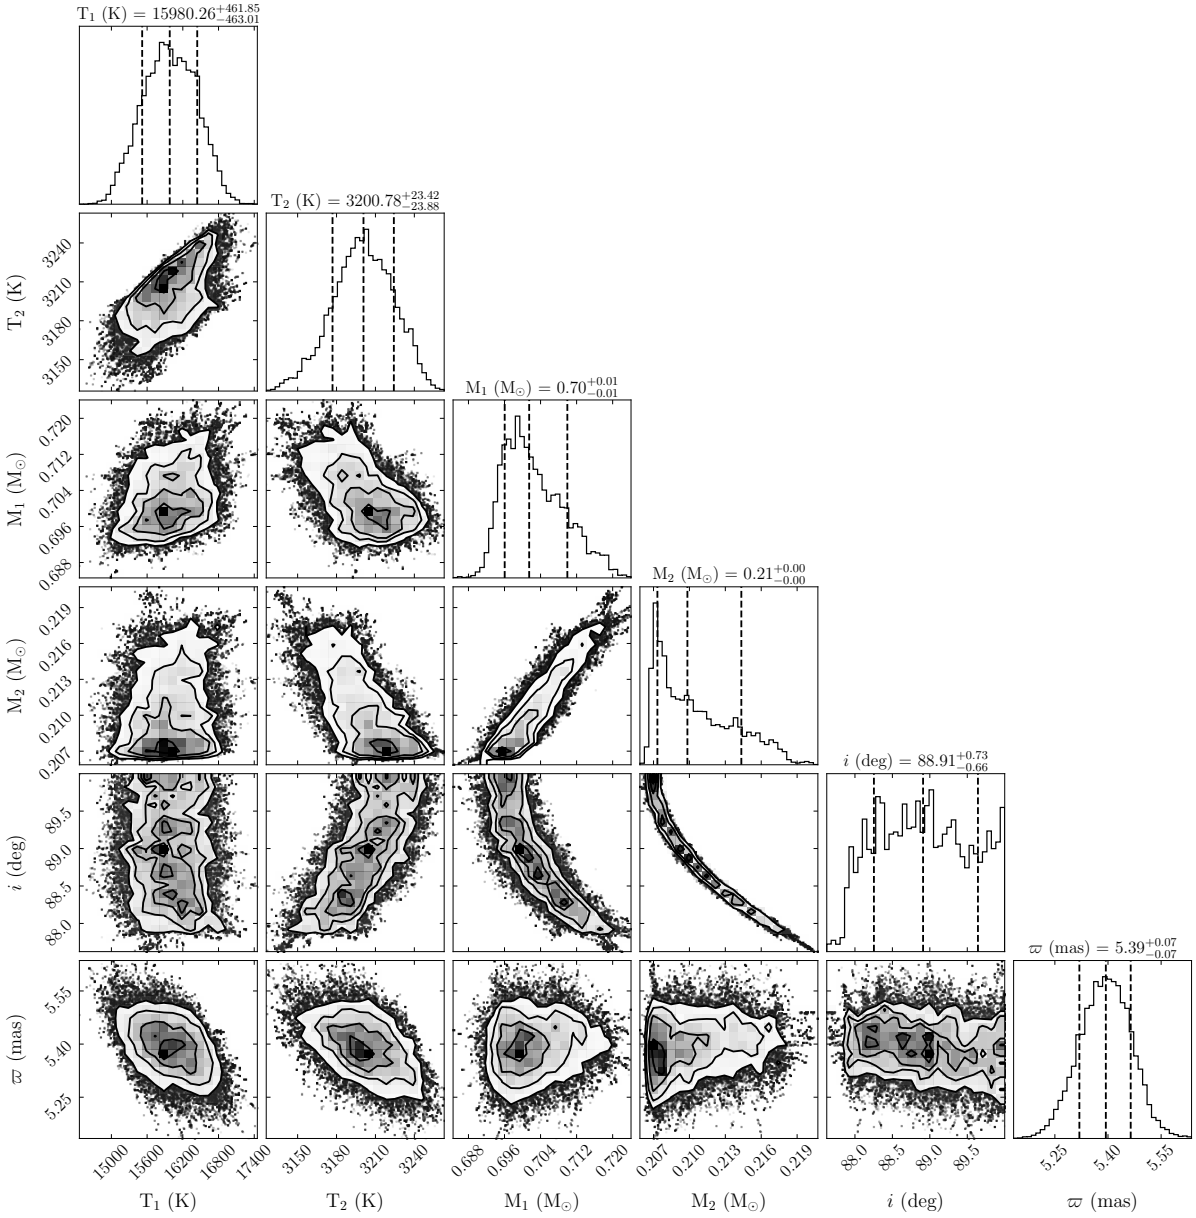

Figure D.7: As in Figure D.1 but for ZTF J0639+1919.

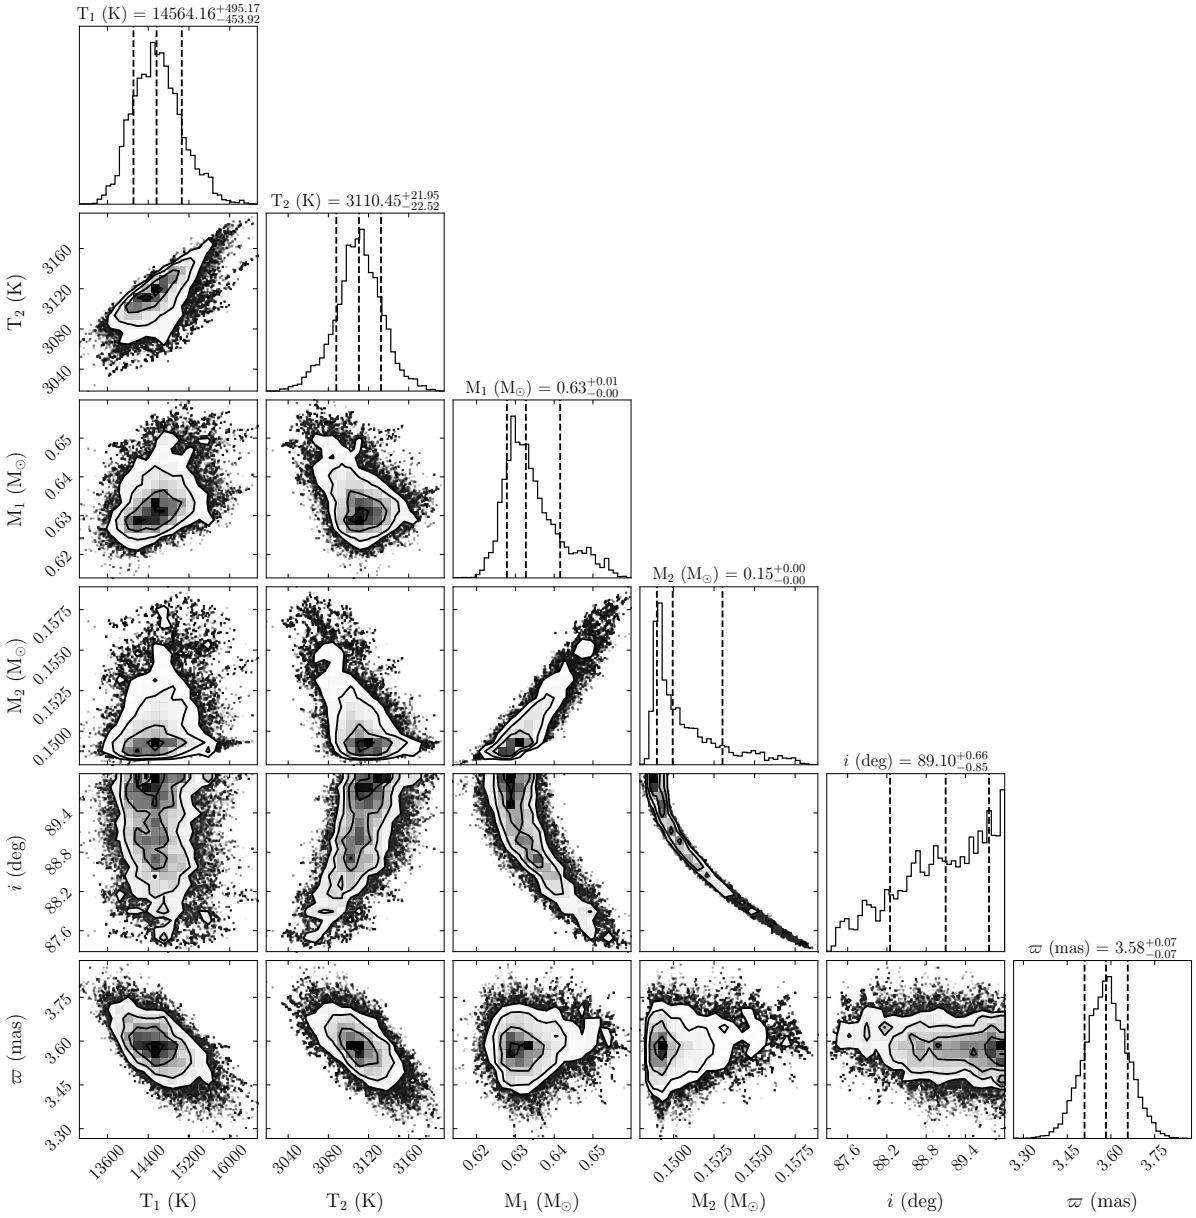

Figure D.8: As in Figure D.1 but for ZTF J0642+1314.

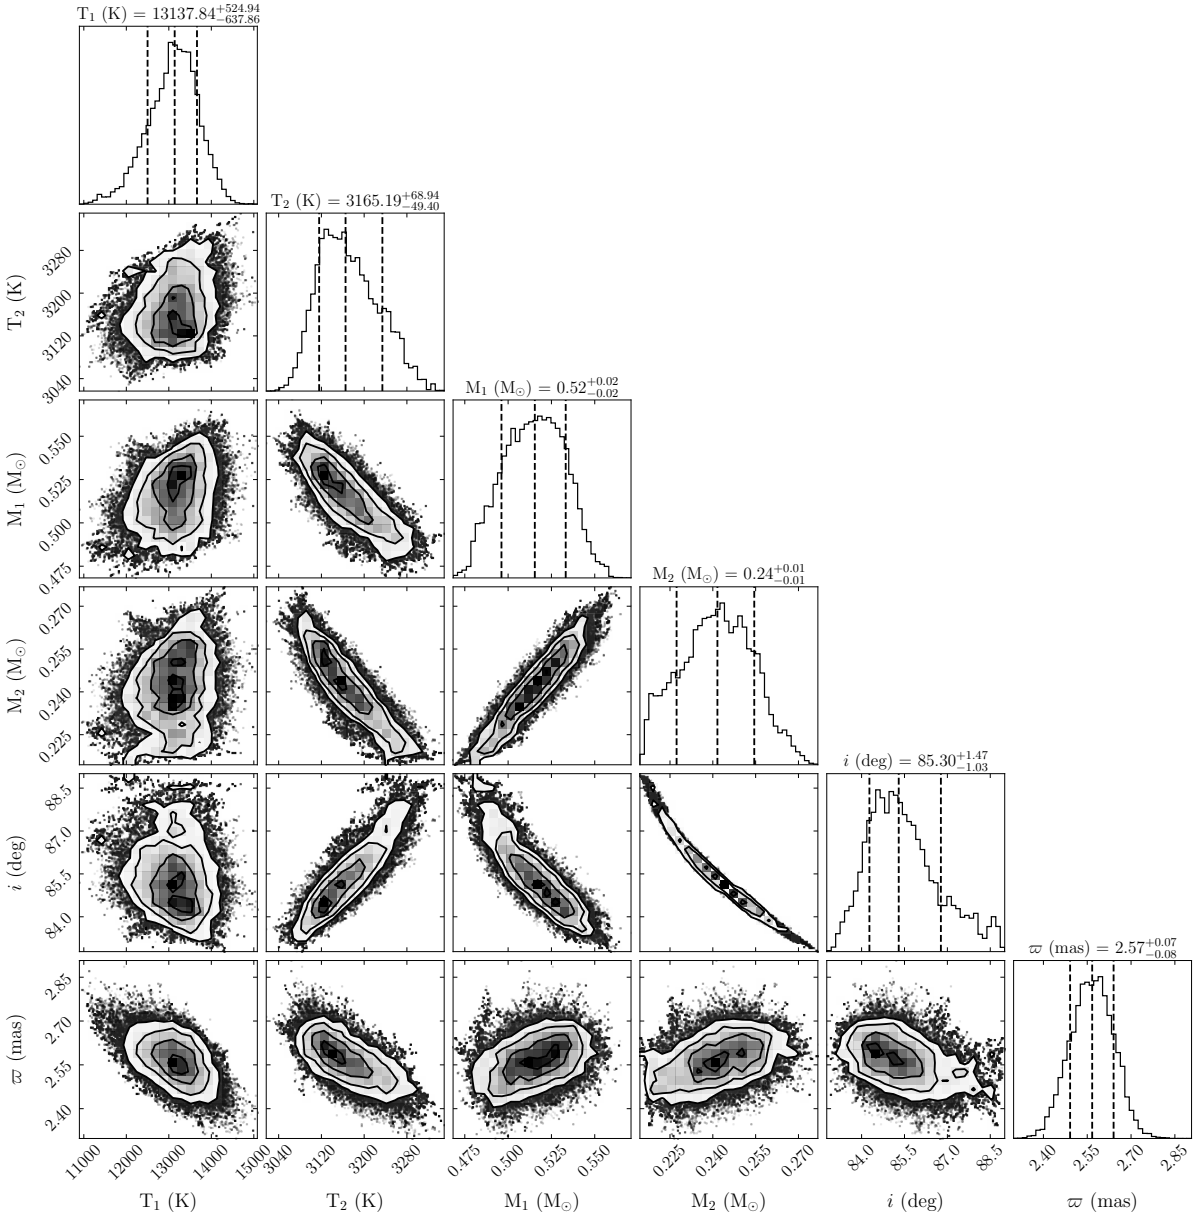

Figure D.9: As in Figure D.1 but for ZTF J0651+1452.

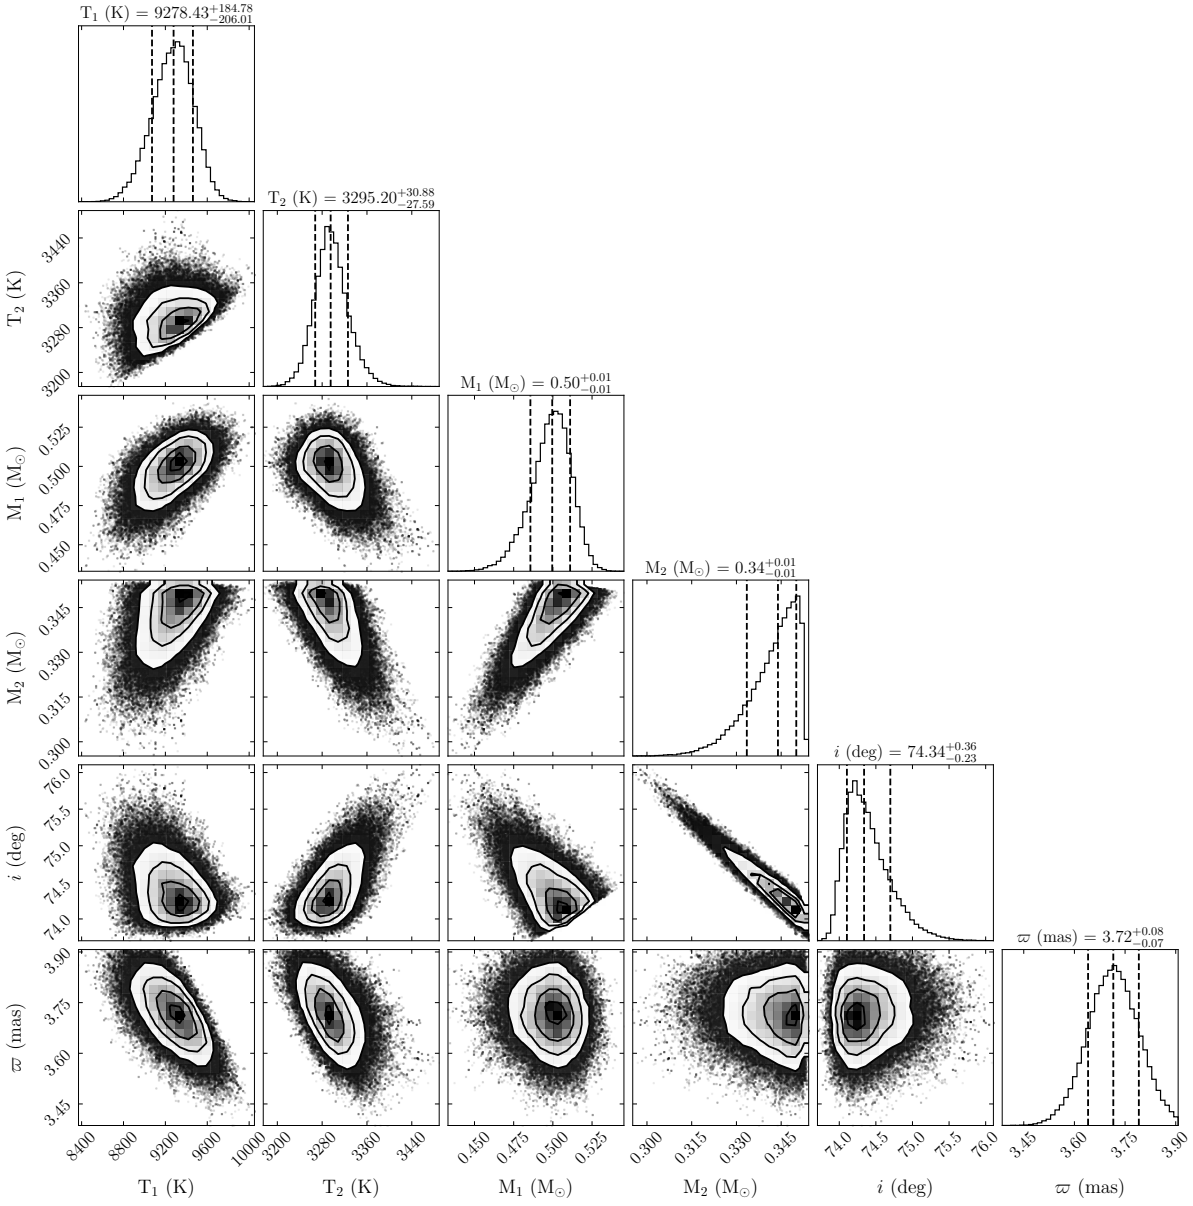

Figure D.10: As in Figure D.1 but for ZTF J0704−0201.

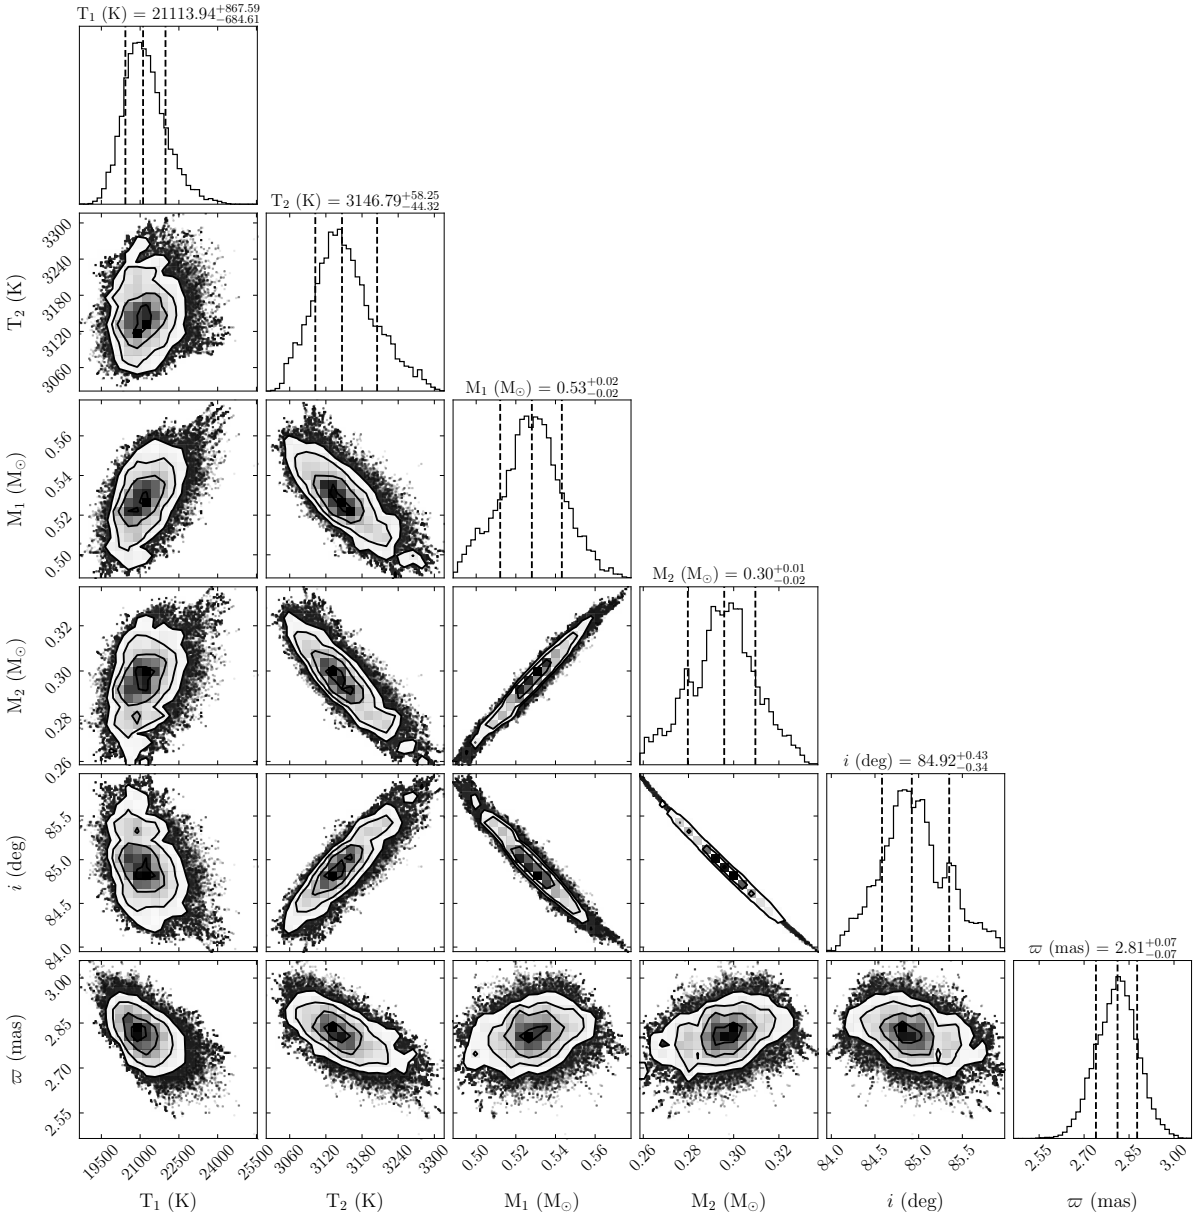

Figure D.11: As in Figure D.1 but for ZTF J0717+1136.

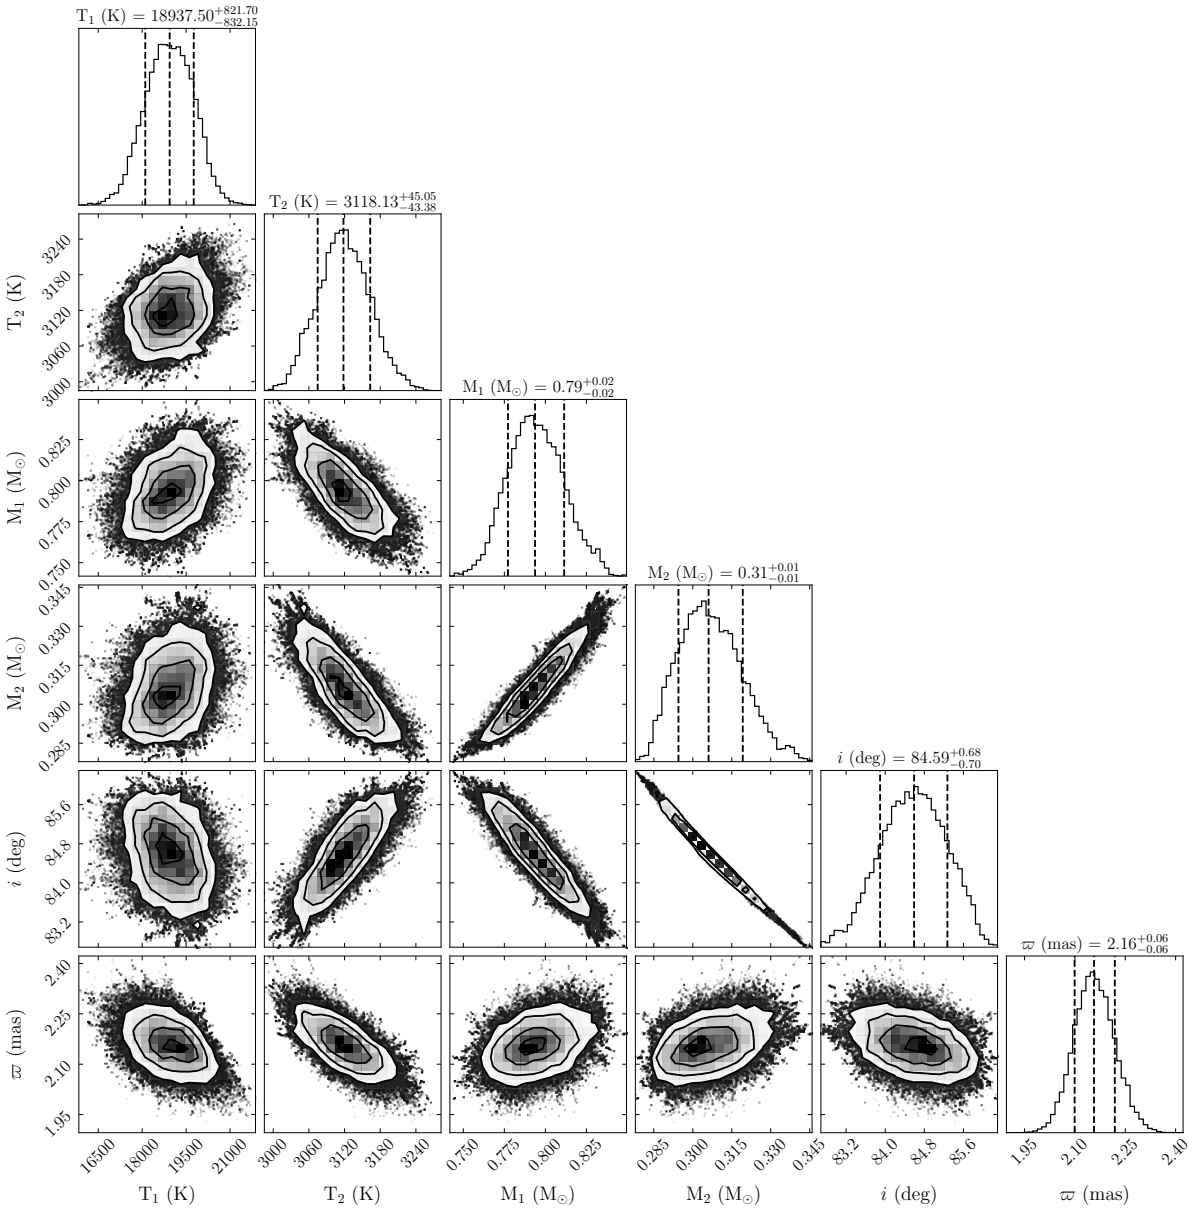

Figure D.12: As in Figure D.1 but for ZTF J0718−0852.

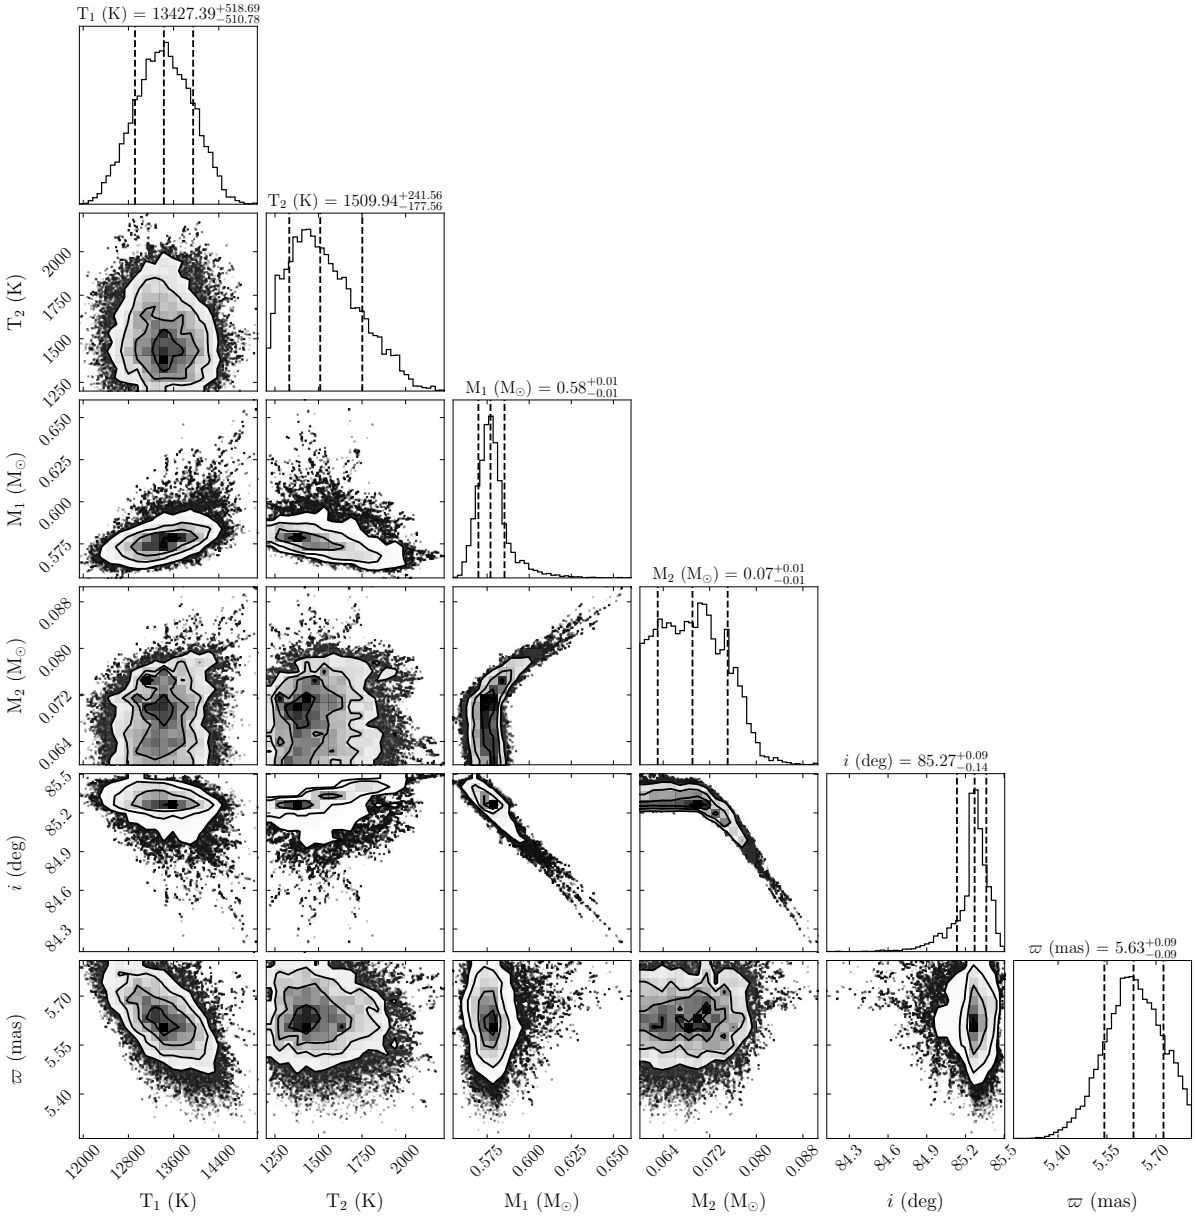

Figure D.13: As in Figure D.1 but for ZTF J0804−0215.

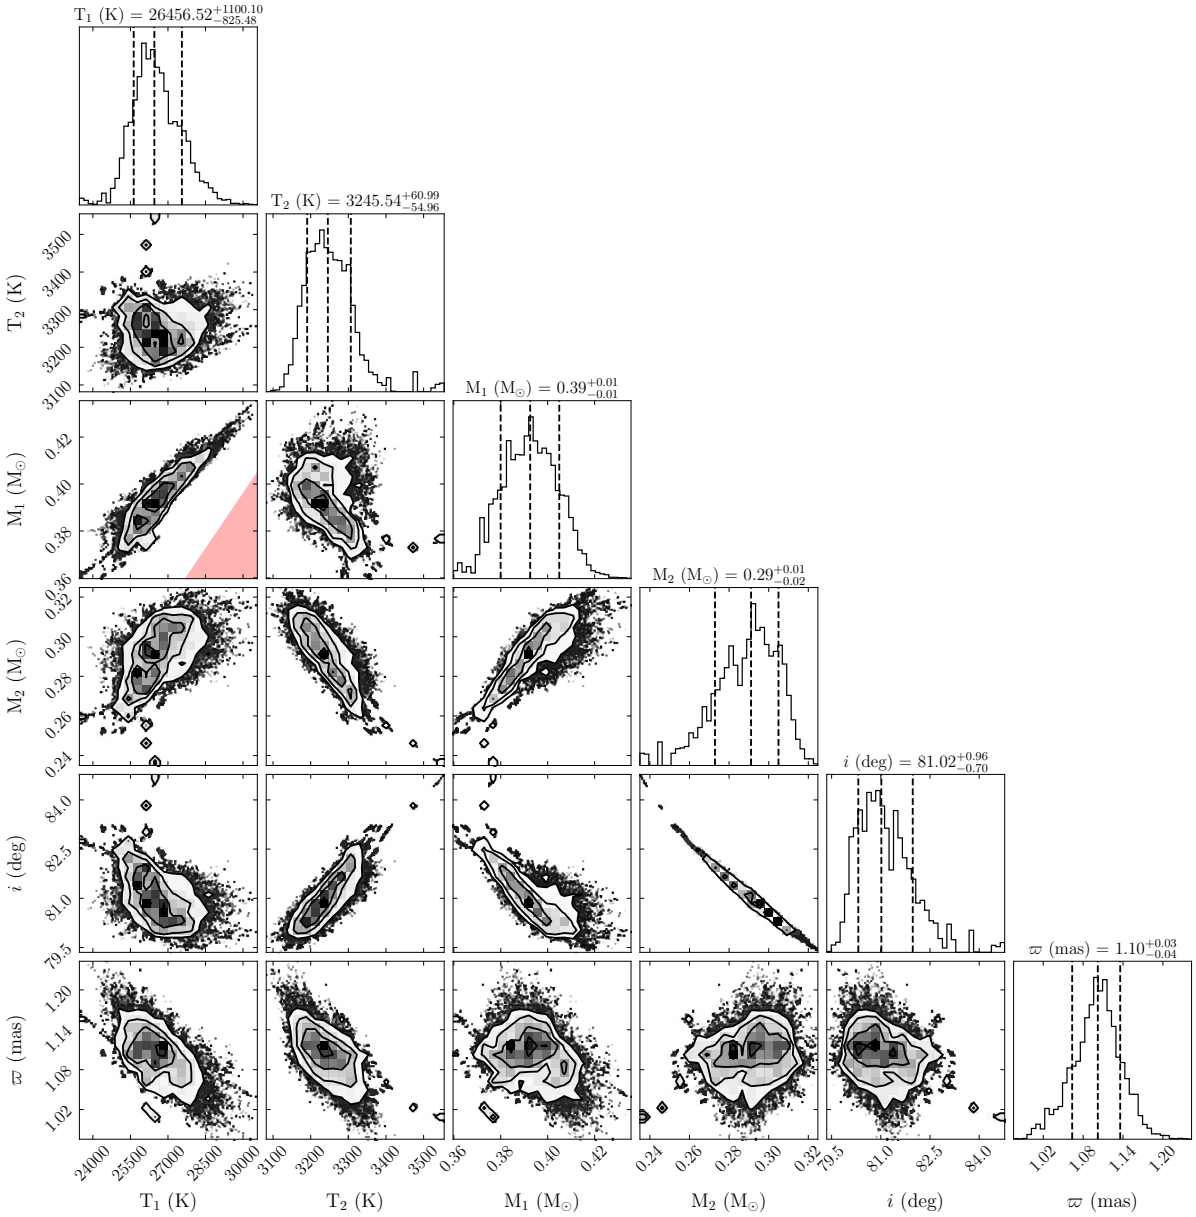

Figure D.14: As in Figure D.1 but for ZTF J0805−1430. The red region shows the limits of the He-core models of (Panei et al., 2007).

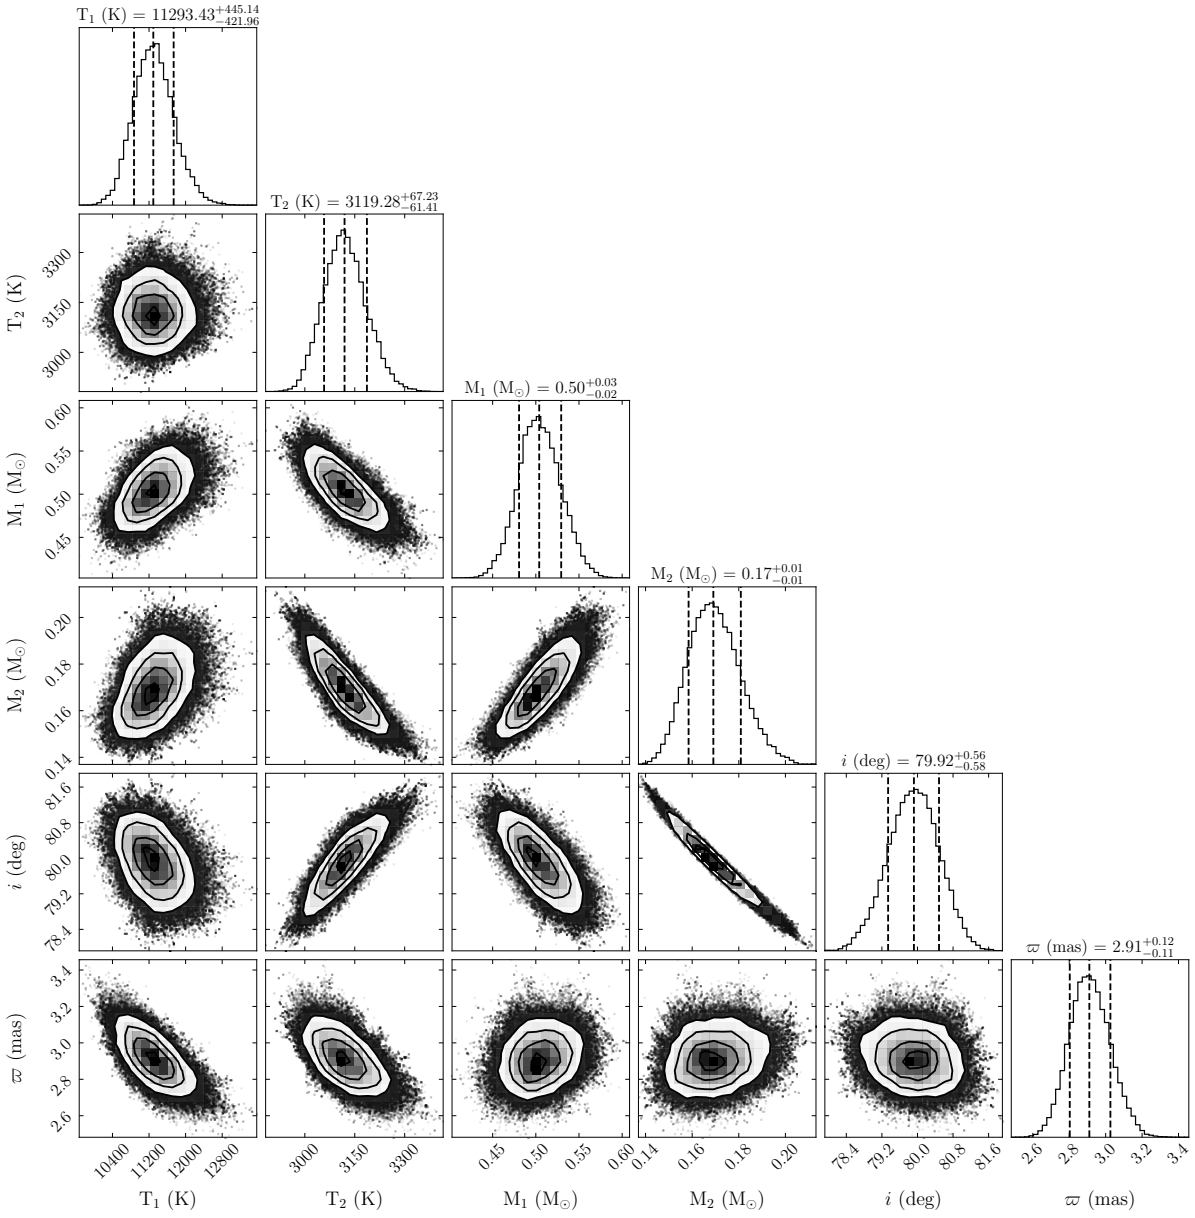

Figure D.15: As in Figure D.1 but for ZTF J0948+2538.

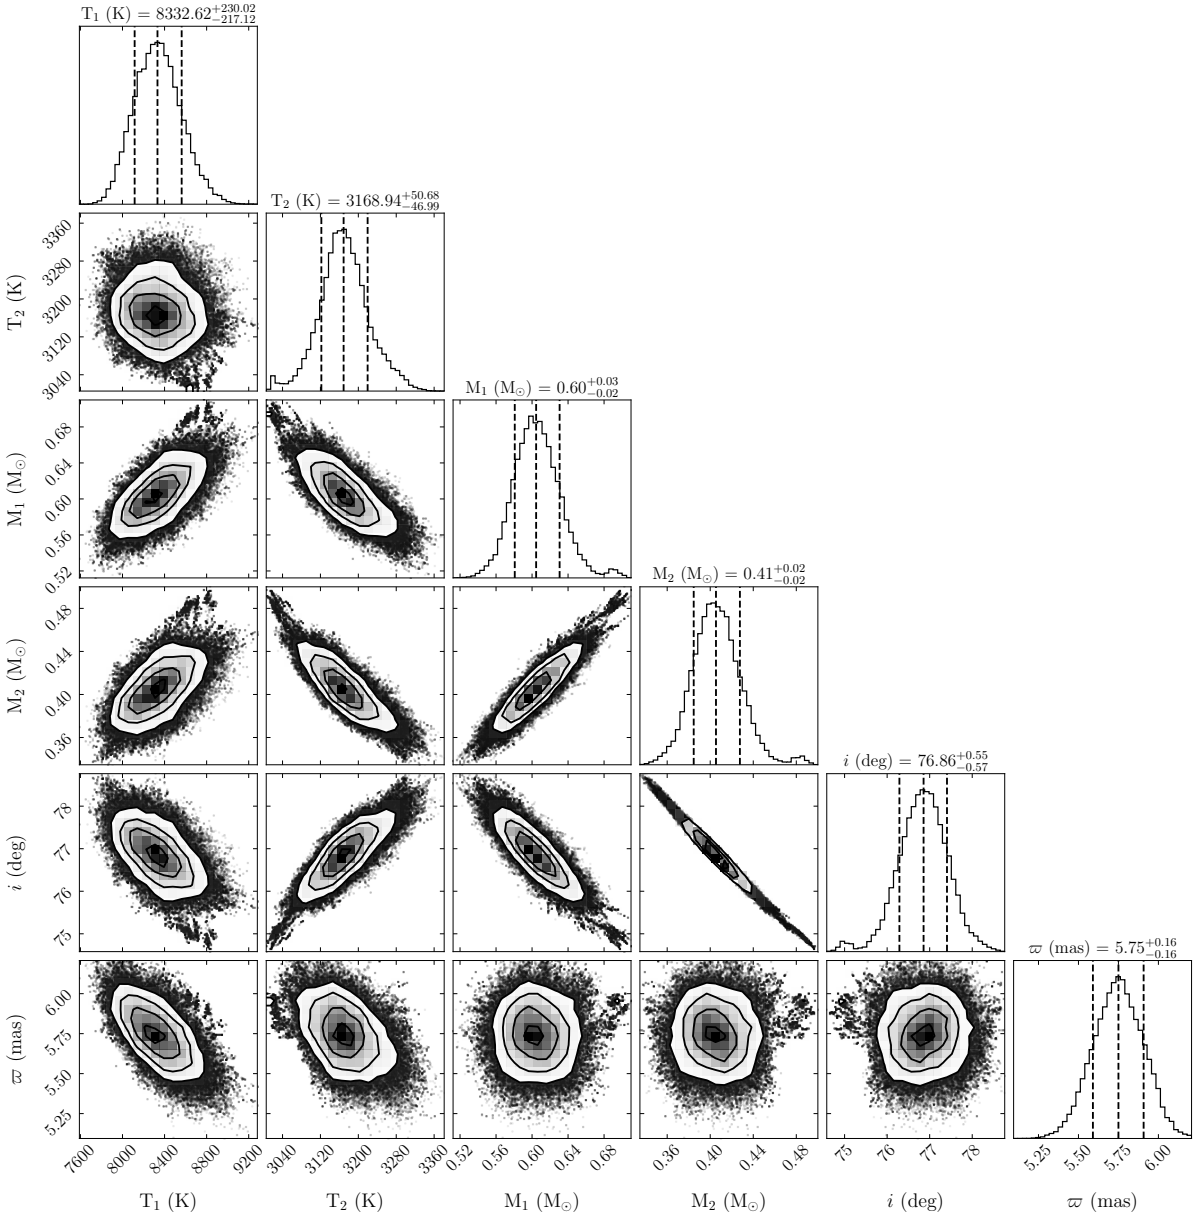

Figure D.16: As in Figure D.1 but for ZTF J1022−0803.

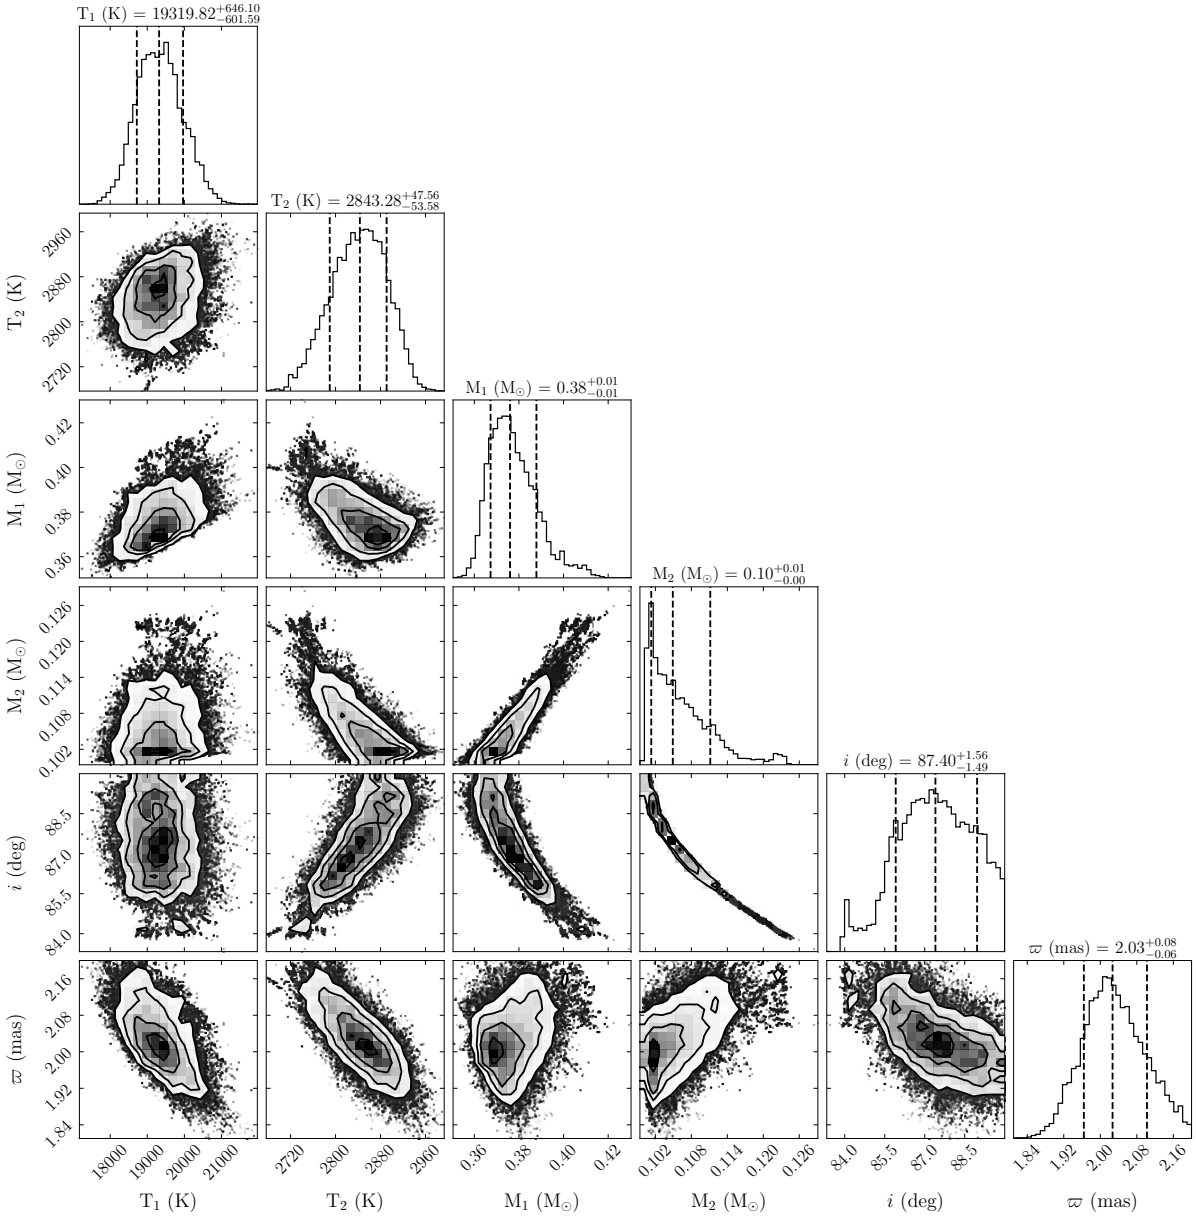

Figure D.17: As in Figure D.1 but for ZTF J1026−1013.

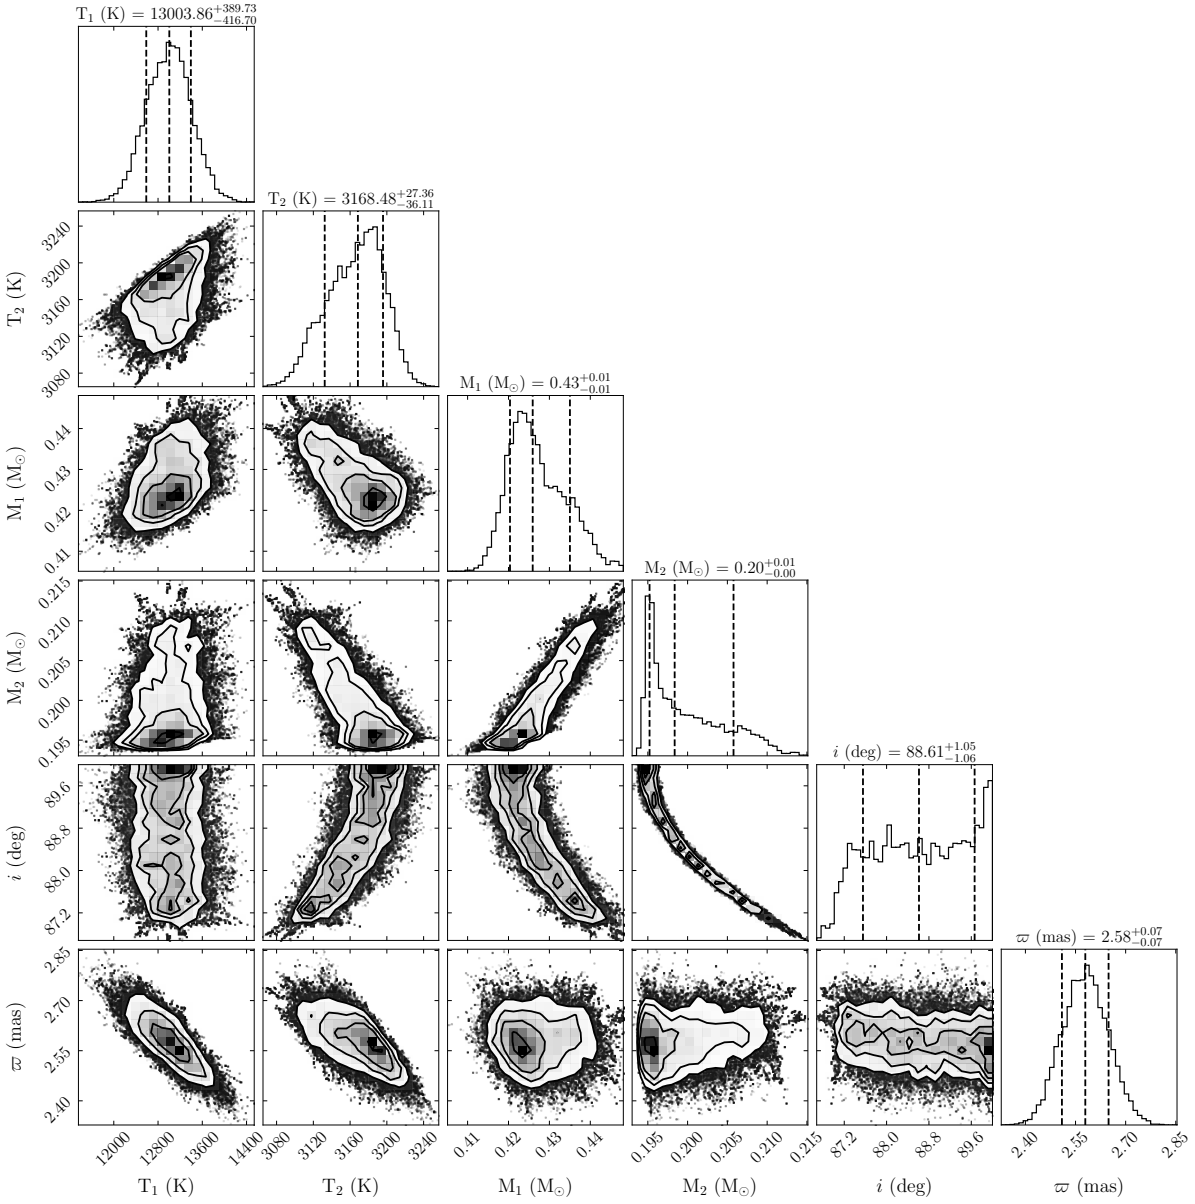

Figure D.18: As in Figure D.1 but for ZTF J1049−1755.

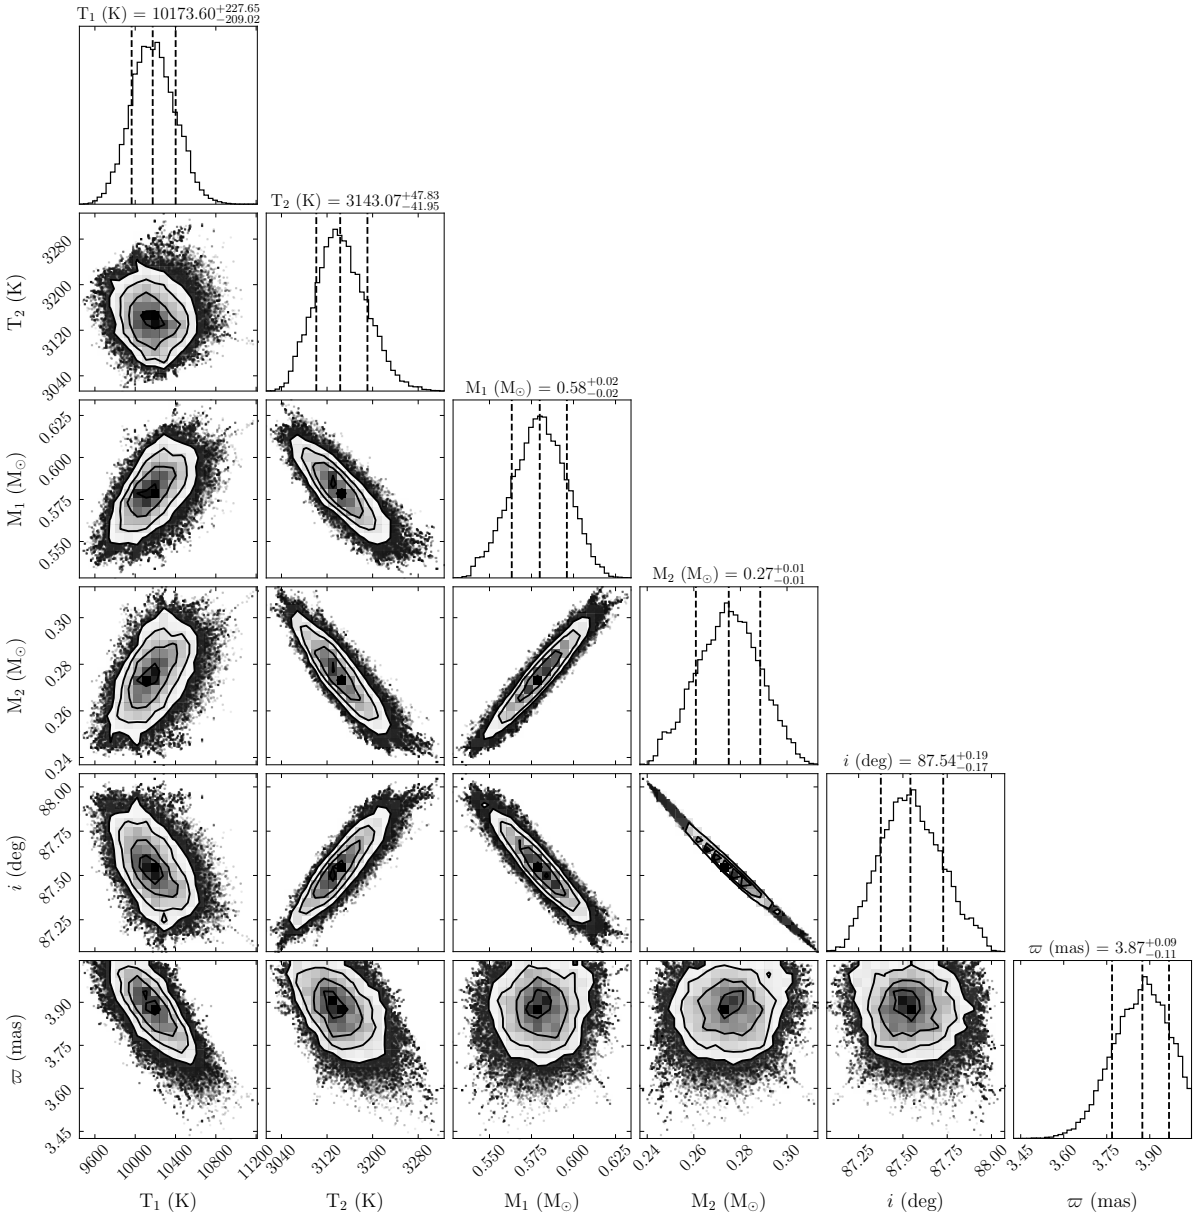

Figure D.19: As in Figure D.1 but for ZTF J1220+0821.

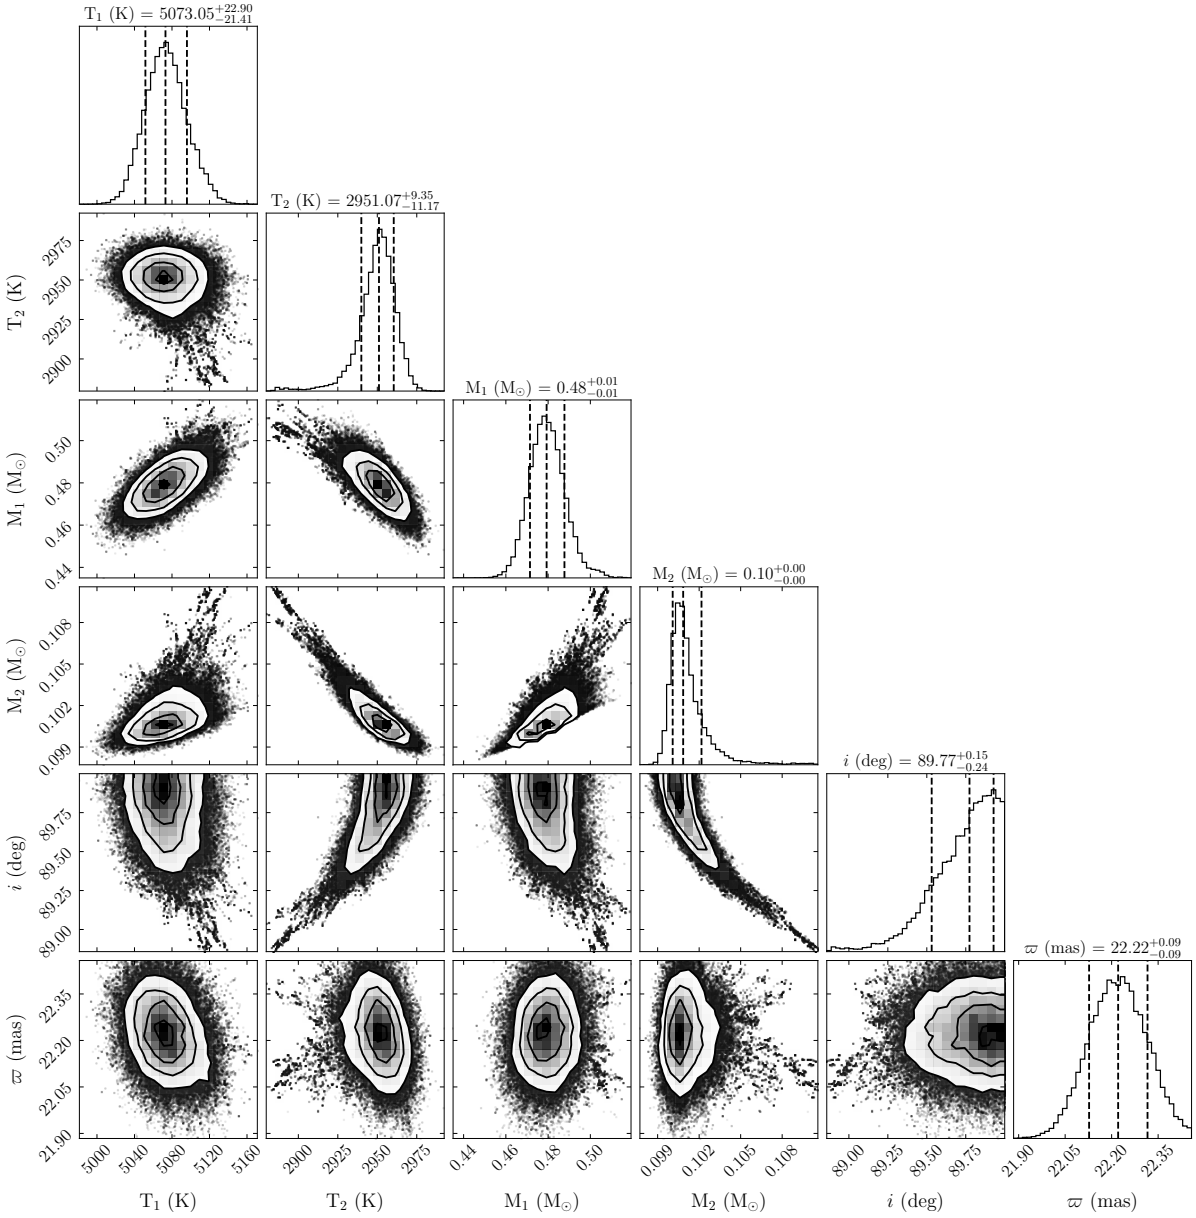

Figure D.20: As in Figure D.1 but for ZTF J1256+2117.

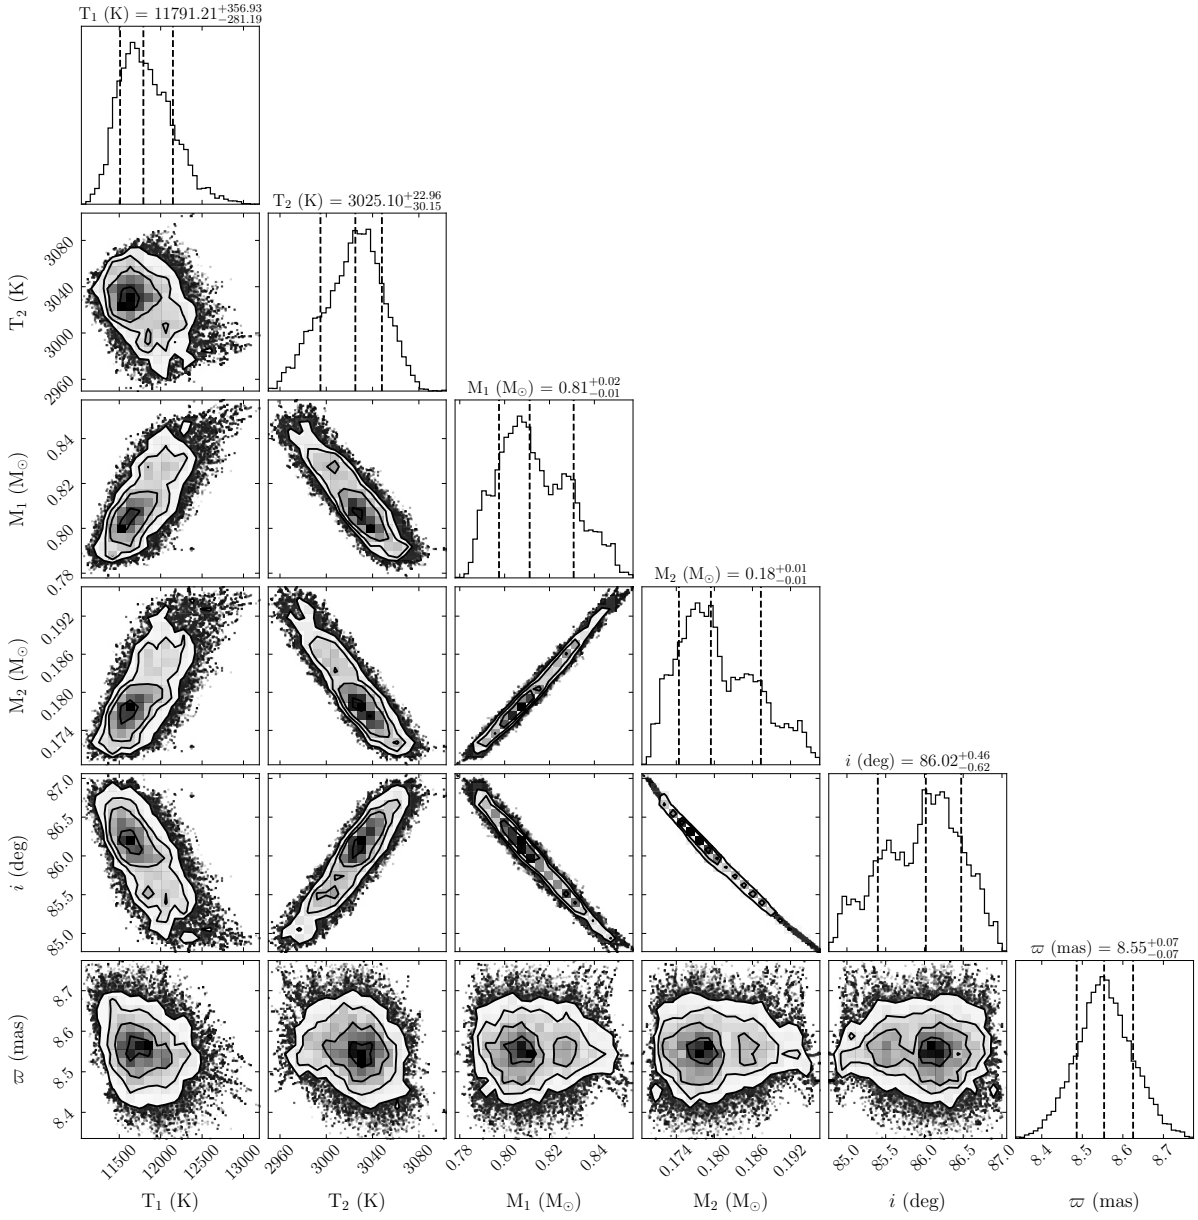

Figure D.21: As in Figure D.1 but for ZTF J1302−0032.

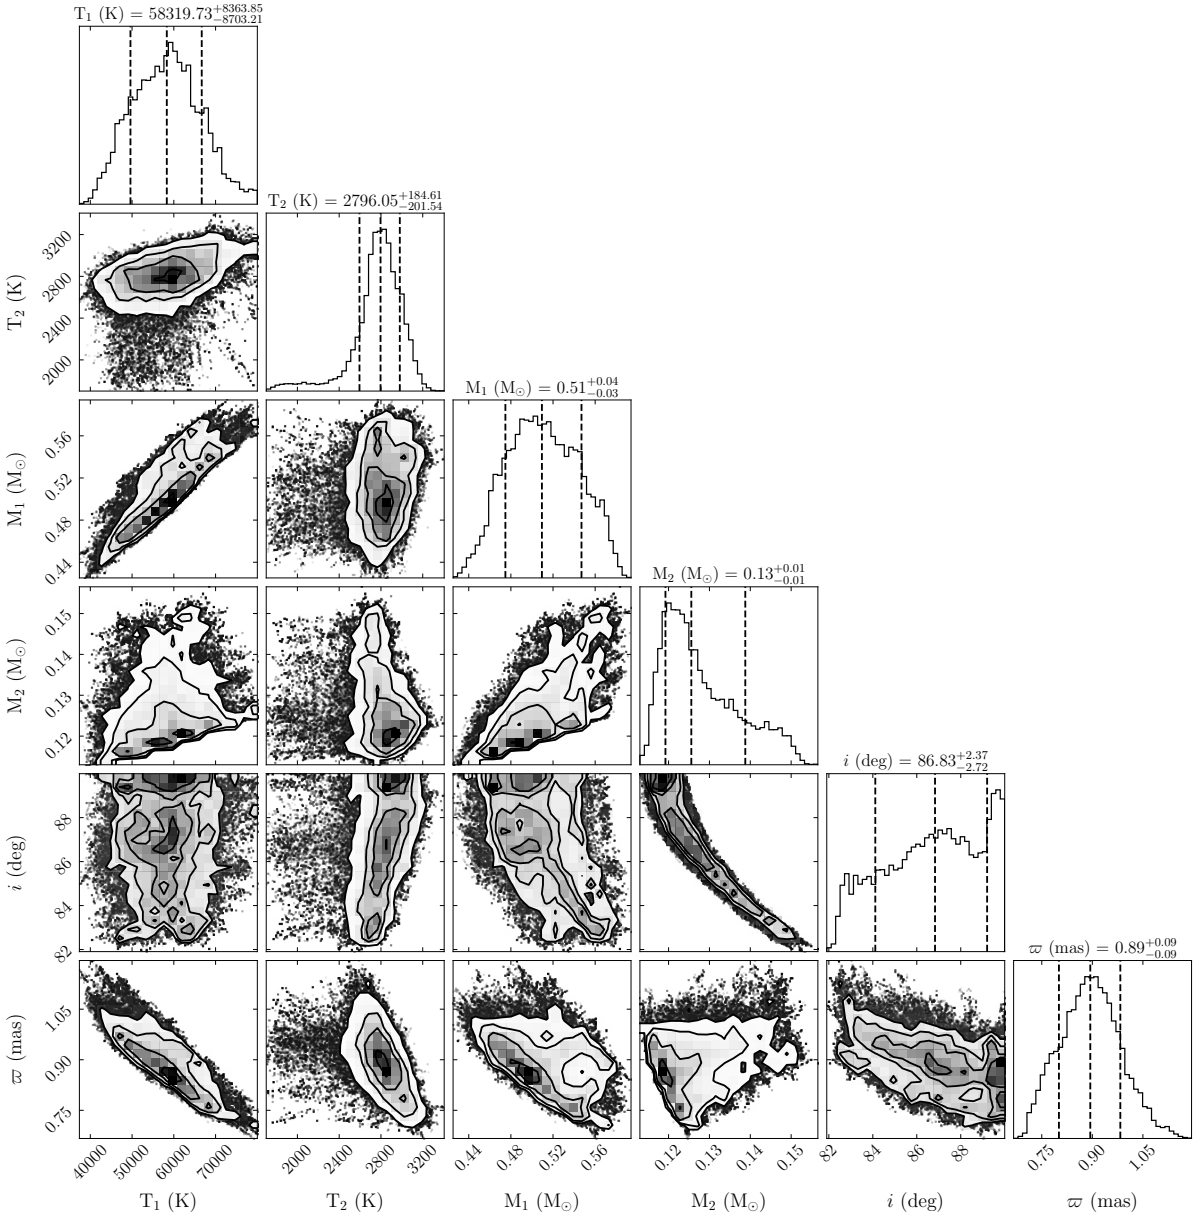

Figure D.22: As in Figure D.1 but for ZTF J1341−0626.

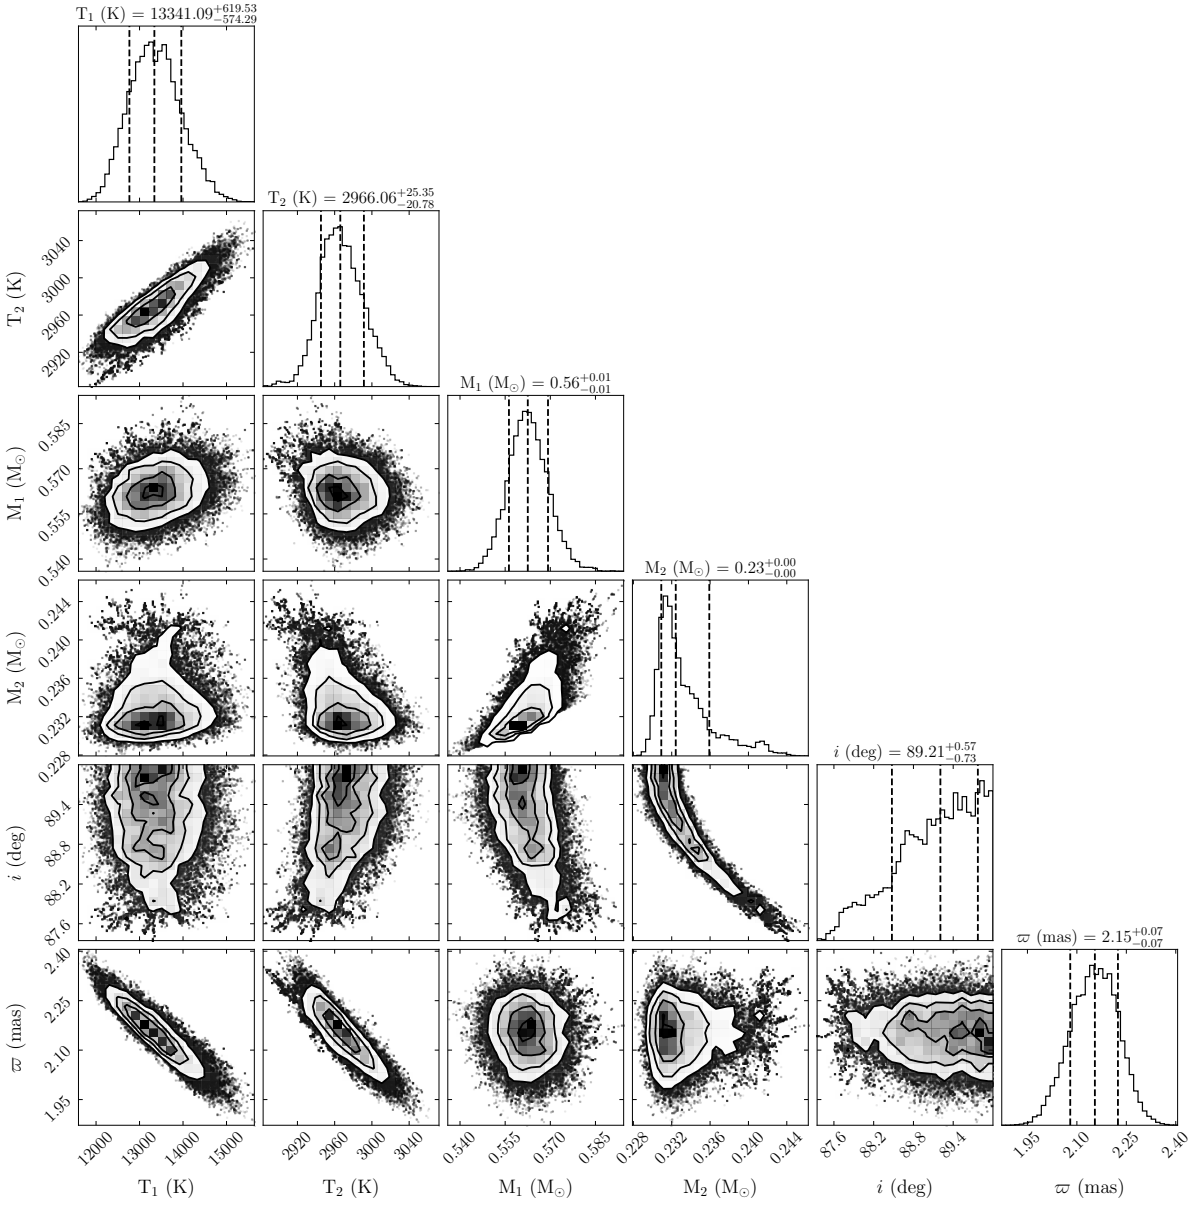

Figure D.23: As in Figure D.1 but for ZTF J1400+0814.

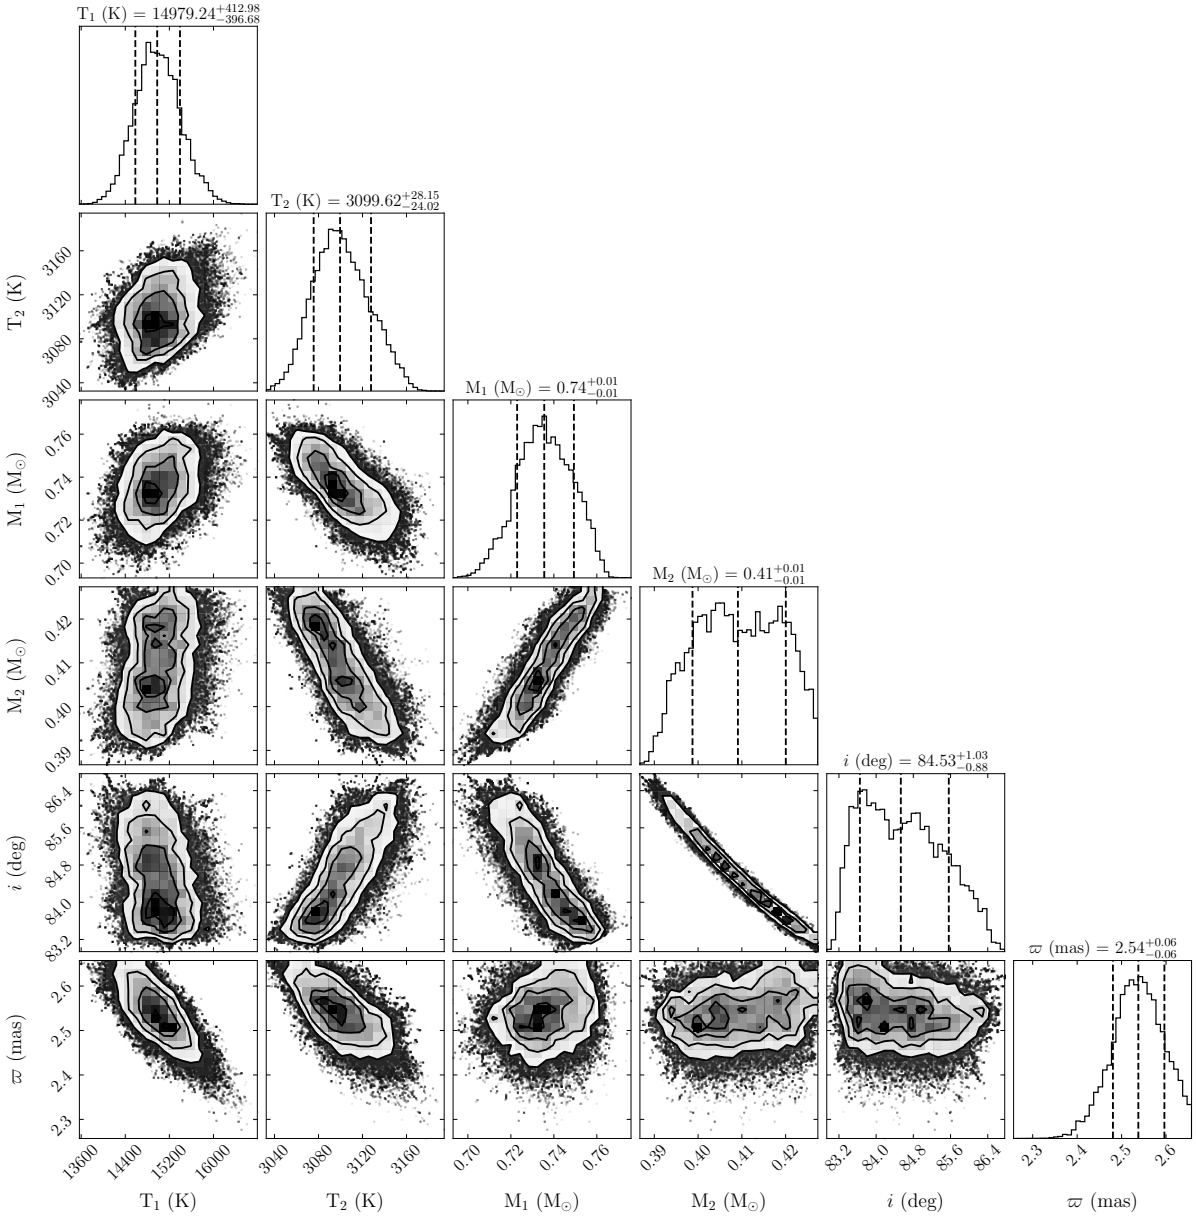

Figure D.24: As in Figure D.1 but for ZTF J1404+0655.

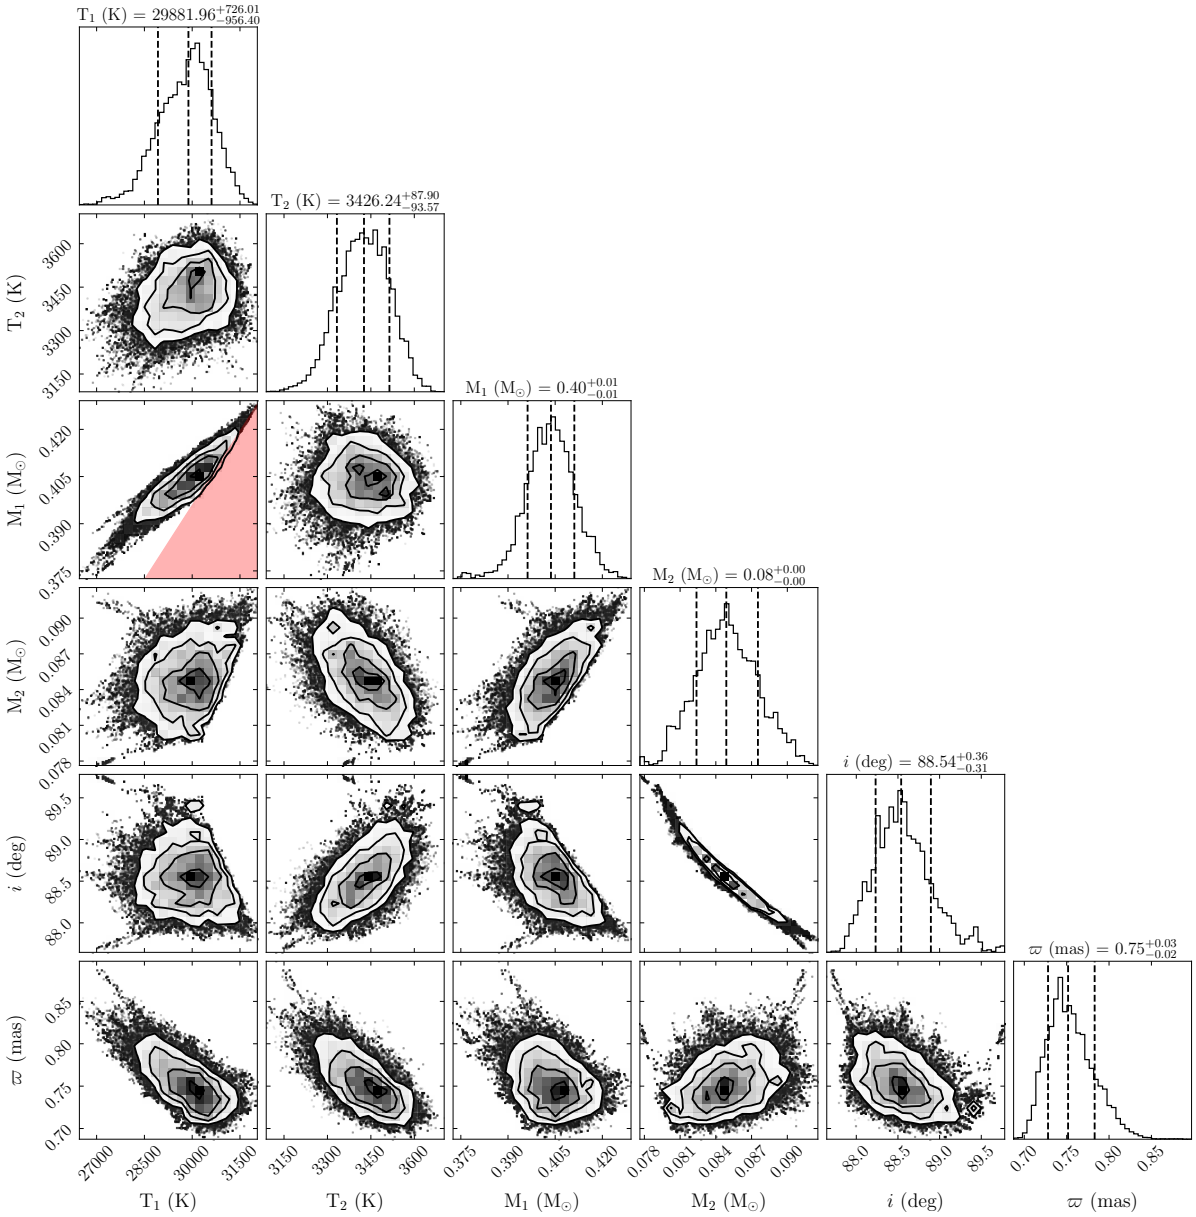

Figure D.25: As in Figure D.1 but for ZTF J1405+1039. The red region shows the limits of the He-core models of (Panei et al., 2007).

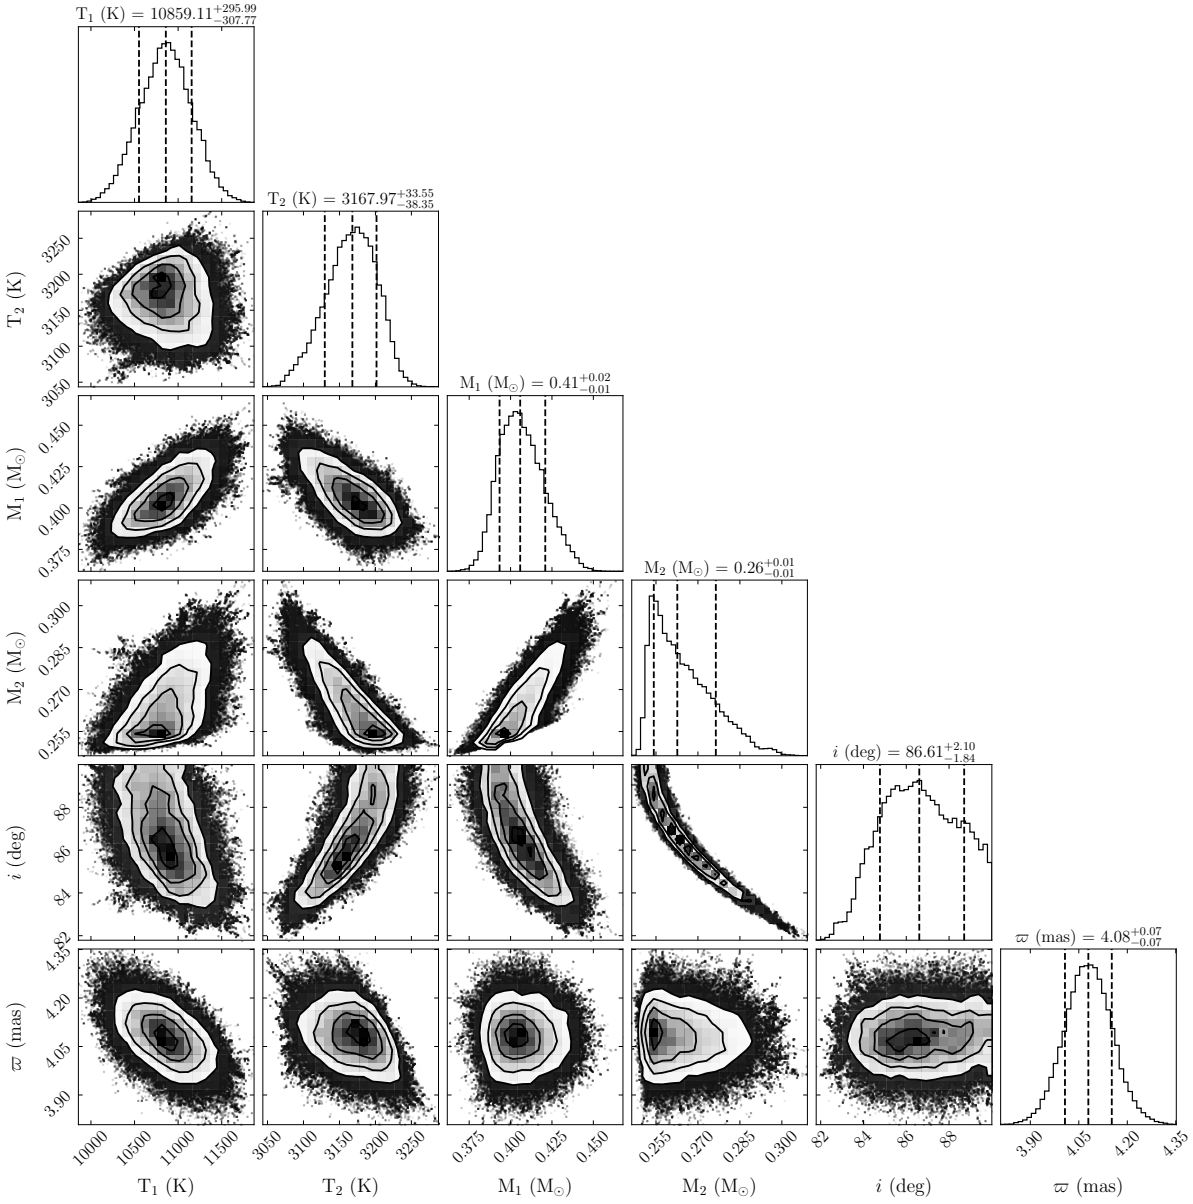

Figure D.26: As in Figure D.1 but for ZTF J1407+2115.

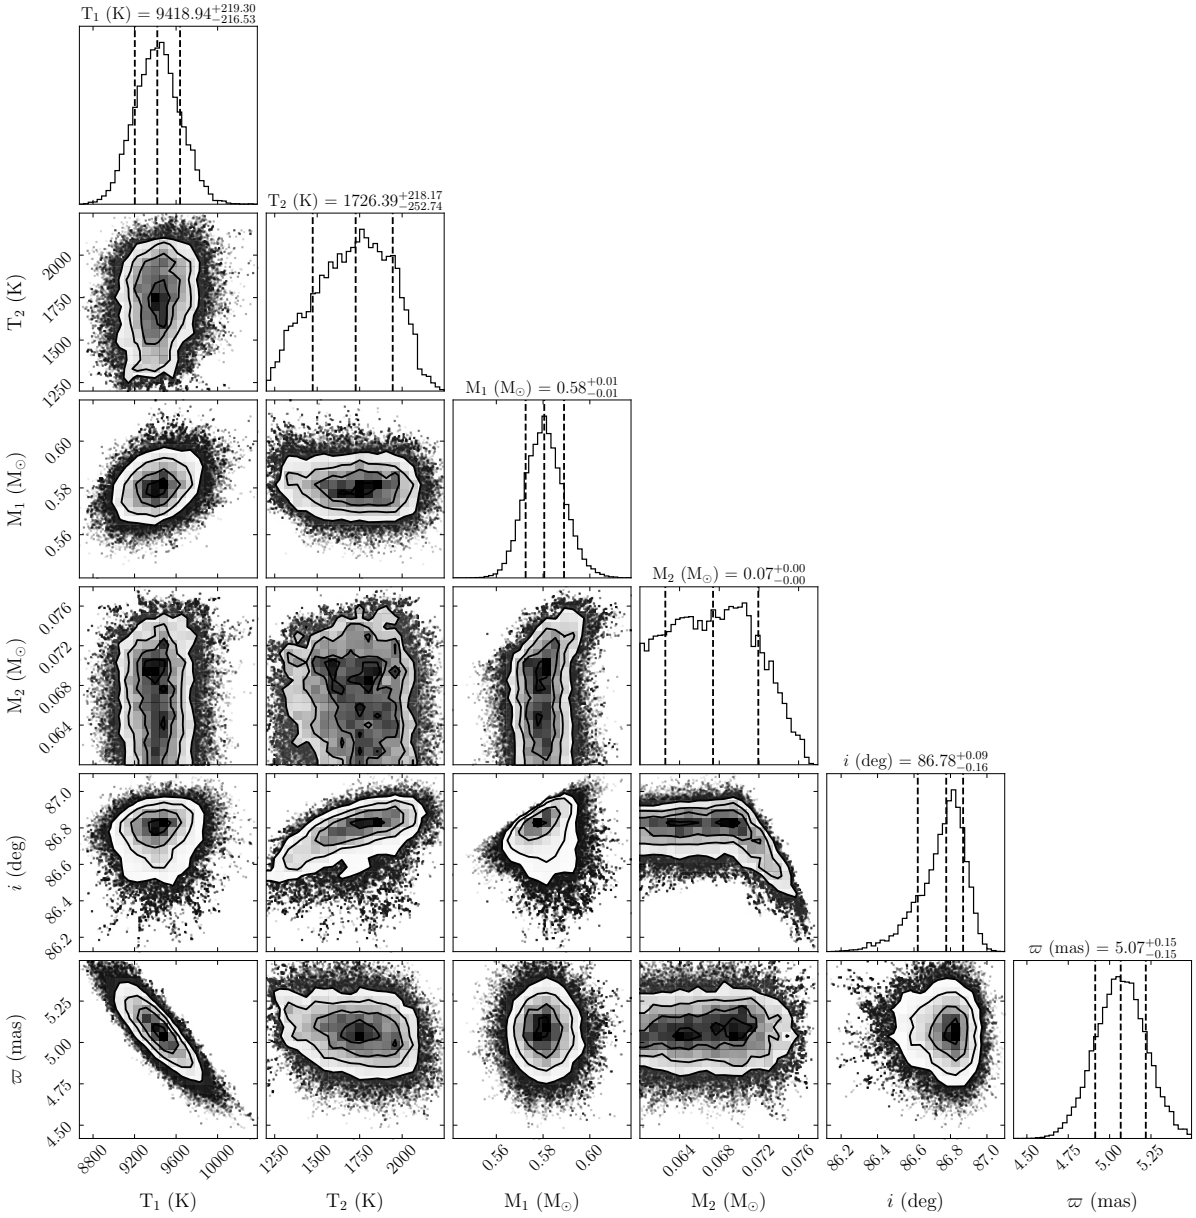

Figure D.27: As in Figure D.1 but for ZTF J1458+1313.

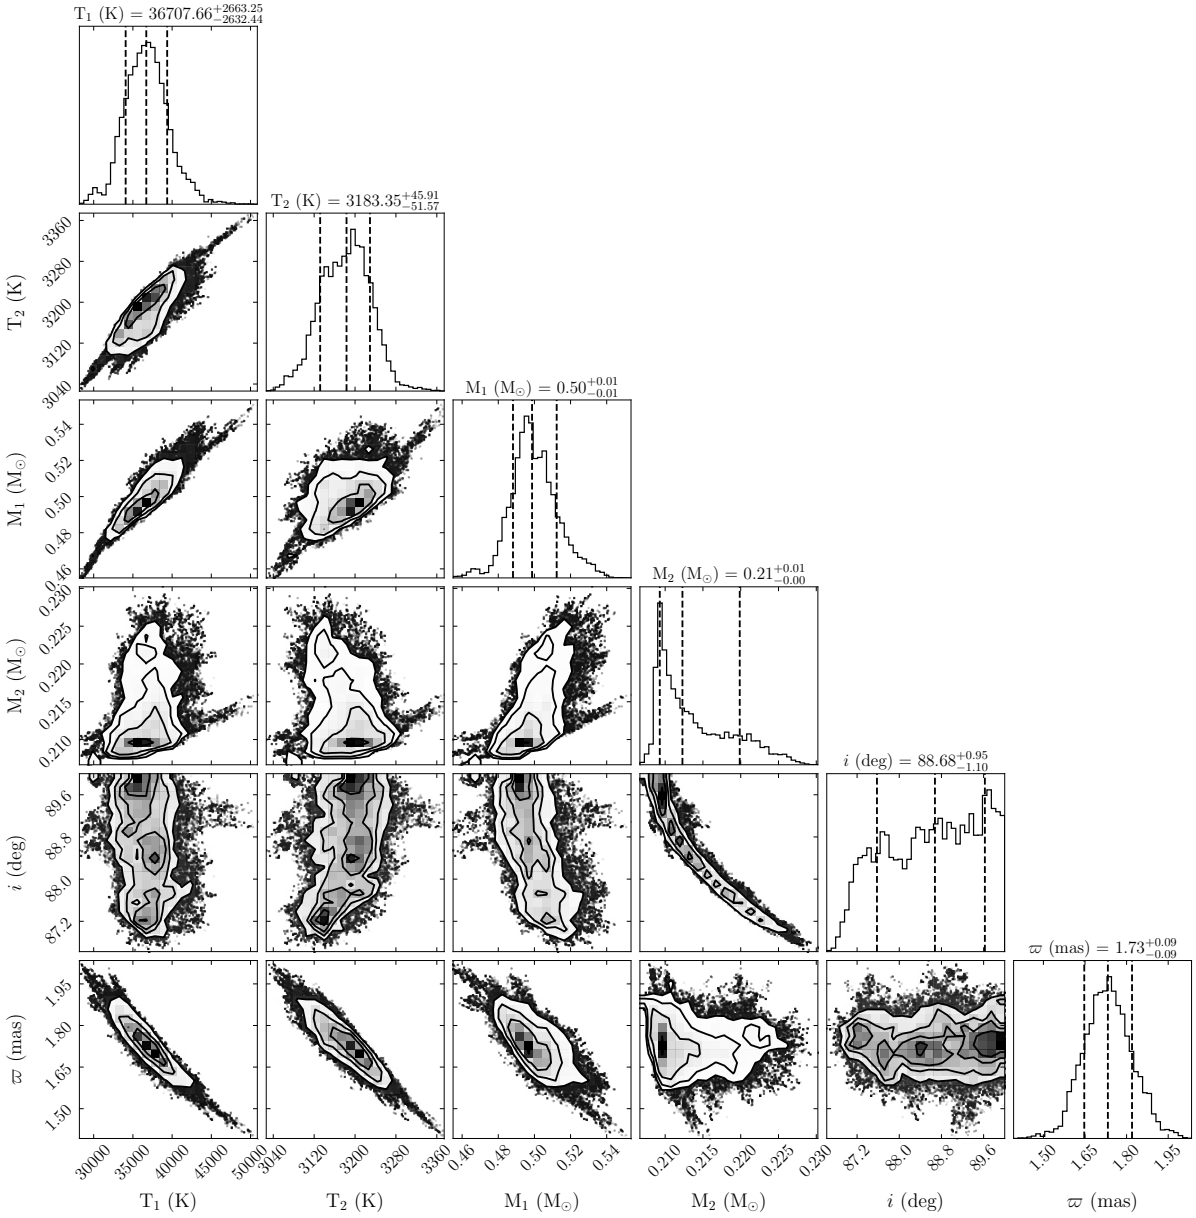

Figure D.28: As in Figure D.1 but for ZTF J1626−1018.

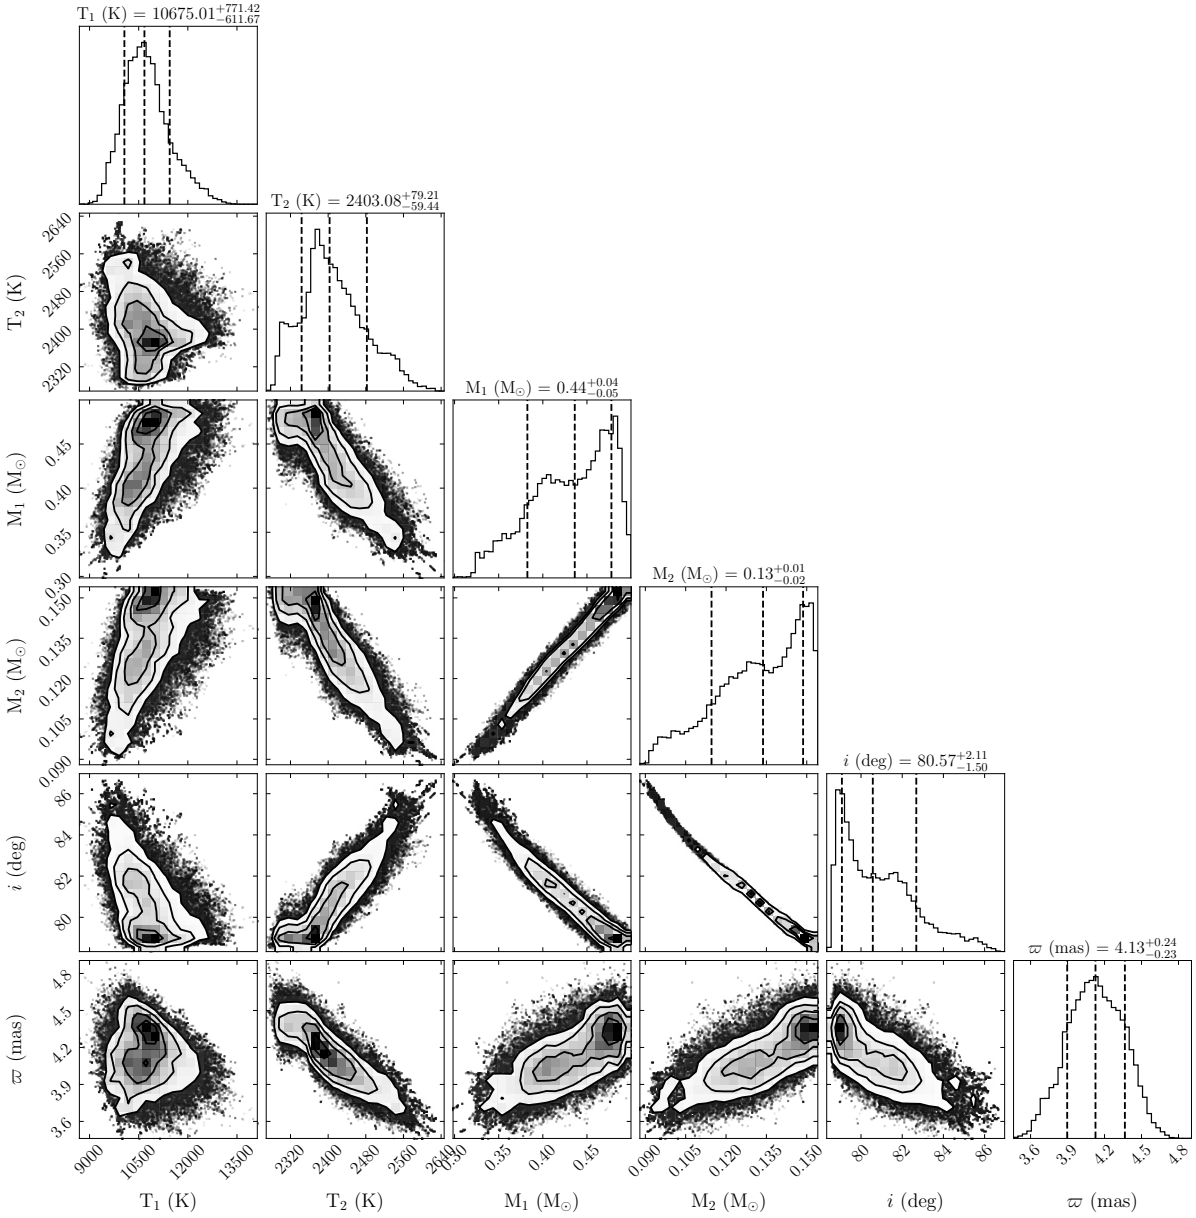

Figure D.29: As in Figure D.1 but for ZTF J1634−2713.

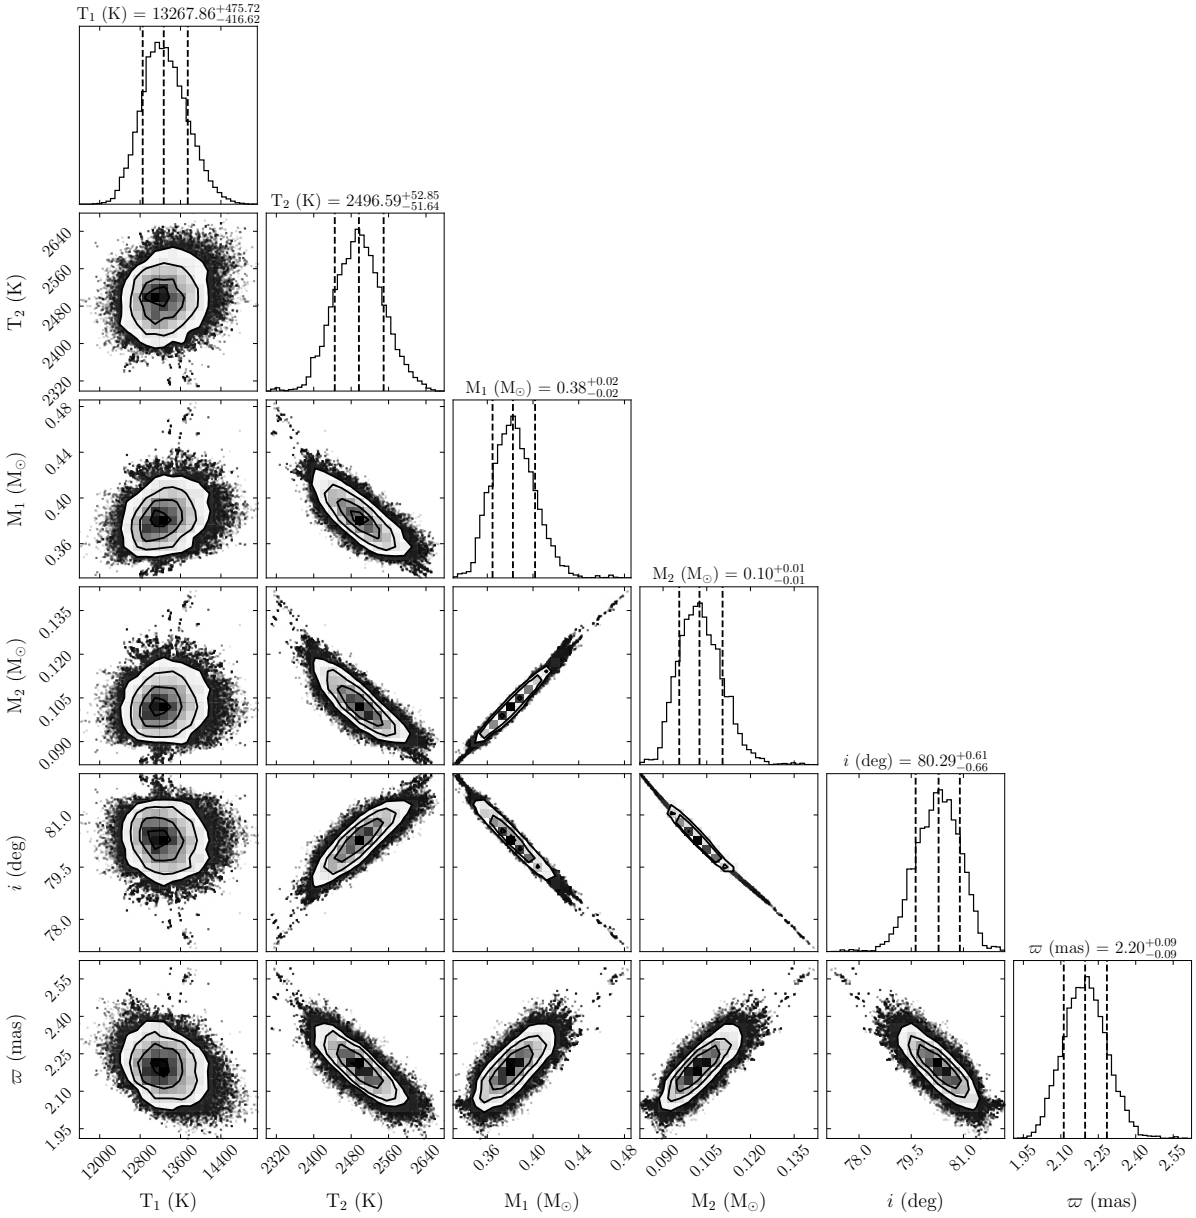

Figure D.30: As in Figure D.1 but for ZTF J1644+2434.

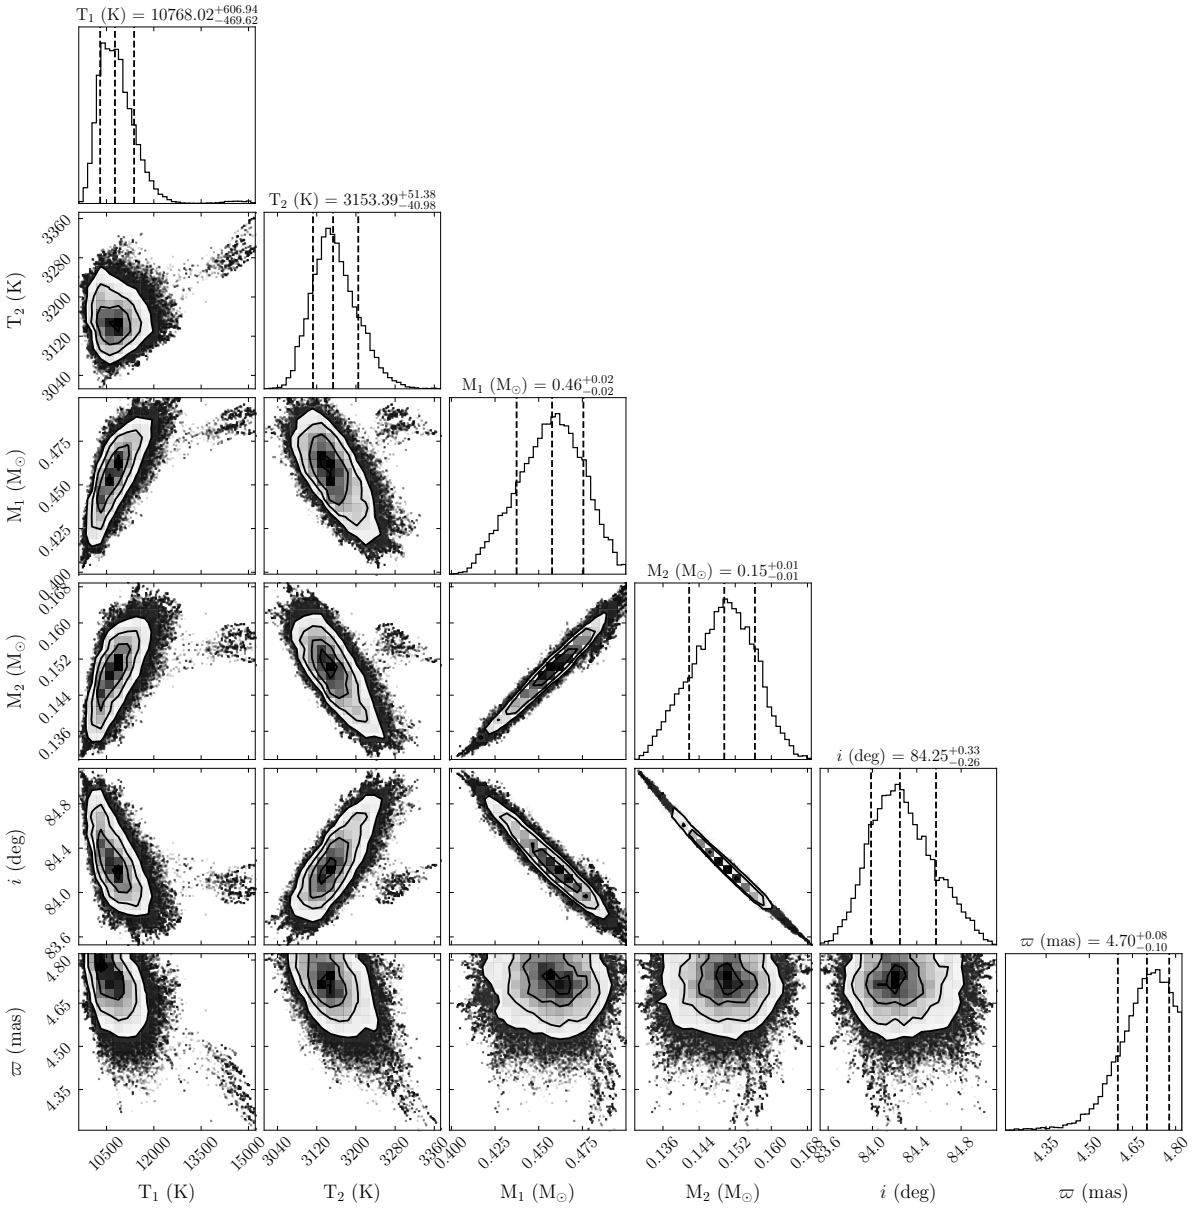

Figure D.31: As in Figure D.1 but for ZTF J1802−0054.

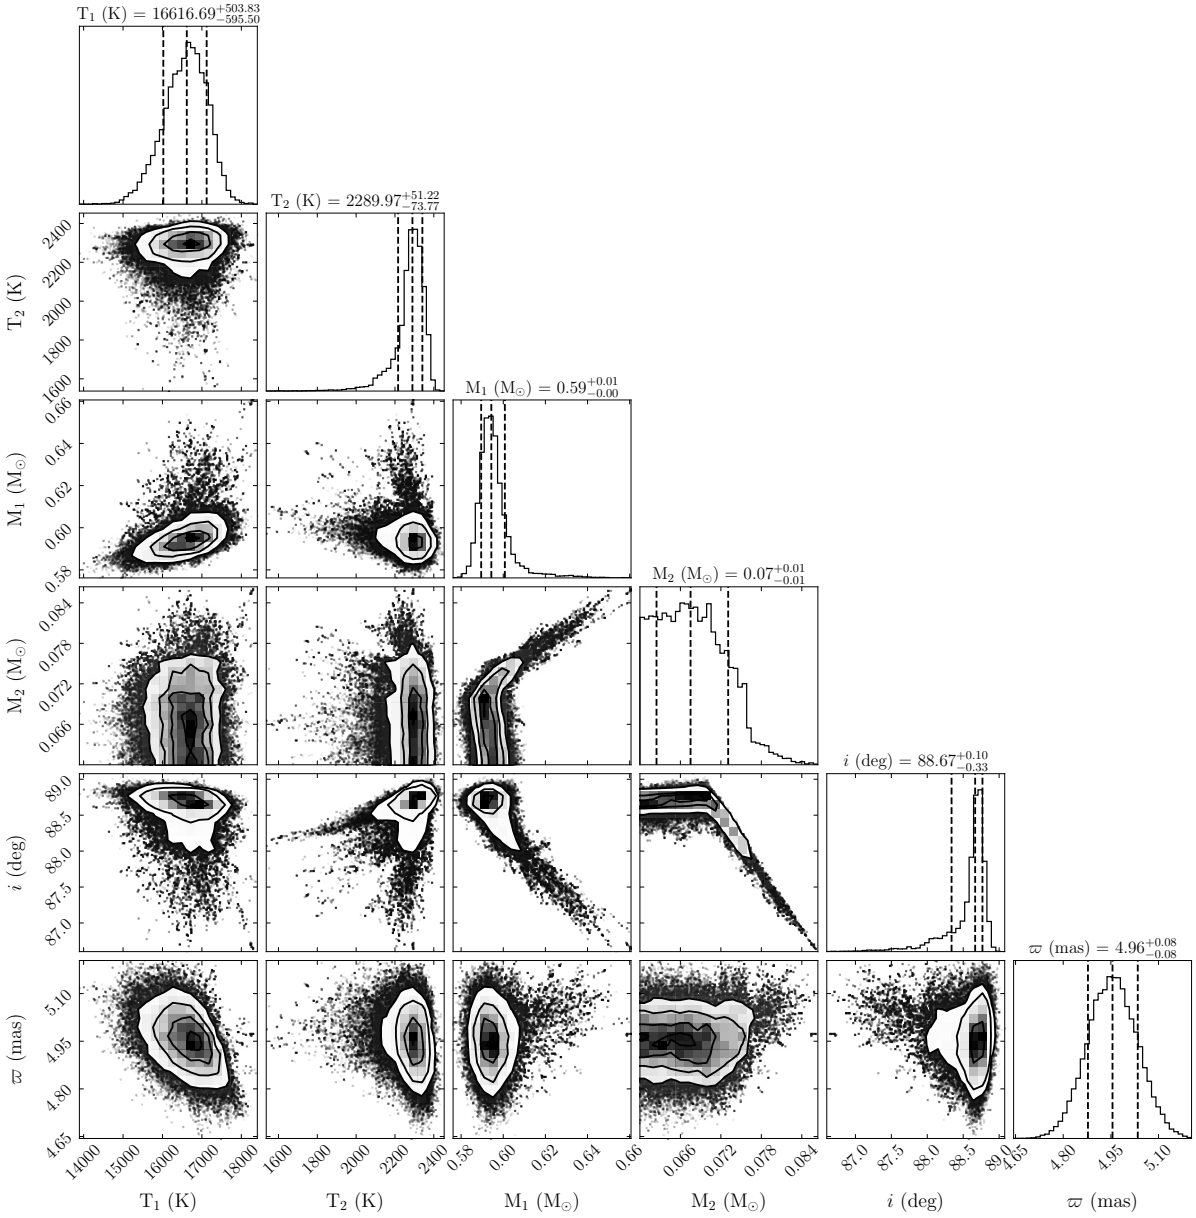

Figure D.32: As in Figure D.1 but for ZTF J1828+2308.

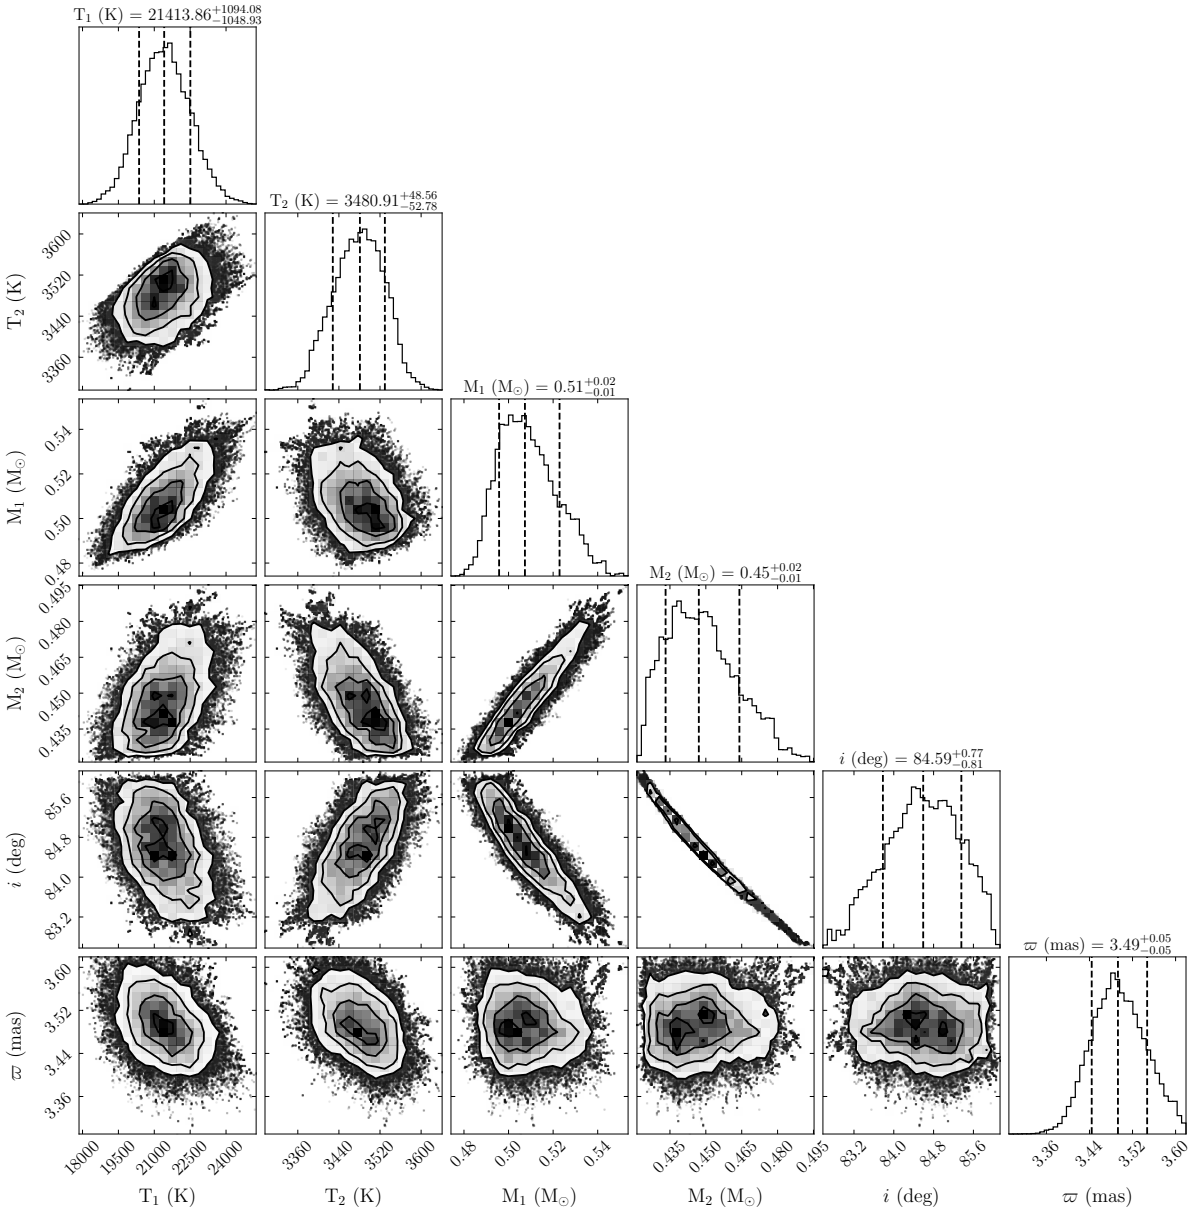

Figure D.33: As in Figure D.1 but for ZTF J1954+1019.

“

And all that is now
And all that is gone
And all that’s to come
And everything under the sun is in tune
But the sun is eclipsed by the moon

”

— Eclipse - Pink Floyd