On The Evolution, Detection, and Characterization of Small Planets in the Habitable Zones of M Dwarfs

Rodrigo Luger

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2017

Reading Committee:

Rory Barnes, Chair

Eric Agol

Victoria Meadows

David Catling

Program Authorized to Offer Degree:
Astronomy

University of Washington

Abstract

On The Evolution, Detection, and Characterization
of Small Planets in the Habitable Zones of M Dwarfs

Rodrigo Luger

Chair of the Supervisory Committee:
Professor Rory Barnes
Astronomy

As the technology behind instrumentation in astronomy improves, so too does our ability to detect and characterize worlds outside our solar system. We are currently witnessing a revolution in exoplanet science: for the past three decades, the number of known planets orbiting other stars has grown exponentially, showing no signs of tapering off. We now know of dozens of small planets in the habitable zones of their stars, and this number is expected to grow with upcoming survey missions such as the Transiting Exoplanet Survey Satellite (TESS) and the PLAnetary Transits and Oscillations telescope (PLATO). Improving commensurately with our capacity to detect these planets is our ability to characterize them. Missions such as the James Webb Space Telescope (JWST) and subsequent generations of space-based telescopes will be capable of characterizing these planets' atmospheres and searching for molecular signatures of habitability and life. Given the large number of potentially habitable planets we will soon discover, knowing which targets to prioritize for follow-up observations is paramount to furthering our goal of understanding the potential for habitability of exoplanets. Once data becomes available, its interpretation will rely heavily on a physical understanding of the processes that contribute to making a planet habitable (or not). Models of the evolutionary processes of potentially habitable planets can therefore improve target selection for biosignature searches and enhance the science return from terrestrial planet characterization.

In this dissertation, I develop theoretical models of the evolution of the atmospheres and

surface water inventories of planets in the habitable zones of low mass stars. While these stars currently offer the best opportunity to characterize potentially habitable planets, my work shows that vigorous atmospheric escape from these planets due to intense stellar activity could render many of them uninhabitable. I discuss observational signatures of the escape process and best case scenarios for planets around low mass stars, including the possibility that planets that form with substantial primordial atmospheres of hydrogen and helium could weather the active phase of the host star without substantial devolatilization.

I also refine existing techniques to detect and characterize exoplanets, with particular emphasis on small planets in the habitable zones of low mass stars. I introduce `EVEREST`, a pipeline to remove instrumental noise from photometric datasets and enable the detection of planet transit signals that would otherwise be hidden in the noise. Furthermore, I develop two novel techniques for the detection and characterization of potentially habitable exoplanets: the exo-auroral method, which relies on the spectroscopic detection of auroral emission from terrestrial planets, and planet-planet occultations, wherein an exoplanet occults another planet in the same system, imparting a small photometric signal on the system's light curve. I show how the next generation of telescopes may enable the application of both techniques to planets in the habitable zones of low mass stars, uncovering detailed information about their orbits and surface/atmospheric properties.

I discuss all of my results in the context of TRAPPIST-1, a nearby low mass star hosting seven transiting planets, three of which are in the habitable zone. This and similar soon-to-be discovered systems will likely revolutionize our understanding of exoplanets, habitability, and astrobiology in general.

Contents

List of Figures

List of Tables

Acknowledgments

The work I present here would not be possible without the support and guidance of many people—faculty, graduate students, friends, and family, all of whom believed in me and taught me to believe in myself.

In particular, I would like to thank Rory Barnes and Vikki Meadows for their constant support and for always having my back through the ups and downs of graduate school. Their care and insight over the past five years have proven invaluable, and I owe so many of my accomplishments to their unfaltering guidance. I also want to thank Eric Agol, who entertained my crazy ideas and helped me turn them into papers, and who was always available and willing to help with his keen insight into all sorts of astronomy problems. A big thanks to David Catling for his advice over the years and for agreeing to be on my dissertation committee (alongside Rory, Vikki, and Eric), and to Andy Becker, who introduced me to data processing and *Kepler* data and (more recently than I'm willing to admit) how to use a Linux terminal. And of course thank you to Dan Foreman-Mackey, who taught me to overcome my fear of matrices and embrace the gospel of Bayesian statistics.

I would also like to thank the graduate students at UW, in particular Jacob Lustig-Yaeger and David Fleming for all our crazy idea brainstorming sessions, many over matches of foosball; Russell Deitrick and Michael Tremmel for letting me pick their brains so many times and on so many different things over the past five years; Phoebe Upton Sanderbeck for all the hikes and trips to IHOP that got me away from the office when things got stressful; and all the other grads who made my time at UW so much fun and supported me through thick and thin. I would also like to thank Brooke Hubbard for her love and support and Van and Cici for their meows and cuddles that kept me sane while writing this dissertation.

And finally, I would like to thank my parents, Peter and Rose, for their love, guidance, support, and friendship, and for believing in me and helping me—morally, mentally, financially, and in so many other ways—to get to where I am today.

Dedication

To Mom & Dad

Introduction

Introduction

The Tip of the Iceberg

The end of my first year in graduate school—2013—marked the 50th anniversary of the discovery of the planetary system around Barnard's star (van de Kamp, 1963), comprised of two mutually inclined sub-Jupiter size planets in 12- and 26-year orbits around a nearby M dwarf, the first exoplanet system to be discovered (van de Kamp, 1969). Incidentally, 2013 was also the year of the most conclusive *undiscovery* of those planets (Choi et al., 2013). While Choi et al. (2013) definitively ruled out the existence of van de Kamp's planets with high precision radial velocity monitoring of the star, the astronomy community had long regarded the initial detection as spurious (e.g., Gatewood and Eichhorn, 1973; Benedict et al., 1998; Kürster et al., 2003). In fact, the putative planetary system around Barnard's star has been the butt end of jokes for quite some time. A few years ago, at a conference talk discussing the large occurrence rate of exoplanets around low mass stars inferred from *Kepler* data, I heard an astronomer quip that "every star in the sky except for Barnard's star has a planet in orbit around it."

It is quite unfortunate that Peter Van de Kamp's legacy is the Barnard's star affair. Van de Kamp was a brilliant astronomer, a pioneer in the field of astrometry, and the well-respected telescope director at the nation's top liberal arts college[1]. However, there are three lessons to be learned from it. In one way or another, these lessons, taken together, capture the essence of this dissertation.

First, regarding the quote above: according to the latest occurrence rates, low mass stars host on average 2.5 small planets in orbits within 200 days (Dressing and Charbonneau, 2015). It is therefore more likely than not that Barnard's star—a low mass M dwarf and the second closest neighbor to the Sun—hosts a planetary system, along with most of the M dwarfs in the galaxy. It was only in 1995 that the first planet was conclusively discovered around a main-sequence star (Mayor and Queloz, 1995), but since then the number of known exoplanets has grown exponentially: as of this writing, we know of 2,950 confirmed exoplanets, with another 2,337 candidates in the works[2]. And the exponential rate of discoveries shows no signs of slowing down: upcoming missions

[1] Swarthmore College, this author's beloved alma mater.
[2] Numbers taken from `exoplanets.org` (Han et al., 2014).

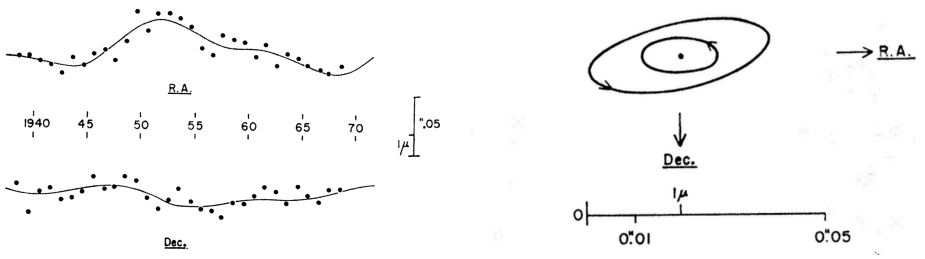

Figure 1: Astrometric measurements of Barnard's star from photographic plates by Peter van de Kamp from 1938–1969, reproduced from van de Kamp (1969). The left panel shows the data in right ascension (top) and declination (bottom) along with the model fit for the two planets. The panel on the right shows their sky-projected orbits. ©*AAS. Reproduced with permission.*

such as TESS and PLATO are expected to unveil several thousands more in the next few years (Sullivan et al., 2015; Rauer et al., 2014). While van de Kamp may have been wrong about his particular planetary system, he was keenly aware of this imminent revolution:

> *Could this be the first 'tip (or tips) of the iceberg'? How many planetary systems may be ready for discovery, and are patiently awaiting astrometric studies, by conventional photographic or other techniques?* — van de Kamp (1986)

Second, unbeknownst to van de Kamp at the time, M dwarfs offer the best opportunities to both detect and characterize potentially habitable exoplanets. M dwarfs are a class of stars smaller than the Sun, with masses ranging from about 0.08 $M_\odot$—the minimum mass for hydrogen fusion to take place at a star's core—to about 0.6 $M_\odot$. M dwarfs are also extremely faint; so faint, in fact, that even though they are the most common stars in the galaxy, *none* can be seen with the naked eye[3]. Because of their low luminosities, the habitable zone (HZ)—the region around a star where liquid water is possible on the surface of a terrestrial planet—is closer in, which increases the detectability of planets therein (e.g., Baraffe et al., 1998; Reid and Hawley, 2005). In particular, it increases the geometric probability that the planet will transit the face of its star and be photometrically detectable. In contrast to planets in the HZ of Sun-like stars, whose orbital periods (like Earth's) are several hundreds of days, planets in the HZ of M dwarfs have orbital periods ranging from just a couple to a few tens of days. In any given time interval, they will thus be seen to transit ten to one hundred times more than a corresponding planet around a

[3]Barnard's star, the fourth closest star to the Sun, is about six magnitudes fainter than the faintest star visible to the naked eye.

Sun-like star. Moreover, the smaller sizes of M dwarfs—at the lower end, these are barely larger than Jupiter—mean that during a transit, a terrestrial planet will block a much larger fraction of the star's light, resulting in a larger photometric signature. On the characterization front, M dwarf planets are similarly more amenable to being studied with transit transmission or secondary eclipse spectroscopy, for the same reason: a smaller star means a larger relative signal due to the planet. This is a point that I will return to throughout this dissertation: M dwarf planets are currently our best shot for understanding potentially habitable exoplanets, so we must tirelessly pursue them on both observational and theoretical fronts.

Third, and finally, the spurious detection of the planetary system around Barnard's star was almost certainly due to unmodeled instrumental noise. Van de Kamp used observations made with photographic plates to infer the astrometric motion of Barnard's star, claiming variations in the position of the star on the plate at the level of a few microns. Disassembly of the telescope for cleaning in the 1940s—halfway through the astrometric dataset—changed its focus, leading to an offset in subsequent measurements at that level (Kent, 2001). Additional instrumental effects contributed to what appeared to be astrometric signals (Hershey, 1973). In fact, van de Kamp's astrometric solution based on the original dataset suggested a single companion 1.6 times the mass of Jupiter (van de Kamp, 1963); the presence of a second companion had to be invoked six years later after the collection of more data (van de Kamp, 1969). Additional observations required even more changes to the orbital solutions (van de Kamp, 1975). In hindsight, this was a clear red flag suggesting that van de Kamp was merely fitting noise. Fifty years later, as we continue the search for planets around nearby stars, it is more important than ever that we understand the sources of noise in our datasets and learn to properly account for and remove them. This will not only ensure we are not tricked into detecting signals that are not there, but if done carefully, can also increase our sensitivity to smaller signals, pushing the frontiers of exoplanet detection and characterization towards our ultimate goal of detecting a habitable, Earth-like exoplanet. Again, van de Kamp said it best:

> *We shall always be confronted with threshold limitations, a natural condition of research, very much so of the present long-focus photographic astrometric approach. But should we not come to the rescue of a cosmic phenomenon which may be trying to reveal itself in a sea of errors?* — van de Kamp (1986)

Dissertation Overview

This dissertation focuses on the detection, characterization, and evolution of potentially habitable planets orbiting M dwarf stars. The topics I will cover below are quite diverse and are by no means a comprehensive analysis of all that we know about these planets and their detectability. Rather, I focus on specific topics within the overarching theme of M dwarf habitability, and connect them with applications to specific planetary systems. The dissertation is divided into three parts, described below.

Part I: The Evolution of Small Planets in the Habitable Zones of M Dwarfs

Part I concerns topics in the evolution of planets in the HZs of M dwarfs. Modeling the physical processes that shape planets in the HZ is paramount to understanding their present-day potential for habitability and guiding efforts to select targets for detailed follow-up observations. As efforts to detect and characterize these planets improve, it is more important than ever that we have a solid physical understanding of the evolutionary pathways these planets undergo so that observations can be correctly interpreted and placed in context. It is now clear that the environment in which M dwarf planets form and evolve is drastically different from that of Earth or other planets orbiting Sun-like stars. Planets around M dwarfs are generally believed to form more quickly; those that form *in situ* may also be smaller and drier than those orbiting Sun-like stars (Lissauer, 2007; Raymond et al., 2007). Recent evidence based on results from the *Kepler* mission suggests the planet formation process around these stars is quite different, typically resulting in the formation of larger numbers of smaller planets, which is at odds with the predictions of classic planet formation models (Mulders et al., 2015). Once formed, these planets also undergo markedly different evolution, as they are exposed to strong high energy photon and particle radiation from their magnetically active host stars (e.g., Scalo et al., 2007; Segura et al., 2010) and may experience strong tidal forces that can shape their orbits and even lead to runaway heating of the planet surface (Barnes et al., 2013). In this part of my dissertation, I focus on yet another important difference: the extended pre-main sequence phase of M dwarfs, which causes elevated stellar fluxes in the habitable zone for up to several hundreds of Myr. I show how this can drive terrestrial planets into runaway greenhouses, during which their oceans boil and can quickly escape to space, and how it can lead gas-rich planets to shed their envelopes and transform into volatile-rich planets

with solid or liquid water surfaces. I discuss how these vigorous atmospheric escape processes can either severely limit or potentially enhance these planets' potential for habitability. I finally argue that a detailed understanding of the history of atmospheric escape on M dwarf planets is crucial to selecting the best targets for future characterization efforts and interpreting the results of these observations.

Part II: The Detection of Planets in the Habitable Zones of M Dwarfs via the Transit Method

Part II discusses the detection of planets via the transit method, with application to TRAPPIST-1, a nearby M dwarf hosting three terrestrial-size planets in the habitable zone. I introduce `EVEREST`, a photometric pipeline to remove instrumental noise from observations made with *K2*, the repurposed *Kepler* spacecraft. Because of two failed reaction wheels, the spacecraft is unable to achieve the fine pointing precision required for the detection of small transiting planets. Software-based techniques such as `EVEREST` are needed to de-trend the raw data so that it can be used to search for and characterize exoplanets. I show how `EVEREST` light curves enable the refinement of exoplanet parameters derived from light curves and the detection of small exoplanets whose transit signals are below the noise floor of other *K2* pipelines. This is particularly important for terrestrial planets in the habitable zone, whose transit signals are small. I apply my results to TRAPPIST-1, using its *K2* light curve to study the rich chain of planet resonances in the system and confirm the existence of its outermost planet, TRAPPIST-1h, which I find is likely too cold to be habitable under a terrestrial $N_2/CO_2/H_2O$ atmosphere. Finally, while I discuss the implementation of `EVEREST` for *K2*, the techniques I develop are not specific to this mission or even to the detection of exoplanets, but can be applied to refine the precision of data collected by any CCD photometer subject to detector sensitivity variations or other sources of instrumental noise.

Part III: Novel Techniques to Detect and Characterize Planets in the Zones of M Dwarfs

In the final part of this dissertation, I continue my work on the detection of small planets, but this time turn to novel methods that have until recently not been feasible due to the faintness of the signals involved. As van de Kamp put it, the sky is filled with "cosmic phenomen[a] which may be

trying to reveal [themselves] in a sea of errors." I discuss the production and detectability of auroral signals from terrestrial planets orbiting nearby M dwarfs, with particular emphasis on the claimed planet in the habitable zone of Proxima Centauri b, the closest star to the Sun (Anglada-Escudé et al., 2016). Because of the extreme magnetic activity of the star, auroral emission from this planet could be 100–10,000 times stronger than on Earth, a signal I show is potentially detectable with the next generation of telescopes. Such a detection would not only confirm the existence of the planet, but could be used to map out the three-dimensional structure of its orbit, its mass, and its atmospheric composition. I also discuss the detectability of planet-planet occultations (PPOs) in nearby multi-planetary systems, with application to TRAPPIST-1. I show how PPOs are extremely frequent and will soon be detectable by the James Webb Space Telescope and other next generation instruments. Like exo-aurorae, these can be leveraged to map out the three-dimensional structure of the planetary system, but can also be used to infer crude surface maps of planets, including those in the habitable zone.

Relevance to Astrobiology

One of the primary goals of the field of astrobiology is to "determine the distribution of habitable environments and life in the Universe" (2015 Astrobiology Strategy Document; Hays, 2015). The Strategy Document further highlights the importance of studies that "observe and characterize potentially habitable exoplanets." The three parts of this dissertation fulfill these and other objectives presented in Section 5 (*Identifying, exploring, and characterizing environments for habitability and biosignatures*) and Section 6 (*Constructing habitable worlds*) of the Strategy Document.

Part I sheds light on the evolution of planets in the habitable zone, outlining processes that can both help and hinder their habitability, addressing key points in Section 6.5.II (*What are the exogenic factors in the formation of a habitable planet?*) and Section 6.5.V (*How does habitability change through time?*) of the Document.

Part II refines current methods to detect exoplanets, with applications to small planets in or near the habitable zones of low mass stars. The techniques I develop specifically address Section 5.4.IV (*How can we identify habitable planets and search for life beyond the solar system?*). In the final chapter of Part II, I apply my detection techniques to the TRAPPIST-1 system, which hosts at least three potentially habitable planets.

Part III presents novel methods to characterize exoplanets, paying specific attention to the potentially habitable exoplanet Proxima Centauri b and the ones in the TRAPPIST-1 system. The tools I develop in this part further help address the points in Section 5.4.IV of the Strategy Document by enabling the characterization of these planets' atmospheres and surfaces and hence probing their potential for habitability.

As a whole, this dissertation tackles several of the big picture aspects of astrobiology beyond the Solar System, satisfying one of the principal objectives in the Strategy Document: "to understand how habitable worlds and environments form and evolve, better understand the range of parameters that influence habitability, and determine how to detect, confirm, and characterize habitable environments."

Part I:

The Evolution of Small Planets in the Habitable Zones of M Dwarfs

1 Overview

Portions of this chapter were originally published in collaboration with R. Barnes in the journal Astrobiology (Luger and Barnes, 2015, Astrobiology, ©2015 Mary Ann Liebert, Inc.) and in collaboration with R. Barnes, E. Lopez, J. Fortney, B. Jackson, and V. S. Meadows in the journal Astrobiology (Luger et al., 2015, Astrobiology, ©2015 Mary Ann Liebert, Inc.), and are reproduced below with the permission of Mary Ann Liebert, Inc.

The first part of this thesis concerns the evolution of planets in the habitable zones of M dwarfs. As I mentioned previously, these low mass stars currently offer the best opportunity for the detection and characterization of potentially habitable planets, and yet the true potential of their planets to host life is all but unknown. In fact, recent studies suggest planets orbiting M dwarfs must deal with a whole slew of challenges that could preclude their habitability, including potentially dry formation conditions (e.g., Lissauer, 2007), prolonged exposure to high energy photon and particle irradiation (e.g., Scalo et al., 2007; Segura et al., 2010), and detrimental tidal heating of the surface (Barnes et al., 2013). Understanding these challenges in detail is therefore critical to the field of astrobiology, especially as we begin to probe these planets' atmospheres in search of habitability markers.

In this section, I review aspects of stellar evolution, atmospheric evolution, and tidal evolution relevant to planets orbiting M dwarfs. I define the habitable zone and its evolution in time and discuss the problems that all of these processes pose to habitability. Next, in Chapter 2, I discuss how vigorous atmospheric escape could lead to the loss of a terrestrial planet's surface water inventory and the buildup of oxygen in its atmosphere, turning the planet into a "mirage" Earth. In Chapter 3 I adopt a more optimistic view: the same atmospheric escape processes that can desiccate terrestrial planets could also remove the gaseous envelopes from gas-rich "super-Earths" and "mini-Neptunes"[4], which I argue could form potentially habitable evaporated cores.

[4]The terms "super-Earth" and "mini-Neptune" have rather loose definitions and are often used to refer to the same kind of planet: a planet with a mass in between that of Earth and Neptune, of which there is no analogue in our solar system. In general, "super-Earth" connotes a primarily terrestrial planet–in some sense, a scaled-up version of Earth. By contrast, a "mini-Neptune" suggests a largely gaseous body and thus a scaled-down version of Neptune; in this and subsequent chapters, I adopt this convention when referring to these planets. While these planets are extremely common in the sample of transiting exoplanets discovered by *Kepler*, their actual compositions are not yet well understood.

1.1 Stellar Evolution

Following their formation from a giant molecular cloud, stars contract under their own gravity until they reach the main sequence (MS), at which point the core temperature is high enough to ignite hydrogen fusion (e.g., Hayashi, 1961). While the Sun is thought to have spent $\lesssim 50$ Myr in this pre-main sequence (PMS) phase (Baraffe et al., 1998), M dwarfs take hundreds of Myr to fully contract; the lowest mass M dwarfs reach the main sequence only after ~ 1 Gyr (e.g., Reid and Hawley, 2005). During their contraction, M dwarfs remain at a roughly constant effective temperature (Hayashi, 1961), so that their luminosity is primarily a function of their surface area, which is significantly larger than when they arrive on the main sequence. M dwarfs therefore remain super-luminous for several hundred Myr, with total (bolometric) luminosities higher than the main sequence value by up to two orders of magnitude. As I discuss in later sections, this will significantly affect the atmospheric evolution of any planets these stars may host.

XUV emissions ($1\text{Å} \lesssim \lambda \lesssim 1000\text{Å}$) from M dwarfs also vary strongly with time (West et al., 2008). This is because the XUV luminosity of M dwarfs is rooted in the vigorous magnetic fields generated in their large convection zones (Scalo et al., 2007). Stellar magnetic fields are largely controlled by rotation (Parker, 1955), which slows down with time due to angular momentum loss (Skumanich, 1972), leading to an XUV activity that declines with stellar age. However, due to the difficulty of accurately determining both the XUV luminosities (usually inferred from line proxies) and the ages (often determined kinematically) of M dwarfs, the exact functional form of the evolution is still very uncertain (for a review, see Scalo et al., 2007). Further complications arise due to the fact that the process(es) that generate magnetic fields in late M dwarfs may be quite different from those in earlier type stars. The rotational dynamo of Parker (1955) involves the amplification of toroidal fields generated at the radiative-convective boundary; late M dwarfs, however, are fully convective, and have no such boundary. Instead, their magnetic fields may be formed by turbulent dynamos (Durney et al., 1993), which may evolve differently in time from those around higher mass stars (Reid and Hawley, 2005).

In a comprehensive study of the XUV emissions of solar-type stars of different ages, Ribas et al. (2005) found that the XUV evolution is well modeled by a simple power law with an initial

short-lived "saturation" phase:

$$\frac{L_{\rm XUV}}{L_{\rm bol}} = \begin{cases} \left(\frac{L_{\rm XUV}}{L_{\rm bol}}\right)_{\rm sat} & t \leq t_{\rm sat} \\ \left(\frac{L_{\rm XUV}}{L_{\rm bol}}\right)_{\rm sat} \left(\frac{t}{t_{\rm sat}}\right)^{-\beta} & t > t_{\rm sat}, \end{cases} \tag{1}$$

where $L_{\rm XUV}$ and $L_{\rm bol}$ are the XUV and bolometric luminosities, respectively, and $\beta = -1.23$. Prior to $t = t_{\rm sat}$, the XUV luminosity is said to be "saturated," as observations show that the ratio $L_{\rm XUV}/L_{\rm bol}$ remains relatively constant at early times.

Jackson et al. (2012) find that $t_{\rm sat} \approx 100\,{\rm Myr}$ and $(L_{\rm XUV}/L_{\rm bol})_{\rm sat} \approx 10^{-3}$ for K dwarfs. Similar studies for M dwarfs, however, are still being developed (e.g., Engle and Guinan, 2011), but it is likely that the saturation timescale is much longer for late-type M dwarfs. Wright et al. (2011) show that X-ray emission in low mass stars is saturated for $P_{\rm rot}/\tau \lesssim 0.1$, where $P_{\rm rot}$ is the stellar rotation period and τ is the convective turnover time. The extent of the convective zone increases with decreasing stellar mass; below $0.35{\rm M}_\odot$, M dwarfs are fully convective (Chabrier and Baraffe, 1997), resulting in larger values of τ (see, e.g., Pizzolato et al., 2000). Low mass stars also have longer spin-down times (Stauffer et al., 1994); together, these effects should lead to significantly longer saturation times compared to solar-type stars. This is consistent with observational studies; West et al. (2008) find that the magnetic activity lifetime increases from $\lesssim 1\,{\rm Gyr}$ for early (i.e., most massive) M dwarfs to $\gtrsim 7\,{\rm Gyr}$ for late (least massive) M dwarfs, possibly due to the onset of full convection around spectral type M3. Finally, Stelzer et al. (2013) find that for M dwarfs, $\beta = -1.10 \pm 0.02$ in the X-ray and $\beta = -0.79 \pm 0.05$ in the FUV (far ultraviolet), suggesting a slightly shallower slope in the XUV compared to the value from Ribas et al. (2005).

1.2 The Habitable Zone

The habitable zone (HZ) is classically defined as the region around a star where an Earth-like planet can sustain liquid water on its surface (Hart, 1979; Kasting et al., 1993). Interior to the HZ, high surface temperatures lead to the runaway evaporation of a planet's oceans, which increases the atmospheric infrared opacity and reduces the ability of the surface to cool in a process known as the *runaway greenhouse* (Ingersoll, 1969; Kasting, 1988; Nakajima et al., 1992). Surface temperatures rise, causing more water evaporation and higher infrared opacity, until the temperature exceeds the critical point of water and the atmosphere becomes predominantly water vapor. Eventually,

the planet heats up to the point that its blackbody emission peaks in the near infrared, where atmospheric windows allow the planet to cool and establish radiation balance at a balmy ∼1500 K.

Exterior to the HZ, greenhouse gases—gases, like water vapor, that absorb strongly in the infrared—are unable to warm the planet. The elevated abundances of these gases needed to warm the surface via greenhouse heating actually lead to increased Rayleigh scattering, which prevents the downwelling stellar flux from reaching the surface in the first place. At a certain distance from the star, these gases are therefore unable to maintain surface temperatures above the freezing point, regardless of their abundance. Recently, Kopparapu et al. (2013) re-calculated the location of the HZ boundaries as a function of stellar luminosity and effective temperature using an updated one-dimensional, radiative-convective, cloud-free climate model. Their five boundaries are the (1) Recent Venus, (2) Runaway Greenhouse, (3) Moist Greenhouse, (4) Maximum Greenhouse, and (5) Early Mars habitable zones.

The first and last limits can be considered "optimistic" empirical limits, since prior to ∼ 1 and ∼ 3.8 Gyr ago, respectively, Venus and Mars may have had liquid surface water. The ability of a planet to maintain liquid water and to sustain life at these limits is still unclear and probably depends on a host of properties of its climate. Conversely, the other three limits are the "pessimistic" theoretical limits, corresponding to where a cloud-free, Earth-like planet would lose its entire water inventory due to the greenhouse effect (2 and 3) and to where the addition of CO_2 to the atmosphere would be incapable of sustaining surface temperatures above freezing (4).

It is important to note, however, that considerable uncertainty exists on the limits discussed above. Classical calculations of the HZ boundaries (Kasting, 1988; Kopparapu et al., 2013, 2014) rely on one-dimensional models tailored to reproduce conditions on Earth and are unable to capture changes in the atmospheric circulation and cloud formation mechanisms that occur as a planet begins to warm. Recent studies such as those of Abe et al. (2011), Leconte et al. (2013), Yang et al. (2014), and Kopparapu et al. (2016) used global climate models to show that the threshold for a runaway greenhouse can be significantly higher than that predicted by 1D models and may be quite sensitive to factors such as the planet's rotation rate and surface water content.

Moreover, because stellar luminosities vary with time, the location of the HZ is not fixed. While for K and G dwarfs the change in the location of the HZ is modest during the pre-main sequence

phase, the HZ of M dwarfs moves in substantially during that time, changing by nearly an order of magnitude for the least massive stars. Due to this evolution, planets observed in the HZ of M dwarfs today likely spent a significant amount of time interior to the inner edge of the HZ, provided they either formed *in situ* or migrated to their current positions relatively early. I discuss the implications of this point in subsequent chapters.

Finally, it is important to note that the location of the HZ is also a function of the orbital eccentricity e and the planet mass. The former is due to the fact that at a fixed semi-major axis a, the orbit-averaged flux $\langle \mathcal{F} \rangle$ is higher for more eccentric orbits (Williams and Pollard, 2002):

$$\langle \mathcal{F} \rangle = \frac{L_{\rm bol}}{4\pi a^2 \sqrt{1-e^2}}. \tag{2}$$

The dependence on mass is due to the fact that the pressure at the emission level of a saturated atmosphere scales with surface gravity; at a higher pressure, the temperature of the emission level is higher and the planet is able to cool more effectively, delaying the runaway state (Pierrehumbert, 2010).

1.3 Atmospheric Escape

Planetary atmospheres constantly evolve. Several mechanisms exist through which planets can lose significant quantities of their atmospheres to space, leading to dramatic changes in composition and in some cases complete atmospheric erosion (for a review, see Catling and Kasting, 2017, Chapter 5). The early Earth itself could have been rich in hydrogen, with mixing ratios as high as 30% in the prebiotic atmosphere (Tian et al., 2005; Hashimoto et al., 2007). A variety of processes subsequently led to the loss of most of this hydrogen; Watson et al. (1981) argue that on the order of 10^{23} g ($\sim 10^{-5} M_\oplus$) of hydrogen could have escaped in the first billion years. Similarly, Kasting and Pollack (1983) calculated that early Venus could have lost an Earth ocean equivalent of water in the same amount of time. Currently, observational evidence for atmospheric escape exists for two "hot Jupiters," HD 209458b (Vidal-Madjar et al., 2003) and HD 189733b (Lecavelier Des Etangs et al., 2010), and one "hot Neptune," GJ 436b (Kulow et al., 2014), whose proximity to their parent stars leads to the rapid hydrodynamic loss of hydrogen.

Atmospheric escape mechanisms fall into two major categories: thermal escape, in which the heating of the upper atmosphere accelerates the gas to velocities exceeding the escape velocity, and

nonthermal escape, which encompasses a variety of mechanisms and may involve energy exchange among ions or erosion due to stellar winds. While nonthermal processes certainly do play a role in the atmospheric escape of planets in the HZ, the high exospheric temperatures resulting from strong XUV irradiation probably make thermal escape the dominant mechanism for planets around M dwarfs at early times (Kislyakova et al., 2013). However, the escape can be greatly enhanced by flares and coronal mass ejections, which can completely erode the atmosphere of a planet lacking a strong magnetic field (Lammer et al., 2007; Scalo et al., 2007). For a review of the nonthermal mechanisms of escape, the reader is referred to Hunten (1982).

In this section I discuss the escape of two different but closely related types of atmospheres: hydrogen-dominated atmospheres, like those of mini-Neptunes and gas-rich super-Earths, and water vapor-dominated atmospheres, which one may expect for a water-rich terrestrial planet undergoing a runaway greenhouse.

1.3.1 Escape of Hydrogen

Below I review three regimes of thermal escape relevant to hydrogen-rich atmospheres: Jeans escape, energy-limited escape, and radiation/recombination-limited escape. While I focus on the escape of hydrogen atoms, the formalism described below applies in principle to any atmosphere whose dominant species is the one escaping.

1.3.1.1 Jeans Escape

In the low temperature limit, atmospheric mass loss occurs via the ballistic escape of individual atoms from the collisionless exosphere, where the low densities ensure that atoms with outward velocities exceeding the escape velocity will escape the planet. Originally developed by Jeans (1925), who assumed a hydrostatic, isothermal atmosphere, the mass loss rate of a pure hydrogen atmosphere is given by (Öpik, 1963)

$$\dot{M}_{\rm J} = 4\pi R_{\rm exo}^2 n m_H v_t \frac{(1+\lambda_J)e^{-\lambda_J}}{2\sqrt{\pi}} \tag{3}$$

where $R_{\rm exo}$ is the radius of the exobase, n is the number density of hydrogen atoms at the exobase, m_H is the mass of a hydrogen atom, v_t is the thermal velocity of the gas, and λ_J is the Jeans escape parameter, defined as the ratio of the potential energy to the kinetic energy of the gas and

given by

$$\lambda_J \equiv \frac{GM_p m_H}{kT_{\rm exo}R_{\rm exo}}, \tag{4}$$

where G is the gravitational constant, M_p is the mass of the planet, and $T_{\rm exo}$ is the temperature in the (isothermal) exosphere. Since in the Jeans regime the thermal velocity of the gas is less than the escape velocity, the bulk of the gas remains bound to the planet, and only atoms in the tail of the Maxwell-Boltzmann distribution are able to escape. Jeans escape is thus slow. As an example, the present Jeans escape flux for hydrogen on Venus is on the order of 10 atoms $\mathrm{cm}^{-2}\mathrm{s}^{-1}$ (Lammer et al., 2006), corresponding to the feeble rate of $\sim 10^{-4}$ g/s, which would remove only one part in 10^{11} of the atmosphere per billion years.

For higher exospheric temperatures and/or larger values of $R_{\rm exo}$, however, corresponding to low values of λ_J, the atmosphere enters a regime in which the velocity of the average atom in the exosphere approaches the escape velocity of the planet. In this regime, the bulk of the upper atmosphere ceases to be hydrostatic (and isothermal), (4) is no longer valid, and the escape rate must be calculated from hydrodynamic models.

1.3.1.2 Hydrodynamic Escape

One of the primary mechanisms for heating the exosphere and decreasing λ_J is via strong XUV irradiation. XUV photons are absorbed close to the base of the thermosphere, where they deposit energy and heat the gas via ionization (usually of atomic hydrogen). This heating is balanced by the adiabatic expansion of the upper atmosphere, downward heat conduction, and cooling by recombination radiation. For sufficiently high XUV fluxes, the expansion of the atmosphere accelerates the gas to supersonic speeds, at which point a hydrodynamic wind is established similar to the solar Parker wind (Parker, 1965). Once the gas reaches the exosphere, it will escape the planet provided its kinetic energy exceeds the energy required to lift it out of the planet's gravitational well.

Since the kinetic energy of a hydrogen atom at the exobase is $\frac{3}{2}kT_{\rm exo}$, the condition $\lambda_J < 1.5$ implies that the kinetic energy of the gas is greater than the absolute value of its gravitational binding energy, and it should therefore begin to escape in bulk in a process commonly referred to as "blow-off." Unlike in the Jeans escape regime, where the mass loss occurs on a per-particle basis,

blow-off leads to the rapid loss of large portions of the upper atmosphere, irrespective of particle species, as atoms and molecules heavier than hydrogen are carried along by the hydrodynamic wind. However, contrary to what Öpik (1963) suggests, the mass loss in this stage is not *arbitrarily* high, since once blow-off begins the upper atmosphere can no longer be treated as isothermal. As the exosphere expands it also cools, so that in the absence of an energy source the value of λ_J will tend to increase, thereby moderating the blow-off. The mass loss is, in this sense, "energy-limited," and may be calculated by equating the energy input to the energy required to drive the escape.

Originally proposed by Watson et al. (1981), the energy-limited model assumes that the XUV flux is absorbed in a thin layer at radius $R_{\rm XUV}$ where the optical depth to stellar XUV photons is unity. Recently updated to include tidal effects by Erkaev et al. (2007), this model approximates the mass loss as

$$\dot{M}_{\rm EL} \approx \frac{\epsilon_{\rm XUV} \pi \mathcal{F}_{\rm XUV} R_p R_{\rm XUV}^2}{G M_p K_{\rm tide}(\xi)} \tag{5}$$

where $\epsilon_{\rm XUV}$ is the heating efficiency parameter (see below), $\mathcal{F}_{\rm XUV}$ is the incident XUV flux, R_p is the planetary radius, and $K_{\rm tide}$ is a tidal enhancement factor, accounting for the fact that for sufficiently close-in planets, the stellar gravity reduces the gravitational binding energy of the gas such that it need only reach the Roche radius to escape the planet. Erkaev et al. (2007) show that

$$K_{\rm tide}(\xi) = \left(1 - \frac{3}{2\xi} + \frac{1}{2\xi^3}\right), \tag{6}$$

where the parameter ξ is defined as

$$\xi \equiv \frac{R_{\rm Roche}}{R_{\rm XUV}} \tag{7}$$

with

$$R_{\rm Roche} \equiv \left(\frac{M_p}{3M_\star}\right)^{1/3} a, \tag{8}$$

where $\rm M_\star$ is the mass of the star and a is the semi-major axis. The XUV absorption radius $R_{\rm XUV}$ is typically no more than 10–20% larger than the planet radius R_p, and is often taken to be equal to R_p for simplicity (Murray-Clay et al., 2009; Lopez et al., 2012).

Since the input XUV energy is balanced in part by cooling radiation (via Lyman α emission in the case of hydrogen) and heat conduction, only a fraction of it goes into the adiabatic expansion that drives escape. Rather than running complex hydrodynamic and radiative transfer models to determine the net heating rate, many authors (Jackson et al., 2010; Leitzinger et al., 2011; Lopez et al., 2012; Koskinen et al., 2013a; Lammer et al., 2013) choose to fold the balance between XUV heating and cooling into an efficiency parameter, $\epsilon_{\rm XUV}$, defined as the fraction of the incoming XUV energy that is converted into PdV work. Because of the complex dependence of $\epsilon_{\rm XUV}$ on the atmospheric composition and structure, its value is still poorly constrained. Lopez et al. (2012) estimate $\epsilon_{\rm XUV} = 0.2 \pm 0.1$ for super-Earths and mini-Neptunes based on values found throughout the literature. Earlier work by Chassefière (1996a) estimates $\epsilon_{\rm XUV} \lesssim 0.30$ for Venus-like planets but the author argues that the actual value may be closer to 0.15. Recently, Owen and Jackson (2012) found X-ray efficiencies $\lesssim 0.1$ for planets more massive than Neptune, but higher efficiencies (~ 0.15) for terrestrial planets. Moreover, Shematovich et al. (2014) argue that studies that assume efficiencies higher than about 0.2 probably lead to overestimates in the escape rate.

Unlike Jeans escape, hydrodynamic blow-off is fast. Chassefière (1996a) calculates the maximum hydrodynamic escape rate from the early Venusian atmosphere to be $\sim 10^6$ g/s, ten orders of magnitude higher than the present Jeans escape flux (see § 1.3.1.1). Although there has been debate over the validity of the blow-off assumption (see, for instance, the discussion in Tian et al., 2008), recently Lammer et al. (2013) showed that for super-Earths exposed to high levels of XUV irradiation, the energy-limited approximation yields mass loss rates that are consistent with hydrodynamic models to within a factor of about two, which is within the uncertainties of the problem.

1.3.1.3 Jeans Escape or Hydrodynamic Escape?

Since the location of the habitable zone is governed primarily by the total (bolometric) flux incident on a planet, the higher ratio of $L_{\rm XUV}$ to $L_{\rm bol}$ of M dwarfs implies a much larger XUV flux in the HZ compared to solar-type stars. The present-day solar XUV luminosity is $L_{\rm XUV}/L_{\rm bol} \approx 3.4 \times 10^{-6}$ (see Table 4 in Ribas et al., 2005), while for active M dwarfs this ratio is $\sim 10^{-3}$ (e.g., Scalo et al. 2007). Therefore planets in the HZ of M dwarfs likely experience XUV fluxes several orders of magnitude greater than the present Earth level ($F_{\rm XUV_\oplus} \approx 4.64\,{\rm erg/s/cm^2}$), especially at early times. Recent papers (Lammer et al., 2007, 2013; Erkaev et al., 2013) show that terrestrial planets

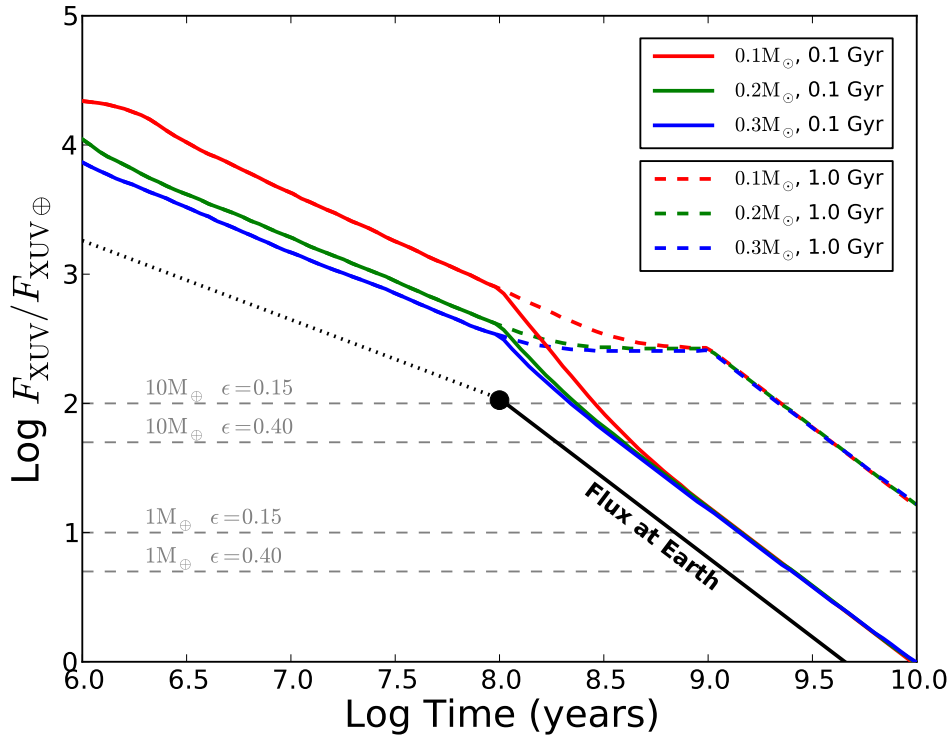

Figure 2: Evolution of the XUV flux received by planets close to the inner edge of the HZ (at 1 Gyr) for stars of mass 0.1, 0.2, and 0.3 $M_\odot$. Solid lines correspond to an XUV saturation time of 0.1 Gyr; dashed lines correspond to 1 Gyr. The flux at Earth is indicated by the black line. The dot corresponds to the earliest time for which Ribas et al. (2005) has data for solar-type stars; this is also roughly the time at which Earth formed. An XUV luminosity saturated at $10^{-3}L_{\rm bol}$ is roughly indicated by the dotted line. Finally, the dashed gray lines indicate the minimum XUV fluxes required to sustain blowoff according to the study of Erkaev et al. (2013). Super-Earths in the HZs of M dwarfs remain in the blowoff regime for at least a few 100 Myr; Earths undergo blowoff for much longer. For an XUV saturation time of 1 Gyr, blowoff occurs for several Gyr for all planets.

experiencing XUV fluxes corresponding to 10× and 100× $F_{\rm XUV_\oplus}$ are in the hydrodynamic flow regime—it is therefore likely that this is the case for planets in the HZ of active M dwarfs.

In Figure 2 I plot the evolution of the XUV flux received by a planet located close to the inner edge of the HZ (defined at 5 Gyr), for three different M dwarf masses and two different XUV saturation times (see §1.1). The dashed lines correspond to the critical fluxes in Erkaev et al. (2013) above which hydrodynamic escape occurs, for 1 and 10 $M_\oplus$ and two values of $\epsilon_{\rm XUV}$. Earth-mass planets remain in the hydrodynamic escape regime for at least 1 Gyr in all cases and in excess of 10 Gyr for active M dwarfs. The duration of this regime is shorter for 10 $M_\oplus$ planets,

but still on the order of several 100 Myr.

It is important to note that tidal effects can significantly increase the critical value of λ_J below which hydrodynamic escape occurs (see discussion in Erkaev et al., 2007). Hydrodynamic escape ensues when the thermal energy of the gas exceeds its potential energy, which occurs when

$$1.5 \geq \frac{GM_p K_{\rm tide} m_H}{kT_{\rm exo} R_{\rm exo}}$$

or

$$\begin{aligned} \lambda_J &\leq \frac{1.5}{K_{\rm tide}} \\ &\equiv \lambda_{\rm crit}, \end{aligned} \tag{9}$$

provided we maintain the original definition of the Jeans parameter (4). Due to the strong tidal forces acting on planets in the HZ of M dwarfs (see §1.4 below), this effect should greatly increase the critical value of the escape parameter, effectively reducing the value of $\mathcal{F}_{\rm XUV}$ for which hydrodynamic escape occurs.

1.3.1.4 Energy-Limited or Radiation/Recombination-Limited?

Hydrodynamic escape from planetary atmospheres need not be energy-limited. In the limit of high extreme ultraviolet (EUV) flux (low-energy XUV photons with 100Å $\lesssim \lambda \lesssim$ 1000Å), Murray-Clay et al. (2009) showed that the escape is "radiation/recombination-limited," scaling roughly as $\dot{M} \propto (\mathcal{F}_{\rm EUV})^{1/2}$. In this regime, the upper atmosphere thermostats to $T \sim 10^4$ K, photoionization is balanced by radiative recombination (as opposed to PdV work), and a large fraction of the gas energy budget is lost to Lyman α cooling. The mass loss rate is found by solving the mass continuity equation, yielding

$$\dot{M} = 4\pi r_s^2 \rho_s c_s \tag{10}$$

where r_s is the radius of the sonic point (where the wind velocity equals the local sound speed c_s) and ρ_s is the density at r_s. Since the bulk of the flow is ionized, the density is fixed by ionization-recombination balance, scaling roughly as $(\mathcal{F}_{\rm EUV} r_s)^{1/2}$. For a $0.7 M_{\rm J}$ hot Jupiter, the

radiation/recombination limited mass loss rate is (Murray-Clay et al., 2009)

$$\dot{M}_{\rm RR} \approx 4 \times 10^{12}\ {\rm g/s} \left(\frac{\mathcal{F}_{\rm EUV}}{5 \times 10^5\ {\rm erg/cm^2/s}} \right)^{1/2}. \tag{11}$$

Owen and Jackson (2012) re-derive this expression with explicit scalings on several planet properties:

$$\begin{aligned} \dot{M}_{\rm RR} \approx 2.4 \times 10^{11}\ {\rm g/s}\ (1+x)^{3/2}\, f_{\rm parker} \left(\frac{\Phi_\star}{10^{40}\ {\rm s^{-1}}} \right)^{1/2} \\ \times \left(\frac{a}{0.1\ {\rm AU}} \right) \left(\frac{R_p}{10\,{\rm R_\oplus}} \right)^{3/2} \left(\frac{A}{1/3} \right)^{1/2} \left(\frac{c_{\rm EUV}}{10\ {\rm km/s}} \right), \end{aligned} \tag{12}$$

where the quantity $(1+x)$ is the radius of the ionization front in units of R_p, $f_{\rm parker}$ is the Mach number of the flow, $\Phi_\star$ is the stellar EUV luminosity in photons per second, A is a geometrical factor and $c_{\rm EUV}$ is the isothermal sound speed of the gas. Taking $x = 0$, $f_{\rm parker} = 1$, $A = 1/3$, $c_{\rm EUV} = 10$ km/s, and an average EUV photon energy $h\nu = 20$ eV, this becomes

$$\dot{M}_{\rm RR} \approx 7.11 \times 10^7\ {\rm g/s} \left(\frac{\mathcal{F}_{\rm EUV}}{\rm erg/cm^2/s} \right)^{1/2} \left(\frac{R_p}{\rm R_\oplus} \right)^{3/2}. \tag{13}$$

The transition from energy-limited to radiation/recombination-limited escape is found by equating the two expressions, namely (5) and (13), and solving for the critical EUV flux. For terrestrial planets, the transition occurs above $\mathcal{F}_{\rm EUV} \sim 10^4\,{\rm erg/cm^2/s}$. Below this flux, the escape is energy-limited; above it, the escape is radiation/recombination-limited and thus increases more slowly with the flux. Planets in the HZ of M dwarfs are exposed to EUV fluxes up to an order of magnitude larger than this critical value. During this period, which lasts on the order of a few tens of Myr, the mass loss rate may be radiation/recombination-limited.

Nonetheless, whether the high-flux mass loss rate is more accurately described by (5) or (13) will depend on whether the flux is dominated by X-ray or EUV radiation. Owen and Jackson (2012) show that the mass loss rate for X-ray driven hydrodynamic winds scales linearly with the X-ray flux; this is because the sonic point for X-ray flows tends to occur *below* the ionization front. Even though recombination radiation still removes energy from the flow, it does so once the gas is already supersonic and thus causally decoupled from the planet, such that it cannot bottleneck the escape. Although the authors caution that the dependence of the mass loss rate on planet mass

and radius must be determined numerically, this regime is analogous to the energy-limited regime and can be roughly approximated by (5).

For X-ray luminosities $L_{\rm X} \gtrsim 10^{-6}\,{\rm L}_{\odot}$, close-in planets undergo X-ray driven hydrodynamic escape (see Figure 11 in Owen and Jackson, 2012). If X-rays contribute significantly to the XUV emissions of young M dwarfs, their X-ray luminosities may exceed this value for as long as 1 Gyr, and close-in planets will undergo energy-limited escape. Unfortunately, given the lack of observational constraints on the X-ray/EUV luminosities of young M dwarfs, it is unclear at this point whether the hydrodynamic escape will be EUV-driven (radiation/recombination-limited) or X-ray-driven (energy-limited) for planets in the HZ.

1.3.2 Escape of Water

Under the strong XUV irradiation of planets orbiting M dwarfs, water may also be lost via hydrodynamic escape. On Earth, the loss of surface water to space is inhibited by the presence of the cold trap, a temperature minimum in the lower atmosphere that causes water to condense and rain out before it ever reaches the stratosphere. In the absence of a cold trap, water reaching the stratosphere becomes susceptible to photolysis by stellar ultraviolet radiation, producing free hydrogen atoms, which as we saw above can readily escape to space.

However, as we discussed above, during a runaway greenhouse a planet's atmosphere is predominantly water vapor, even in the stratosphere. As a consequence, planets in runaway greenhouses can quickly lose large quantities of water to space via the photolysis of water and subsequent escape of hydrogen. As we will see in the next chapter, this is particularly problematic for terrestrial planets orbiting M dwarfs, given their prolonged and active pre-main sequence phase, which causes runaway greenhouses to last up to several hundred Myr.

Assuming neither photolysis nor diffusion are limiting factors, the mass escape rate from a water vapor atmosphere is given by the energy-limited or radiation/recombination-limited expressions presented in the previous section. However, the presence of photolytically-produced oxygen atoms can modify the escape rate in two important ways: (1) if the escape is vigorous enough, oxygen atoms can become entrained in the flow, a process that negatively feeds back on the particle escape flux due to the high mass of oxygen; and (2) as oxygen accumulates in the upper atmosphere, hydrogen may have to diffuse through a static background gas in order to escape, causing the flow

to become capped at the diffusion limit of the gas. In the next few sections, we derive escape rates accounting for the drag imposed by oxygen atoms on the escaping flow of hydrogen.

1.3.2.1 Hydrodynamic Drag

Strong hydrodynamic flows are capable of dragging heavier species along with them. In a water vapor-dominated atmosphere, photolytically-produced oxygen atoms can thus escape to space in the flow. Because oxygen atoms are much heavier than hydrogen atoms, the escape of the former tends to moderate the flow when the escape is energy-limited. Hunten et al. (1987) studied mass fractionation during hydrodynamic escape, demonstrating that an escaping species can efficiently drag a heavier species along with it provided the mass of the latter is smaller than the *crossover mass* m_c, equal to

$$m_c = m_H + \frac{kTF_H}{bgX_H}, \tag{14}$$

in the case of a background flow of atomic hydrogen. Here, m_H is the mass of a hydrogen atom, T is the temperature of the flow, F_H is the planet-averaged upward H particle flux, b is the binary diffusion coefficient for the two species, g is the acceleration due to gravity, and X_H is the H molar mixing ratio at the base of the flow. In the specific case of hydrodynamic drag on oxygen atoms, the oxygen particle flux F_O is given by (Hunten et al., 1987; Chassefière, 1996b; Lammer et al., 2011b; Erkaev et al., 2014)

$$F_O = \frac{X_O}{X_H} F_H \left(\frac{m_c - m_O}{m_c - m_H} \right), \tag{15}$$

where X_O is the oxygen mixing ratio and m_O is the mass of an oxygen atom. The expression above is valid provided $m_c \geq m_O$; otherwise, $F_O = 0$. In the limit of large m_c, the ratio of the particle escape rates is simply the ratio of the abundances at the base of the flow.

It is important to note that both F_H and m_c are indirect functions of F_O; as oxygen begins to escape, the hydrogen particle flux decreases (at fixed energy input), decreasing the crossover mass and in turn reducing the rate of oxygen escape. In order to solve for the individual escape rates, it is convenient to define a reference particle flux F_H^{ref}, equal to the energy-limited particle escape

flux of H in the absence of oxygen (Chassefière, 1996b):

$$F_{\rm H}^{\rm ref} = \frac{\epsilon_{\rm XUV}\mathcal{F}_{\rm XUV}R_{\rm p}}{4GM_{\rm p}K_{\rm tide}m_{\rm H}}, \tag{16}$$

where we may write

$$m_{\rm H}F_{\rm H}^{\rm ref} = \frac{\dot{M}_{\rm EL}}{4\pi R_{\rm p}^2} = m_{\rm O}F_{\rm O} + m_{\rm H}F_{\rm H}. \tag{17}$$

By combining (15) and (17), we obtain as in Chassefière (1996b) the true hydrogen particle flux in terms of the reference flux:

$$F_{\rm H} = \begin{cases} F_{\rm H}^{\rm ref} & \text{if } m_{\rm c} < m_{\rm O} \\ F_{\rm H}^{\rm ref}\left(1 + \dfrac{X_{\rm O}}{X_{\rm H}}\dfrac{m_{\rm O}}{m_{\rm H}}\dfrac{m_{\rm c}-m_{\rm O}}{m_{\rm c}-m_{\rm H}}\right)^{-1} & \text{if } m_{\rm c} \geq m_{\rm O} \end{cases} \tag{18}$$

Inserting this into (14) and doing a little algebra, we obtain an expression for the crossover mass in terms of the reference flux (16):

$$m_{\rm c} = \begin{cases} m_{\rm H} + \dfrac{3kTF_{\rm H}^{\rm ref}}{2bg} & \text{if } F_{\rm H}^{\rm ref} < 10bgm_{\rm H}/kT \\ \dfrac{43}{3}m_{\rm H} + \dfrac{kTF_{\rm H}^{\rm ref}}{6bg} & \text{if } F_{\rm H}^{\rm ref} \geq 10bgm_{\rm H}/kT \end{cases} \tag{19}$$

In this derivation, I used $m_{\rm O} = 16m_{\rm H}$ as well as $X_{\rm H} = 2/3$ and $X_{\rm O} = 1/3$, assuming that all H and O are photolytically produced and that the dissociation of H_2 and O_2 is fast enough that both species are atomic close to the base of the flow (Chassefière, 1996b).

The condition for oxygen escape is $m_{\rm c} > m_{\rm O}$; using (17), we can write this as $\mathcal{F}_{\rm XUV} \geq \mathcal{F}_{\rm crit}$, where

$$\mathcal{F}_{\rm crit} \equiv 180\left(\frac{M_{\rm p}}{\rm M_\oplus}\right)^2\left(\frac{R_{\rm p}}{\rm R_\oplus}\right)^{-3}\left(\frac{\epsilon_{\rm XUV}}{0.30}\right)^{-1}\ \text{erg cm}^{-2}\ \text{s}^{-1} \tag{20}$$

where I have used $b = 4.8 \times 10^{17}(T/{\rm K})^{0.75}\ {\rm cm^{-1}s^{-1}}$ (Zahnle and Kasting, 1986), an average thermospheric temperature $T = 400$ K (Hunten et al., 1987; Chassefière, 1996b), and $K_{\rm tide} = 1$, corresponding to no tidal enhancement. For reference, a 1 $\rm M_\oplus$, 1 $\rm R_\oplus$ planet at the RV limit of a 0.1 $\rm M_\odot$ star (for which tidal effects are strongest), has $K_{\rm tide} \approx 0.88$; setting $K_{\rm tide} = 1$ leads to a

maximum underestimate of the mass loss rate of about 10%.

Thus, for an Earth-size terrestrial planet with $\mathcal{F}_{\rm XUV} > 180$ erg cm^{-2} s^{-1} (equivalent to $\sim 39\mathcal{F}_{\oplus}$), oxygen should begin to escape. The particle escape flux is determined by inserting (19) into (15):

$$F_{\rm O} = \frac{\eta}{2} F_{\rm H}, \tag{21}$$

where I define the oxygen escape parameter η as

$$\eta \equiv \begin{cases} 0 & \text{if } x < 1 \\ \dfrac{x-1}{x+8} & \text{if } x \geq 1 \end{cases} \tag{22}$$

with

$$x \equiv \frac{kTF_{\rm H}^{\rm ref}}{10bgm_{\rm H}}. \tag{23}$$

The parameter η is simply the ratio of the oxygen particle flux to its production rate (compare Equations 25 and 28 in the next section). It is thus limited to the range $0 \leq \eta \leq 1$. For $\eta \to 1$ (large x, high $\mathcal{F}_{\rm XUV}$), the oxygen flux approaches one-half the hydrogen flux, resulting in no accumulation of oxygen in the atmosphere. For $\eta \to 0$ ($x \to 1$, low $\mathcal{F}_{\rm XUV}$), oxygen escape tapers off and is zero for $x < 1$. Note that the condition $x \geq 1$ is mathematically equivalent to both $\mathcal{F}_{\rm XUV} \geq \mathcal{F}_{\rm crit}$ and $m_{\rm c} \geq m_{\rm O}$ in (19). In the next section, I derive simple expressions for the mass loss rates of hydrogen (29) and oxygen (30), the rate of oxygen buildup (31), and the rate of ocean loss (32) in terms of η and $\dot{M}_{\rm EL}$.

1.3.2.2 Rate of Ocean Loss and Oxygen Buildup in the Energy-Limited Regime

The energy-limited escape rate $\dot{M}_{\rm EL}$ is equal to the sum of the upward mass escape rates of hydrogen and oxygen:

$$\dot{M}_{\rm EL} = \dot{m}_{\rm H}^{\uparrow} + \dot{m}_{\rm O}^{\uparrow}. \tag{24}$$

Assuming all of the hydrogen and oxygen comes from photolysis of water, the mass production rate of oxygen is eight times that of hydrogen:

$$\dot{m}_{\rm O} = 8\dot{m}_{\rm H}. \tag{25}$$

Assuming further that all of the hydrogen escapes hydrodynamically, the total mass escape rate of hydrogen is $\dot{m}_{\rm H} = \dot{m}_{\rm H}^{\uparrow}$. Since the oxygen either escapes or accumulates in the atmosphere, the total mass escape rate of oxygen is $\dot{m}_{\rm O} = \dot{m}_{\rm O}^{\uparrow} + \dot{m}_{\rm O}^{\rm atm}$, where $\dot{m}_{\rm O}^{\rm atm}$ is the rate of oxygen buildup in the atmosphere. We thus have

$$\dot{m}_{\rm O}^{\uparrow} + \dot{m}_{\rm O}^{\rm atm} = 8\dot{m}_{\rm H}^{\uparrow}. \tag{26}$$

Given that the particle and mass fluxes are related by

$$\frac{F_{\rm O}}{F_{\rm H}} = \frac{1}{16}\frac{\dot{m}_{\rm O}^{\uparrow}}{\dot{m}_{\rm H}^{\uparrow}}, \tag{27}$$

it follows from expression (21) that

$$\dot{m}_{\rm O}^{\uparrow} = 8\dot{m}_{\rm H}^{\uparrow}\eta, \tag{28}$$

where η is given by (22). Combining expressions (24), (26) and (28), we have

$$\dot{m}_{\rm H}^{\uparrow} = \left(\frac{1}{1+8\eta}\right)\dot{M}_{\rm EL} \tag{29}$$

$$\dot{m}_{\rm O}^{\uparrow} = \left(\frac{8\eta}{1+8\eta}\right)\dot{M}_{\rm EL} \tag{30}$$

$$\dot{m}_{\rm O}^{\rm atm} = \left(\frac{8-8\eta}{1+8\eta}\right)\dot{M}_{\rm EL} \tag{31}$$

with $\dot{M}_{\rm EL}$ given by (5). Finally, the rate at which the ocean is lost (either to H only escape or H+O escape) is

$$\begin{aligned}\dot{m}_{\rm ocean} &= \left(\dot{m}_{\rm H}^{\uparrow} + \dot{m}_{\rm O}^{\uparrow} + \dot{m}_{\rm O}^{\rm atm}\right) \\ &= \left(\frac{9}{1+8\eta}\right)\dot{M}_{\rm EL},\end{aligned} \tag{32}$$

which approaches $\dot{M}_{\rm EL}$ in the limit $\eta \to 1$ (H+O escape) and $9\dot{M}_{\rm EL}$ in the limit $\eta \to 0$ (H only escape).

1.3.2.3 Diffusion Through a Background Gas

The energy-limited mass loss rate (5) is an upper limit to the thermal escape rate from a planetary atmosphere, as it assumes that all of the available XUV energy goes into driving the escape (after accounting for an efficiency factor, $\epsilon_{\rm XUV}$). However, the escape of hydrogen depends on the availability of hydrogen atoms at the base of the flow. If not all of the oxygen escapes or is absorbed by the surface, it can eventually become a major constituent of the atmosphere. Once this happens, hydrogen will have to diffuse through a static background of oxygen before it reaches the base of the hydrodynamic wind. The rate at which hydrogen can do so is given by the *diffusion limit* (Hunten, 1973; Zahnle et al., 1990):

$$F_{\rm H}^{\rm diff} \equiv \frac{bg(m_{\rm O} - m_{\rm H})}{kT(1 + X_{\rm O}/X_{\rm H})}, \tag{33}$$

which can be significantly lower than energy-limited escape flux (18), especially at early times.[5] In particular, note that when $X_{\rm O}/X_{\rm H} = 1/2$, $F_{\rm H}^{\rm diff} = 10bgm_{\rm H}/kT$, which is precisely the value of the reference flux at which oxygen begins to escape in (19). The diffusion limit is, in fact, defined as the maximum upward flux of a gas for which the background gas is static (Hunten, 1973). If the hydrogen particle flux is greater than this limit, oxygen *must* escape. The converse is also true; if hydrogen diffusion through oxygen in the lower atmosphere limits the supply of H atoms at the base of the flow, the escaping flux will be insufficient to drag away any oxygen.

Thus, for planets that build up significant amounts of oxygen in their atmospheres, the H

[5] In (33), I (conservatively) assumed diffusion through atomic oxygen. Diffusion through molecular oxygen is faster, but if the oxygen is photolyzed below the base of the hydrodynamic wind, it is diffusion through atomic oxygen that will bottleneck the escape.

particle escape flux is given by the smaller of (18) and (33), the O particle escape flux is zero, and the rate at which the ocean is lost is 9 times the H escape rate.

1.4 Tidal Evolution

Most of us are familiar with the concept of tides, particularly in Earth's ocean: the differential force of gravity due to the Moon (and the Sun) induces a tidal bulge in the ocean, which we experience as a passing high tide as the Earth rotates. On Earth, tides are relatively small – on the order of tens of centimeters – and are really only noticeable in the ocean. However, exoplanets orbiting close to their stars can experience tidal forces orders of magnitude larger than we do on Earth, which can cause deformation of the solid body (in the case of a terrestrial planet) or the envelope (in the case of a gaseous planet). This deformation can dissipate large amounts of energy via frictional heating, removing energy from either the rotation or the orbit (or both). The action of tidal forces can therefore lead to heating of the planet and evolution of its orbit. Planets in the habitable zones of M dwarfs, at periods of a couple to a few tens of days, are deep enough in the gravitational potential well of their stars that tidal evolution can be quite significant, causing eccentricities to change, orbits to shrink, and planet spins to synchronize with their orbital periods. In some cases, tidal heating can be significant enough to trigger runaway greenhouses (Barnes et al., 2013).

Studying tidal evolution is therefore essential to understanding the habitability of planets orbiting M dwarfs. Below, I discuss two common formalisms adopted to model the effect of tides: the constant phase lag (CPL) and the constant time lag (CTL) models. Both make strong assumptions about the nature of the tidal response of planets and rely heavily on approximations, but at low eccentricity the models generally agree and match solar system observations.

1.4.1 Constant Phase Lag

Classical tidal theory predicts that torques arising from interactions between tidal deformations on a planet and its host star lead to the secular evolution of the orbit and the spin of both bodies. The "equilibrium tide" model, developed by Sir George Darwin (Charles Darwin's son), approximates the tidal potential as a superposition of Legendre polynomials representing waves propagating along the surfaces of the bodies; these add up to give the familiar tidal "bulge" (Darwin, 1880). Because of viscous forces in the bodies' interiors, the tidal bulges do not instantaneously align with

the line connecting the two bodies. Instead, the N^{th} wave on the i^{th} body will lag or lead by an angle $\varepsilon_{N,i}$, assumed to be constant in the *constant phase lag* (CPL) model. In general, different waves may have different lag angles, and it is unclear how the $\varepsilon_{N,i}$ vary as a function of frequency. A common approach (see Ferraz-Mello et al., 2008) is to assume that the magnitudes of the lag angles are equal (see Goldreich and Soter, 1966), while their signs may change depending on the orbital and rotational frequencies involved. This allows us to introduce the *tidal quality factor*

$$Q_i \equiv \frac{1}{\varepsilon_{0,i}}, \tag{34}$$

which in turn allows us to express the lags (in radians) as

$$\varepsilon_{N,i} = \pm \frac{1}{Q_i}. \tag{35}$$

The parameter Q_i is a measure of the dissipation within the i^{th} body; it is inversely proportional to the amount of orbital and rotational energy lost to heat per cycle, in analogy with a damped-driven harmonic oscillator. The merit of this approach is that the tidal response of a body can be captured in a single parameter. Planets with high values of Q_p have smaller phase lags, dissipate less energy and undergo slower orbital evolution; planets with low values of Q_p have larger phase lags, higher dissipation rates, and therefore faster evolution. Measurements in the solar system constrain the value of Q_p for terrestrial bodies in the range 10-500, while gas giants are consistent with $Q_p \sim 10^4 - 10^5$ (Goldreich and Soter, 1966). Values of $Q_\star$ for the Sun and other main sequence stars are poorly constrained but are likely to be $\gtrsim 10^5 - 10^6$ (Schlaufman et al., 2010; Penev et al., 2012). Intuitively, this makes sense, given that the dissipation due to internal friction in rocky bodies should be much higher than that in bodies dominated by gaseous atmospheres. One should bear in mind, however, that the exact dependence of Q_i on the properties of a body's interior is likely to be extremely complicated. Given the dearth of data on the composition and internal structure of exoplanets, it is at this point impossible to infer precise values of Q_p for these planets.

By calculating the forces and torques due to the tides raised on both the planet and the star, one can arrive at the secular expressions for the evolution of the planet's orbital parameters, which are given by a set of coupled nonlinear differential equations; these are reproduced in the next

section.

1.4.1.1 Tidal Evolution Expressions

For zero inclination and zero obliquity, the tidal evolution equations for a, e, and ω_i in a two-body system are (Heller et al., 2011; Barnes et al., 2013)

$$\frac{da}{dt} = \frac{a^2}{4GM_\star M_p} \sum_{i \neq j} Z_i' \left(4\epsilon_{0,i} + e^2 \left[-20\epsilon_{0,i} + \frac{147}{2}\epsilon_{1,i} + \frac{1}{2}\epsilon_{2,i} - 3\epsilon_{5,i} \right] \right) \tag{36}$$

$$\frac{de}{dt} = -\frac{ae}{8GM_\star M_p} \sum_{i \neq j} Z_i' \left(2\epsilon_{0,i} - \frac{49}{2}\epsilon_{1,i} + \frac{1}{2}\epsilon_{2,i} + 3\epsilon_{5,i} \right) \tag{37}$$

$$\frac{d\omega_i}{dt} = -\frac{Z_i'}{8M_i r_{g,i}^2 R_i^2 n} \left(4\epsilon_{0,i} + e^2 \left[-20\epsilon_{0,i} + 49\epsilon_{1,i} + \epsilon_{2,i} \right] \right), \tag{38}$$

where the sums are taken over the two bodies. Here, G is the gravitational constant, R_i is the radius of the i^{th} body (planet or star), and $r_{g,i}$ are the radii of gyration. The parameter Z_i' is defined

$$Z_i' \equiv 3G^2 k_{2,i} M_j^2 (M_i + M_j) \frac{R_i^5}{a^9} \frac{1}{nQ_i}, \tag{39}$$

where $k_{2,i}$ is the Love number of degree 2 of the i^{th} body, which is of order unity and quantifies the contribution of the tidal deformation to the total potential, n is the mean motion, and Q_i are the tidal quality factors. The parameters $\epsilon_{N,i}$ are the signs of the phase lags (assumed equal in magnitude) of the N^{th} wave on the i^{th} body, calculated from

$$\begin{aligned}
\epsilon_{0,i} &= \mathrm{sgn}(2\omega_i - 2n) \\
\epsilon_{1,i} &= \mathrm{sgn}(2\omega_i - 3n) \\
\epsilon_{2,i} &= \mathrm{sgn}(2\omega_i - n) \\
\epsilon_{5,i} &= \mathrm{sgn}(n) \\
\epsilon_{8,i} &= \mathrm{sgn}(\omega_i - 2n) \\
\epsilon_{9,i} &= \mathrm{sgn}(\omega_i).
\end{aligned} \tag{40}$$

Since short-period planets around M dwarfs are likely to be tidally locked, one need not calculate the planetary spin from (38). Instead, the planet's rotation rate may be calculated from (Goldreich, 1966):

$$\omega_{p,\mathrm{eq}}^{\mathrm{CPL}} = n(1 + 9.5e^2), \tag{41}$$

where n is the mean motion.

1.4.1.2 The Typical Case

Because of the fast rotation of M dwarfs at early times, planets in the HZ are likely to be far outside the corotation radius of their parent stars. It is useful at this point to consider as an example the specific case of a tidally-locked planet for which $n \ll \omega_\star$. In this case, the stellar phase lags (40) are all positive and $\epsilon_{N,\star} = +1$. For a tidally locked planet, $\epsilon_{2,p} = \epsilon_{5,p} = \epsilon_{9,p} = 1$, and $\epsilon_{1,p} = \epsilon_{8,p} = -1$. The parameter $\epsilon_{0,p}$, however, is less straightforward to calculate. If tidal locking is to be maintained, the average angular acceleration over one period must be zero. Ferraz-Mello et al. (2008) show that the only self-consistent way to ensure this is if $\epsilon_{0,2}$ has a different magnitude than the other lags, equal to

$$\varepsilon_{0,p} = 12e^2\varepsilon_{2,p}, \tag{42}$$

which in my notation corresponds to

$$\epsilon_{0,p} = +12e^2. \tag{43}$$

Note that if $e = 0$, $\epsilon_{0,p} = 0$, consistent with $\omega_p = n$ in (40).

In the limit $M_\star \gg M_p$ and keeping only terms up to order e^2, Equations (36) – (38) now reduce

to

$$\frac{1}{a}\frac{da}{dt} = \left[3\sqrt{\frac{G}{M_\star}}\frac{k_{2,\star}R_\star^5 M_p}{Q_\star}\left(1+\frac{51}{4}e^2\right) - 21\sqrt{GM_\star^3}\frac{k_{2,p}R_p^5}{Q_p M_p}e^2\right]a^{-13/2} \tag{44}$$

$$\frac{1}{e}\frac{de}{dt} = \left[\frac{57}{8}\sqrt{\frac{G}{M_\star}}\frac{k_{2,\star}R_\star^5 M_p}{Q_\star} - \frac{21}{2}\sqrt{GM_\star^3}\frac{k_{2,p}R_p^5}{Q_p M_p}\right]a^{-13/2} \tag{45}$$

$$\frac{d\omega_\star}{dt} = -\left[\frac{3}{32}\frac{G^2 M_p^2 k_{2,\star}R_\star^3}{r_{g,\star}^2 n^2 Q_\star}\left(1+\frac{15}{2}e^2\right)\right]a^{-9}. \tag{46}$$

There is no need to calculate $d\omega_p/dt$, since I assume the planet's spin is instantaneously set to the equilibrium value. The first term in (44) and (45) is the orbital effect of the tide raised by the planet on the star. In both equations this term is positive, implying that the stellar tide acts to increase both a and e. This makes sense under the assumption that the planet is outside of corotation; the stellar bulge therefore leads the planet in the orbit, torquing the planet such that it speeds up and migrates outwards. For an eccentric orbit, the strongest impulse occurs at pericenter; since the planet must return to that point in the orbit, the pericenter distance is preserved, but the faster orbital speed results in a more distant apocenter, and thus higher eccentricity. A similar analysis for the tide raised by the star on the planet (the second term in each of the above equations) yields the result that such a tide acts to *decrease* both a and e.

The evolution of the orbit will therefore depend on the relative magnitudes of the stellar and planetary tides. A simple order of magnitude calculation shows that the ratio of the rate of change of the semi-major axis due to the tide generated by the star on the planet to that due to the tide generated by the planet on the star is

$$\left|\frac{\dot{a}_{\star\to p}}{\dot{a}_{p\to\star}}\right| \approx 7\left(\frac{M_\star}{M_p}\right)^2\left(\frac{R_p}{R_\star}\right)^5\left(\frac{Q_\star}{Q_p}\right)\left(\frac{e^2}{1+\frac{51}{4}e^2}\right). \tag{47}$$

For a $10\,\mathrm{M}_\oplus$, $2\,\mathrm{R}_\oplus$ mini-Neptune with $Q_p = 10^4$ orbiting a $0.1\,\mathrm{M}_\odot$, $0.15\,\mathrm{R}_\odot$ star with $Q_\star = 10^6$, this becomes

$$\left|\frac{\dot{a}_{\star\to p}}{\dot{a}_{p\to\star}}\right| \approx 3\times 10^4\left(\frac{e^2}{1+\frac{51}{4}e^2}\right), \tag{48}$$

which is $\gg 1$ for all $e \gtrsim 0.01$. For an Earth-like planet with $Q_p = 10^2$, the ratio is $\gg 1$ for all

$e \gtrsim 0.001$. Therefore, for a planet with moderate, nonzero initial eccentricity, the planetary tide should dominate the evolution, such that the planet's orbit will shrink.

1.4.2 Constant Time Lag

Unlike the CPL model, which assumes the phase lag of the tidal bulge is constant, the *constant time lag* (CTL) model assumes that it is the time interval between the bulge and the passage of the perturbing body that is constant. Originally proposed by Alexander (1973) and updated by Leconte et al. (2010), this model allows for a continuum of tidal wave frequencies and therefore avoids unphysical discontinuities present in the CPL model. However, implicit in the CTL theory is the assumption that the lag angles are directly proportional to the driving frequency (Greenberg, 2009), which is likely also an oversimplification. Note, however, that in the low eccentricity limit, both the CPL and the CTL models arrive at qualitatively similar results. At higher eccentricities, the CTL model is probably better suited, given that it is derived to eighth order in e (versus second order in the CPL model).

The tidal quality factors Q_i do not enter the CTL calculations at any point; instead, the dissipation is characterized by the time lags τ_i. Although there is no general conversion between Q_i and τ_i, Leconte et al. (2010) show that provided annual tides dominate the evolution,

$$\tau_i \approx \frac{1}{nQ_i}, \tag{49}$$

where n is the mean motion (or the orbital frequency) of the secondary body (in this case, the planet).

For a planet with $Q_p = 10^4$ in the center of the HZ of a late M dwarf, $\tau_p \approx 10$ s; rocky planets with lower Q_p may have values on the order of hundreds of seconds. Since $\tau \propto n^{-1}$, close-in planets should have much lower time lags. For reference, Leconte et al. (2010) argue that hot Jupiters should have $2 \times 10^{-3}\,\mathrm{s} \lesssim \tau_p \lesssim 2 \times 10^{-2}\,\mathrm{s}$.

The tidal evolution expressions are reproduced in the next section. For a more detailed review of tidal theory, the reader is referred to Ferraz-Mello et al. (2008), Heller et al. (2011), and the Appendices in Barnes et al. (2013).

1.4.2.1 Tidal Evolution Expressions

The expressions for the evolutions of the orbital parameters are (Barnes et al., 2013)

$$\frac{da}{dt} = \frac{2a^2}{GM_\star M_p} \sum_{i \neq j} Z_i \left(\frac{f_2(e)}{\zeta^{12}(e)} \frac{\omega_i}{n} - \frac{f_1(e)}{\zeta^{15}(e)} \right) \tag{50}$$

$$\frac{de}{dt} = \frac{11ae}{2GM_\star M_p} \sum_{i \neq j} Z_i \left(\frac{f_4(e)}{\zeta^{10}(e)} \frac{\omega_i}{n} - \frac{18}{11} \frac{f_3(e)}{\zeta^{13}(e)} \right) \tag{51}$$

$$\frac{d\omega_i}{dt} = \frac{Z_i}{M_i r_{g,i}^2 R_i^2 n} \left(\frac{f_2(e)}{\zeta^{12}(e)} - \frac{f_5(e)}{\zeta^9(e)} \frac{\omega_i}{n} \right), \tag{52}$$

where

$$Z_i \equiv 3G^2 k_{2,i} M_j^2 (M_i + M_j) \frac{R_i^5}{a^9} \tau_i \tag{53}$$

and

$$\begin{aligned}
\zeta(e) &\equiv \sqrt{1 - e^2} \\
f_1(e) &\equiv 1 + \frac{31}{2} e^2 + \frac{255}{8} e^4 + \frac{185}{16} e^6 + \frac{25}{64} e^8 \\
f_2(e) &\equiv 1 + \frac{15}{2} e^2 + \frac{45}{8} e^4 + \frac{5}{16} e^6 \\
f_3(e) &\equiv 1 + \frac{15}{4} e^2 + \frac{15}{8} e^4 + \frac{5}{64} e^6 \\
f_4(e) &\equiv 1 + \frac{3}{2} e^2 + \frac{1}{8} e^4 \\
f_5(e) &\equiv 1 + 3e^2 + \frac{3}{8} e^4.
\end{aligned} \tag{54}$$

As before, assuming the planet is tidally locked, its rotation rate is given by

$$\omega_{p,\mathrm{eq}}^{\mathrm{CTL}} = n \left(\frac{f_2(e)}{\zeta^3(e) f_5(e)} \right). \tag{55}$$

1.4.2.2 The Typical Case

As in the CPL model, the tides raised on the planet usually dominate the evolution. Plugging the result of (55) into (50) and (51), we find that the second (negative) terms dominate and the orbit

should therefore shrink and circularize. Proceeding as before,

$$\left|\frac{\dot{a}_{\star\to p}}{\dot{a}_{p\to\star}}\right| \approx \left(\frac{M_\star}{M_p}\right)^2 \left(\frac{R_p}{R_\star}\right)^5 \left(\frac{\tau_p}{\tau_\star}\right) \left|F\left(e, \frac{\omega_\star}{n}\right)\right|, \tag{56}$$

where

$$\begin{aligned} F\left(e, \frac{\omega_\star}{n}\right) &\equiv \frac{\frac{f_2^2(e)}{f_5(e) f_1(e)} - 1}{\frac{f_2(e)\zeta^3(e)}{f_1(e)} \frac{\omega_\star}{n} - 1} \\ &\approx -\frac{241}{16\left(\frac{\omega_\star}{n} - 1\right)} e^2 \end{aligned} \tag{57}$$

for small e. For the same mini-Neptune considered in the CPL case (48), this becomes

$$\left|\frac{\dot{a}_{\star\to p}}{\dot{a}_{p\to\star}}\right| \approx 5 \times 10^5 \left(\frac{e^2}{\frac{\omega_\star}{n} - 1}\right), \tag{58}$$

which is $\gg 1$ for all $e \gtrsim 0.001\sqrt{\omega_\star/n - 1}$, implying that the CTL model also predicts a net inward migration due to tides.

2 "Mirage" Earths

Portions of this chapter were originally published in collaboration with R. Barnes in the journal Astrobiology (Luger and Barnes, 2015, Astrobiology, © 2015 Mary Ann Liebert, Inc.), and are reproduced below with the permission of Mary Ann Liebert, Inc.

2.1 Introduction

Planets around M dwarfs are subject to an array of processes that could negatively impact their habitability. M dwarfs are extremely active (Reid and Hawley, 2005; Scalo et al., 2007), emitting large fractions of their luminosity in the X-ray and extreme ultraviolet (XUV). XUV photons are not only biologically harmful, but can drive fast atmospheric escape that leads to the erosion of planetary atmospheres, as I discussed in the previous chapter (Watson et al., 1981; Lammer et al., 2003; Yelle, 2004; Erkaev et al., 2007; Tian, 2009; Lammer et al., 2009; Owen and Jackson, 2012; Lammer et al., 2013; Erkaev et al., 2013; Koskinen et al., 2013a,b). Moreover, the HZs of these stars are significantly closer in, exposing planets to potentially catastrophic flaring events (Segura et al., 2010) and strong, detrimental tidal effects (Barnes et al., 2013).

However, an issue that has not received as much attention in the literature is the fact that M dwarfs can take up to 1 Gyr to settle onto the main sequence (see §1.1). During the contraction phase following their formation, these stars can be one or even two orders of magnitude more luminous than when they reach the MS. Since terrestrial planets probably form between 10 and 100 Myr after the formation of the star (Chambers, 2004; Raymond et al., 2007; Kleine et al., 2009; Raymond et al., 2013), planets in the HZs of these stars are subject to extreme levels of insolation early on, a fact first pointed out by Lissauer (2007). As a consequence, these planets are likely to be in a runaway greenhouse provided they have sufficient surface water (see, e.g., Kasting, 1988; Kopparapu et al., 2013).

Many papers have explored the effects of a runaway greenhouse on Venus, arguing that it may have lost one or more Earth oceans of water as a consequence of an early runaway (Kasting et al., 1984; Kasting, 1988; Chassefière, 1996a,b; Kulikov et al., 2006; Gillmann et al., 2009). During a runaway greenhouse, water vapor reaches the stratosphere, where it is easily photolyzed by UV radiation. Heating of the upper atmosphere by XUV radiation can then drive a hydrodynamic wind

that carries the hydrogen (and potentially some of the oxygen) to space, leading to the irreversible loss of a planet's surface water, oxidation of the surface, and possible accumulation of oxygen in the atmosphere. Recently, Hamano et al. (2013) extended this idea to exoplanetary systems, arguing for the existence of two fundamentally different types of terrestrial planets: type I planets, which undergo short-lived runaway greenhouses during their formation, and type II planets, which form interior to a critical distance and can remain in runaway greenhouses for as long as 100 Myr. The former type of planet, like Earth, retains most of its water inventory and may evolve to become habitable. The latter, similarly to Venus, undergoes complete surface desiccation during the runaway.

In this chapter I show that because of the early evolution of the star, many terrestrial planets within the HZ of M dwarfs could be similar to type II planets and may therefore be uninhabitable. My work builds on that of Barnes and Heller (2013), Heller and Barnes (2013), and (Barnes et al., 2013), who studied water loss during early runaway greenhouses on planets orbiting white dwarfs and brown dwarfs and on exomoons orbiting giant planets. I further build on the results of Wordsworth and Pierrehumbert (2013b), who showed that significant water loss can occur for planets near the inner edge of the HZ of M dwarfs; however, those authors considered a constant stellar luminosity and thus did not account for the early runaway greenhouse state, which I show can result in water loss rates that are orders of magnitude higher. In a follow-up paper, Wordsworth and Pierrehumbert (2014) showed that water loss can lead to the buildup of abiotic O_2 in the atmospheres of planets in the HZ. I extend this mechanism and demonstrate that hundreds to thousands of bars of abiotic oxygen are possible for planets throughout the HZs of M dwarfs. While a large fraction of this oxygen may be removed by surface processes, some exoplanets could retain detectable amounts of O_2 in their atmospheres for extended periods of time. This validates the predictions of Schindler and Kasting (2000), concerning oxygen atmospheres on Venus-like planets, and of Lammer et al. (2011b), Lammer et al. (2011a), Lammer (2013), and Fossati et al. (2014), who argued that oxygen-rich atmospheres could develop on G dwarf planets in the HZ, in particular on super-Earths. The present work could also strengthen the results of Wordsworth and Pierrehumbert (2014), Tian et al. (2014), and Meadows et al. (2017) that oxygen is not a reliable biosignature; in fact, planets with such elevated quantities of O_2 may be uninhabitable. Because of their potential lack of water and presence of atmospheric oxygen (which is potentially

a biosignature), I refer to these planets as "mirage" Earths.

The chapter is organized as follows: in §2.2 I discuss some preliminary considerations about the evolution of the water and atmospheric oxygen content of terrestrial planets. In §2.3 I describe my model, followed by my results in §2.4 and §2.5. I discuss the stability of O_2-rich atmospheres and implications for habitability in §2.6.

2.2 Preliminary Considerations

2.2.1 Planet Formation & Initial Water Content

The terrestrial planets in the solar system are thought to have formed *in situ* between 10 and 100 Myr after the formation of the Sun (Chambers, 2004; Kleine et al., 2009; Raymond et al., 2013). However, whether or not the bulk of Earth's oceans formed during this accretion period is still up to debate. A recent isotopic study by Hartogh et al. (2011) suggests that a large fraction of Earth's water may have been delivered by comets, possibly a result of scattering by the giant planets prior to and during the Late Heavy Bombardment (Gomes et al., 2005). Nevertheless, simulations show that during the final stages of Earth's assembly, the planet's feeding zone encompassed enough water-rich material to supply $15-70$ terrestrial oceans (TO) of water in the first 70 Myr (Morbidelli et al., 2000; Raymond et al., 2006; Chassefière et al., 2012), and thus a wet *in situ* formation for Earth-like planets is entirely plausible. In this chapter, I define 1 TO $\equiv 1.39 \times 10^{24}$ g ($\sim$ 270 bars) of H_2O, the total amount of water in Earth's oceans (Kasting, 1988; Kulikov et al., 2006).

Whether or not this applies to M dwarfs is unclear. Studies by Raymond et al. (2007) and Lissauer (2007) suggest that planets forming *in situ* in the HZs of M dwarfs are likely to form quickly ($\sim$ 10 Myr after the formation of the star), to be relatively small ($\lesssim 0.3 M_\oplus$) and to have water contents smaller than Earth's. If that is the case, a potential mechanism for forming a wet, Earth-size planet in the HZ is early formation beyond the snow line (the region of the circumstellar disk beyond which water and other volatiles are able to condense into ices and serve as building blocks for planets) followed by disk-driven migration into the HZ (Carter-Bond et al., 2012).

Planets that form prior to the dissipation of the gaseous circumstellar disk experience strong torques that can induce rapid inward migration, especially for planets in the terrestrial mass range (Ward, 1997). In particular, planets that form beyond the snow line, where accretion is orders of magnitude faster due to the higher density of solids, can potentially migrate into the HZ (Ida and

Lin, 2008a,b; Ogihara and Ida, 2009; Cossou et al., 2013). Since disk lifetimes are typically quite short, ranging from ~ 1 to ~ 10 Myr (Walter et al., 1988; Strom et al., 1989), planets migrating in this fashion will settle into their new orbits relatively early. As Raymond and Cossou (2014) show, the abundance of short-period planets with masses $\lesssim 10$ $M_\oplus$ is strong evidence for the ubiquity of this mechanism, since it is highly unlikely that these systems formed *in situ*. Because of their formation beyond the snow line, these planets will likely have large ice mass fractions and therefore much larger initial water contents than Earth (see, e.g., Kuchner, 2003).

It is important to note that not all of a planet's water may be at its surface (or in its atmosphere), particularly at early times. During terrestrial planet formation, giant impacts can deliver enough energy to partially or completely melt a planet's mantle; as a consequence, many of the terrestrial bodies in the solar system may have experienced a magma ocean phase (Matsui and Abe, 1986; Zahnle et al., 1988; Elkins-Tanton and Seager, 2008; Elkins-Tanton, 2008, 2011, 2012; Lammer, 2013; Lebrun et al., 2013; Hamano et al., 2013). Since water is highly soluble in magma, a large fraction of the planet's water content will initially be trapped in the mantle. As the planet cools and the mantle begins to solidify from the bottom up, large amounts of water (between $\sim$ 60 and 99% of the total amount in the mantle) are exsolved to form a massive steam atmosphere, which may eventually collapse to form an ocean (Elkins-Tanton, 2011). Typically, this process occurs within a few Myr of the end of the accretion phase (Elkins-Tanton, 2008, 2011), but the exact timescale for solidification depends on the stellar flux. Lebrun et al. (2013) find that while Earth's magma ocean lasted for ~ 1.5 Myr, it may have lasted as long as 10 Myr on Venus due to the blanketing effect of a runaway greenhouse. Furthermore, Hamano et al. (2013) argue that above a certain stellar flux (close to that received by Venus), a magma ocean may take as long as 100 Myr to solidify. While these planets never develop massive steam atmospheres—since the bulk of the water is always in the mantle—a few tens of bars of water vapor are always maintained in the atmosphere due to a feedback cycle (Matsui and Abe, 1986; Zahnle et al., 1988). Large quantities of water may thus still be lost via hydrodynamic escape from these planets, since escape of water to space will be balanced by exsolution from the magma ocean. As Hamano et al. (2013) point out, this could lead to the complete desiccation of a planet's mantle, potentially terminating tectonics and resulting in permanently dry surface conditions.

One might thus expect that because of the high-luminosity PMS phase of M dwarfs, planets in

the HZs of these stars could remain in the magma ocean phase for several to several tens of Myr, though this should be investigated further. While the magma ocean phase does not prevent water loss to space, it could suppress the buildup of atmospheric O_2. I revisit this point in §2.2.2.3.

2.2.2 Oxygen Atmospheres

2.2.2.1 Earth-like Planets

Because of its highly reactive nature, oxygen is not typically stable in terrestrial planet atmospheres and tends to quickly react at the surface by stripping electrons from *reducing* substances in a process known as *oxidation*. Oxygen on Earth is continuously produced by photosynthesis; in the absence of a steady source, continental weathering, volcanic outgassing of reducing gases, and oxidation of basalt via hydrothermal processes at oceanic ridges would quickly remove most of the atmospheric O_2 (Lécuyer and Ricard, 1999). Recently, Catling (2014) compiled a table of all major oxidation processes on Earth, obtaining an estimated total O_2 removal rate of 2.21×10^{13} mol/year, or ~ 150 bar/Gyr. Over two-thirds of this removal is due to weathering of surface rocks in a process that continuously oxidizes the Earth's crust.

However, the absorption of hundreds to thousands of bars of O_2 could eventually lead to the irreversible oxidation of the surface; processes such as plate tectonics, volcanic resurfacing, or volcanic plumes are therefore necessary to subduct the oxidized species and/or supply fresh reductants to the surface. Catling et al. (2001) show that the modern rate of (oxidized) Fe^{3+} subduction is equivalent to a removal rate of $5.0 \times 10^{11} - 1.9 \times 10^{12}$ mol O_2/year, or $3 - 12$ bars/Gyr; volcanic outgassing contributes an extra ~ 15 bars/Gyr (Catling, 2014).

Thus, while a reducing surface can efficiently remove tens and possibly hundreds of bars of O_2 on Gyr timescales, the removal will eventually be bottlenecked by the rate of surface recycling. Given that neither Mars nor Venus possess an active tectonic cycle, it is possible that many exoplanets lack plate tectonics, which could significantly delay the removal of O_2 from their atmospheres.

2.2.2.2 Water Worlds

The large weathering rates discussed above assume continental coverage similar to Earth's. Provided they are not completely desiccated, planets that form with tens to hundreds of TO could have significantly more surface water and less continental coverage than Earth after the runaway

phase, potentially allowing for longer-lived O_2 atmospheres.

But could the oceans act as an oxygen sink? Oxygen is soluble in seawater, saturating at about 8 ml/L (3.8×10^{-4} mol/L) at 0° C and 35‰ salinity (Levitus, 1982). According to Henry's law, the solubility is proportional to the partial pressure of the gas in equilibrium with the liquid; given a partial pressure of 0.21 bar of O_2 on Earth, this corresponds to roughly 0.015 bar of dissolved O_2 per bar of atmospheric O_2 (for a planet with 1 TO of water). Scaling this to different ocean masses $m_{\rm ocean}$, we have

$$\frac{\text{mass of dissolved } O_2}{\text{mass of } O_2 \text{ in atmosphere}} = 0.015 \times \left(\frac{m_{\rm ocean}}{1 \text{ TO}}\right), \tag{59}$$

implying that a terrestrial planet would need roughly 70 TO to absorb half of its atmospheric oxygen into the oceans.

Oxidation of rocks at the seafloor could further deplete the atmospheric O_2, but this would require efficient mixing of the oxygen to great depths, which may be difficult for planets with deep oceans. Moreover, tectonic activity or volcanic resurfacing may still be necessary to subduct the oxidized rocks and sustain a long-term surface reductant flux. As a consequence, water worlds may take longer to remove a given amount of O_2 from their atmospheres.

2.2.2.3 Planets With Molten Surfaces

I pointed out in §2.2.1 that because of the high luminosities of M dwarfs early on, planets in the HZ could have magma oceans for extended periods of time following their formation. While this process should not directly affect water loss to space—since several tens of bars of water vapor remain in the atmosphere during the magma ocean phase (Matsui and Abe, 1986; Zahnle et al., 1988; Hamano et al., 2013)—it may prevent the accumulation of atmospheric oxygen. Based on measurements of oxygen diffusion in magma by Wendlandt (1991), Gillmann et al. (2009) argue that photolytically produced oxygen on Venus could diffuse to a depth of ~ 1 km over 100 Myr—certainly not enough to absorb all of Venus' oxygen, as this would require an oxidation depth of hundreds of km. However, a convecting magma ocean with a vertical mixing scale of order the mantle thickness (~ 3000 km) could effectively remove all of the atmospheric oxygen (on the order of several hundred bars) on Venus during its early runaway period.

However, the mantle solidification process on planets around M dwarfs is probably different

from that on Venus, given the steady decrease in the stellar luminosity with time. While early on these planets receive stellar fluxes several times that received by Earth, during the later stages of the runaway the lower stellar flux could lead to the solidification of most of their mantles, potentially allowing for the buildup of O_2 in the atmosphere.

Nevertheless, it is reasonable to expect that planets with Earth-like compositions may be able to absorb a large fraction of the photolytically produced O_2 into a primitive magma ocean, at least in the early stages of the runaway, provided (i) the surface composition is similar to Earth's, melting at or below ~ 1500 K; (ii) the surface/interior is initially reducing and capable of absorbing large amounts of oxygen; (iii) the magma ocean is deep and convective, with a sufficiently short turnover time. Strong tidal heating could also potentially extend the magma ocean phase and drive rapid resurfacing, which could lead to efficient oxygen removal from the atmosphere.

2.2.2.4 The Case of Venus

Finally, I consider the specific case of Venus, which may have lost one or more TO of water during its early runaway period (Kasting et al., 1984; Kasting, 1988; Chassefière, 1996a,b; Kulikov et al., 2006; Gillmann et al., 2009). This process should have led to the production of several hundred bars of O_2, a large fraction of which may have been deposited in the atmosphere (e.g., Gillmann et al., 2009). Given the negligible oxygen content of the Venusian atmosphere today (Chassefière, 1997), one or more effective oxygen sinks must have existed in the past.

While studies disagree on the process responsible for the removal of this oxygen, many plausible mechanisms exist. Chassefière (1997) showed that a strong primitive solar wind could have heated the upper layers of the Venusian atmosphere, enhancing the thermal escape of hydrogen and facilitating the hydrodynamic drag on the oxygen, potentially carrying all of it to space. Kulikov et al. (2006), on the other hand, argue that nonthermal interactions between O^+ ions and the solar wind could have removed 1 TO of oxygen (~ 240 bars) over 4.6 Gyr. Surface processes may have also contributed, but Rosenqvist and Chassefière (1995) show that tectonic activity ~ 15 times more vigorous than on Earth would be required to subduct all the atmospheric oxygen. Gillmann et al. (2009) argue that vigorous outgassing of reduced gases could remove the equivalent of 1 TO of O_2 in 4 Gyr. Finally, a magma ocean could have removed most or all of the O_2 (Hamano et al., 2013), though rigorous quantitative studies of this process are lacking.

The removal of a few hundred bars of O_2 during/after the runaway greenhouse period is entirely plausible, but highly dependent on properties of the star-planet system. In a review of the evolution of Venus, Chassefière et al. (2012) argue that for different initial conditions, Venus could have had an O_2-rich atmosphere for $\sim$ 1 Gyr, and that such planets may exist around other stars.

2.3 Model Description

For varying planet masses, initial water content, formation times, and HZ limits, I evolve planet/star systems forward in time for 5 Gyr on a grid of stellar mass and semi-major axis, keeping track of the duration of the runaway greenhouse phase, the amount of water lost and the total O_2 buildup. As I outline below, I run separate cases assuming either energy-limited escape or diffusion-limited escape. For simplicity, I do not consider the case of radiation/recombination-limited escape; the water loss values reported here should therefore be considered upper limits. Integrations are performed using an adaptive time-stepping scheme similar to that in Barnes et al. (2013).

I consider planets with masses 1 and 5 $M_\oplus$, with radii calculated from Fortney et al. (2007) for an Earth-like planet composition (2/3 silicate rock, 1/3 iron). I assume these planets have no significant H/He envelopes and that background atmospheric gases are negligible. The latter is justified by the fact that during a runaway greenhouse on a planet with 1 TO, the atmosphere contains about 270 bars of water vapor (Kasting, 1988); a background atmosphere similar to that on the present Earth is negligible. I further assume zero eccentricity and no tidal evolution for the planets in my runs. I discuss how these assumptions could change my results in §2.6.

I assume a formation time of 10 Myr for all planets. As I discuss in §2.2.1, this is probably an upper limit to the migration timescale for planets that form beyond the snow line. As for planets that form *in situ* in the HZ, 10 Myr is roughly equal to the formation time around a $\sim 0.3 M_\odot$ star in the simulations of Raymond et al. (2007), but is significantly longer than that predicted by the scaling arguments in Lissauer (2007). My adopted value is therefore likely to overestimate the formation time for planets around low mass M dwarfs. Around the highest mass M dwarfs and K dwarfs ($\gtrsim 0.5 M_\odot$), I am likely underestimating the formation time. In both formation scenarios, I vary the initial surface water content between 1 and 100 TO, noting that planets that migrate from beyond the snow line are more likely to have $\gtrsim$ 10 TO.

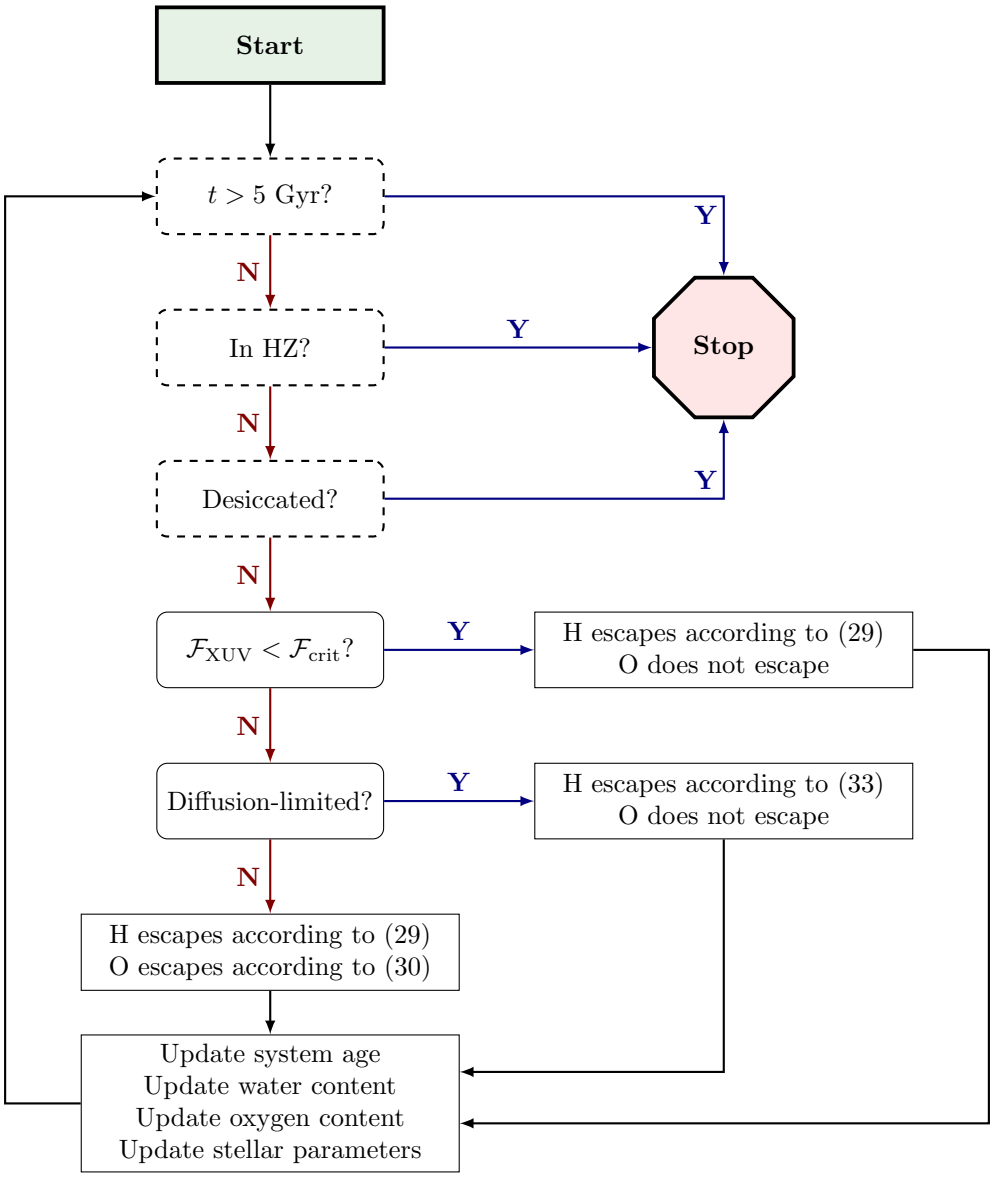

Figure 3: Flowchart representing a single integration of my code. The three halting conditions are represented by dashed boxes. See §2.3 for a detailed description of my model.

At each timestep, I calculate the HZ limits for $1M_\oplus$ and $5M_\oplus$ planets from Kopparapu et al. (2014), using $L_{\rm bol}$ and $T_{\rm eff}$ from the cooling tracks of Baraffe et al. (1998) for solar metallicity. I calculate water loss and O_2 buildup as long as planets are in a runaway greenhouse. Given the variety of mechanisms capable of removing atmospheric O_2 and their complex dependence on an array of planetary properties (§2.2.2), I do not model the details of oxygen absorption by the surface. Instead, I consider two limiting cases regarding the oxygen: **(a)** efficient absorption by the surface, corresponding to an effectively instantaneous removal of O_2 by surface processes; and **(b)** inefficient absorption, corresponding to a rate of oxygen buildup much larger than the rate at which it is removed.

In case (a), which could be the case of a planet with a deep magma ocean, I assume the atmosphere is predominantly H_2O and that diffusion through an oxygen-rich layer does not take place; I therefore use the energy-limited escape equations (29)-(32) to calculate H and O loss rates and O_2 buildup rates. In case (b), corresponding to (say) a planet with a highly oxidized surface and/or one that lacks plate tectonics, I assume that all of the O_2 remains in the atmosphere. This should quickly result in a depletion of H and H_2O relative to oxygen in the upper layers of the atmosphere, such that hydrogen escape will be limited by diffusion. I therefore set the hydrogen escape flux to the *minimum* of the energy-limited escape flux (Equation 18) and the diffusion-limit escape flux (Equation 33). In Equation (18), the ratio $X_{\rm O}/X_{\rm H}$ is calculated from the ratio of H_2O to O_2 in the atmosphere, assuming the two species are well-mixed below the diffusion layer. In this regime, the oxygen escape flux is assumed to be zero.

I report the amount of oxygen retained by the planet (either in the atmosphere or absorbed by the surface) as an equivalent pressure in bars, which I define to be the surface pressure of the oxygen *if it were the only gas in the atmosphere*. Thus, an equivalent pressure of 1 bar corresponds to an amount of oxygen equal in mass to the atmosphere of the current Earth, or about five times the current mass of oxygen in the atmosphere.[6] For a given planet, the *partial* pressure corresponding to this amount will depend on the mixing ratios and the mean molecular weights of other species in the atmosphere. Therefore, to preserve generality, below I present the equivalent O_2 pressure in all of my figures.

In order to capture the uncertainty in the critical stellar flux above which planets go runaway, I

[6]Since the molecular weights of N_2 and O_2 are similar, the equivalent pressure of O_2 on the present-day Earth is very close to its partial pressure.

also run two separate sets of models: one in which the runaway occurs interior to the RG limit (the default case), and one interior to the RV limit. Since terrestrial planets with surface oceans are likely to enter the runaway phase somewhere in between these limits, the two runs should roughly bracket the actual evolution.

For simplicity, I use a saturation time of 0.1 Gyr for K dwarfs ($M_\star > 0.6\mathrm{M}_\odot$) and 1 Gyr for M dwarfs. I adopt an XUV saturation fraction of $\frac{L_{\mathrm{XUV}}}{L_{\mathrm{bol}}} = 10^{-3}$ and a power law slope $\beta = -1.23$ in Equation (1). I assume an XUV absorption efficiency in the range $0.15 \leq \epsilon_{\mathrm{XUV}} \leq 0.30$, typical of hydrogen-rich atmospheres (Chassefière, 1996a; Wordsworth and Pierrehumbert, 2013b). In order to obtain an upper limit to the amount of water lost from these planets, I take $\epsilon_{\mathrm{XUV}} = 0.30$, unless otherwise noted. As in Wordsworth and Pierrehumbert (2014), I further assume that water loss is limited by the escape of H to space rather than by the H_2O photolysis rate.

I do not model the escape of hydrogen after the planet enters the HZ. Even after a planet leaves the runaway greenhouse, diffusion of water vapor into the stratosphere and H escape can still be significant; for instance, Wordsworth and Pierrehumbert (2014) show that about 28% of a TO can be lost from an N_2-poor Earth in the HZ over 4 Gyr.

A flowchart illustrating a sample integration of my code is presented in Figure 3. The integration continues as long as all three halting conditions (represented by the dashed boxes) evaluate to false. The code halts if **(1)** the age of the system reaches 5 Gyr; **(2)** the planet enters the HZ (corresponding to the end of the runaway greenhouse phase); or **(3)** the planet is completely desiccated. Otherwise, the code calculates H and O loss rates and O_2 buildup rates as outlined above. The code then updates stellar and planetary parameters and loops.

2.4 Results: Evolution of the Habitable Zone

2.4.1 Location of the Habitable Zone

In Figure 5 I show the evolution of the HZ as a function of time for M, K, and G dwarfs. In the left panel, I plot the position of the RV limit as a function of stellar age for stars ranging in mass from 0.08 $\mathrm{M}_\odot$ to 1 $\mathrm{M}_\odot$. For solar-type G stars, the HZ evolves inwards primarily in the first 10 Myr, prior to the formation of terrestrial planets. For K dwarfs ($0.6\mathrm{M}_\odot \lesssim M_\star \lesssim 0.9\mathrm{M}_\odot$), only planets that form in the first $\sim$ 10 Myr experience a significant change in the location of the HZ. The HZ of M dwarfs, on the other hand, moves in significantly, even after the formation of terrestrial

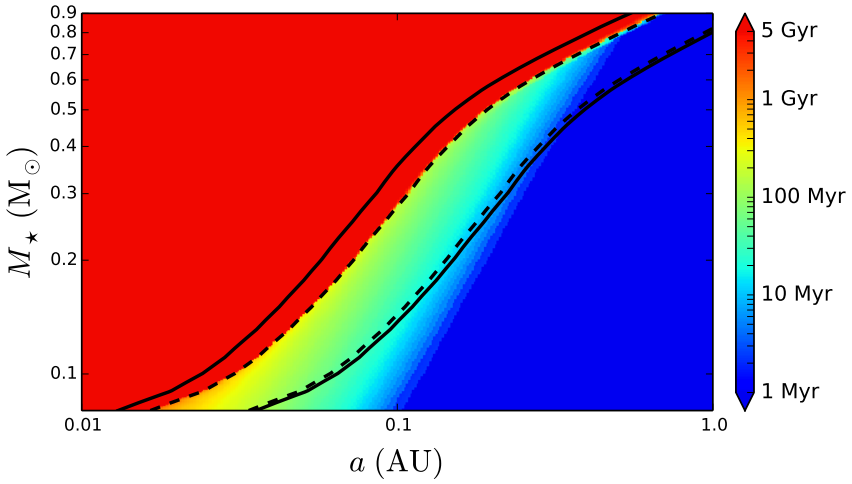

Figure 4: Duration of the runaway greenhouse for planets that formed at 10 Myr with abundant surface water. The solid black lines correspond to the RV (left) and EM (right) limits (the empirical HZ); the dashed lines correspond to the RG (left) and the MG (right) limits (the theoretical HZ). All limits are those at a stellar age of 5 Gyr. Red corresponds to planets that never leave the runaway state, since these are always interior to the HZ. Blue corresponds to planets that spend less than 1 Myr in a runaway greenhouse. Planets throughout the HZs of all M dwarfs spend a substantial amount of time in a runaway greenhouse.

planets. Around the lowest mass M dwarfs, the inner edge of the HZ moves in by nearly an order of magnitude after 10 Myr.

In the right panel, I plot the RV limit at different stellar ages on a classical HZ plot; the extent of the HZ at 1 Gyr is indicated by the blue shading. Note that even planets located at the outer edge of the HZ at $t > 1$ Gyr were interior to the HZ early on, especially for low mass M dwarfs. The change in the location of the HZ of solar-type stars is comparatively small. For stars more massive than about 0.8 $M_\odot$, the HZ moves inward during the PMS phase and then outward due to the steady increase of $L_{\rm bol}$ of these stars once they are on the MS. The HZ of M dwarfs, on the other hand, moves monotonically inward for up to 1 Gyr. After this point, the HZ boundaries remain relatively stationary for up to tens of Gyr due to these stars' slow MS evolution.

2.4.2 Duration of the Runaway Phase

Figure 4 shows the duration of the runaway phase for planets that form at 10 Myr, assuming the runaway occurs interior to the RG limit. The vertical axis (stellar mass) ranges from late M dwarfs (0.08 $M_\odot$) to early K dwarfs (0.9 $M_\odot$). The horizontal axis (semi-major axis) spans the HZ, which

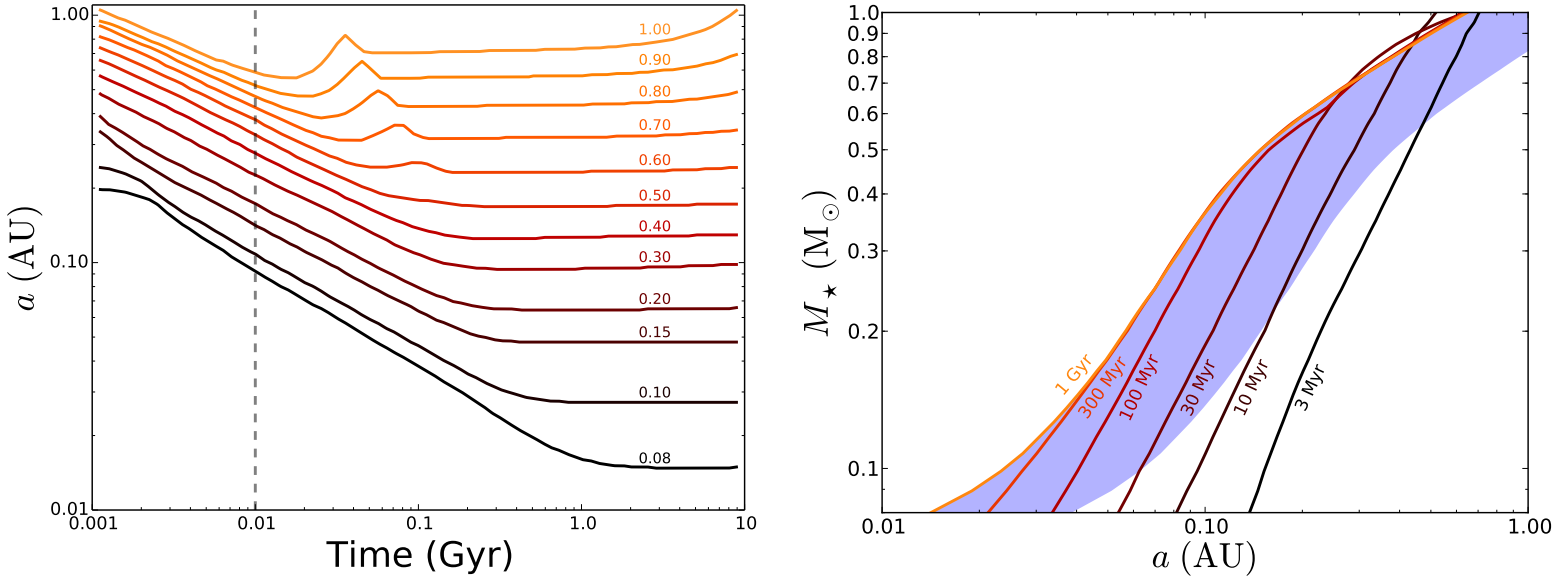

Figure 5: Evolution of the position of the inner edge of the empirical HZ (RV limit) as a function of time. **Left:** Location of the RV limit versus stellar age for stars between 0.08 $M_\odot$ and 1 $M_\odot$. The vertical dashed line corresponds to the 10 Myr formation time assumed in my model. **Right:** Contours of the RV limit on a classical HZ plot between 3 Myr (black line) and 1 Gyr (orange line). The empirical HZ at 1 Gyr is shaded in blue for reference. Note that the inner edge moves in by about an order of magnitude for the lowest mass M dwarfs.

is bounded on the left and right by the RV and EM limits (solid lines); from left to right, the RG and MG limits are indicated by the dashed lines. The limits are those at a stellar age of 5 Gyr.

Unsurprisingly, interior to the RG limit, planets spend the entire age of the system (5 Gyr) in a runaway state. Note that once a planet's water is depleted, it is technically no longer in a runaway greenhouse, since the atmospheric infrared windows will open up and the surface will cool. For the purposes of this figure, I therefore assume an unlimited surface water inventory.

Far to the right of the HZ, planets are never in an insolation-induced runaway greenhouse. However, throughout most of the HZ of M dwarfs, planets spend tens to hundreds of Myr in an early runaway greenhouse phase. Planets in the HZs of K dwarfs may experience a runaway greenhouse for a few tens of Myr, but only if they form early. Above about 0.8 $M_{\odot}$, the duration is negligible, except in the vicinity of the RG boundary; this is the case for planets around solar-type stars.

Because of this early runaway phase, many planets orbiting in the HZs of M dwarfs could potentially remain permanently uninhabitable. In the next section, I investigate the effect that the runaway greenhouse has on water loss and oxygen buildup on these planets.

2.5 Results: Water Loss and O_2 Buildup

2.5.1 Validation Against Venus

Before I present my results for M dwarf planets, I briefly examine water loss and O_2 buildup on Venus. Assuming an XUV saturation fraction $\frac{L_{\rm XUV}}{L_{\rm bol}} = 10^{-3}$, a saturation timescale of 0.1 Gyr, a formation time of 50 Myr, an initial water content of 1 TO (Raymond et al., 2006), and a runaway greenhouse interior to the RG limit, I find that Venus is completely desiccated in both the energy-limited and diffusion-limited escape regimes, accumulating an equivalent O_2 pressure between $\sim$ 120 bars (energy-limited) and $\sim$ 240 bars (diffusion-limited). For a much larger initial water content of 10 TO, it is still completely desiccated in both regimes and builds up between $\sim$ 1800 and $\sim$ 2400 bars of O_2. If, on the other hand, I assume the more optimistic HZ boundary, such that a runaway greenhouse occurs only once Venus is interior to the RV limit, Venus loses a maximum of 0.5 TO and builds up a maximum of 120 bars of O_2 in both regimes.

These figures are broadly consistent with previous estimates (Kasting et al., 1984; Kasting, 1988; Chassefière, 1996a,b; Kulikov et al., 2006; Gillmann et al., 2009), though the large uncertainty

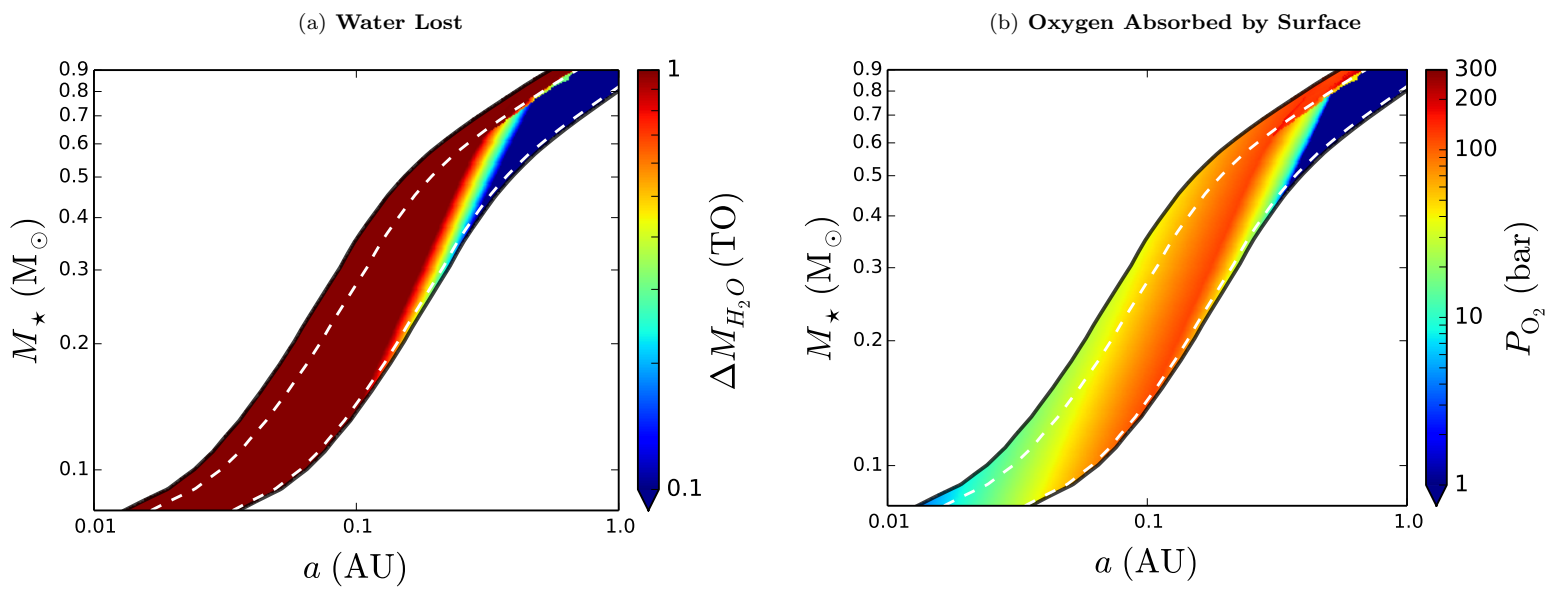

Figure 6: Total amount of water lost and amount of oxygen absorbed at the surface for a $1M_\oplus$ planet formed at 10 Myr with 1 TO of surface water, assuming the planet is in a runaway interior to the RG limit, the oxygen is instantaneously absorbed by the surface, and the escape is energy-limited. The solid lines are the empirical HZ bounds; the dashed lines are the theoretical HZ bounds. See Figure 7 for more details.

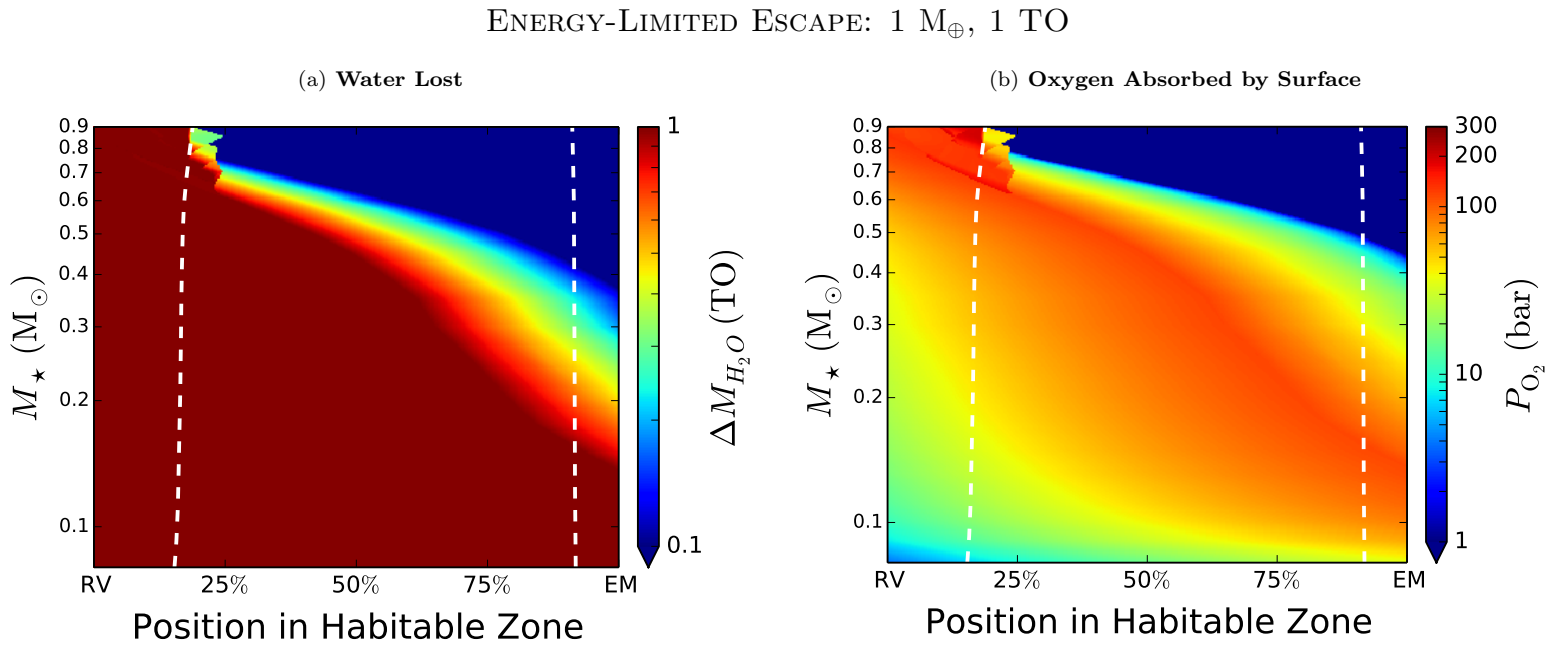

Figure 7: Same as Figure 6 (1 TO, energy-limited), but in an expanded view of the HZ. The axes correspond to the stellar mass (vertical) and the position of the planet within the HZ at 5 Gyr (horizontal). The "position in habitable zone" is the fractional distance between the RV limit and the EM limit (the empirical HZ). The dashed lines once again represent the RG and MG limits. **(a)** Total water lost in TO after 5 Gyr. Dark blue corresponds to less than 0.1 TO; dark red corresponds to complete desiccation. Most planets in the HZ of M dwarfs are completely desiccated; conversely, those close to the outer edge of high mass M dwarfs and throughout most of the HZ of K dwarfs lose little or no water. Interior to the RG limit, planets around stars of all masses are completely desiccated. **(b)** Total amount of oxygen absorbed by the surface in bars. Dark blue corresponds to insignificant O_2 buildup; dark red corresponds to 200 bars of oxygen. Planets that lose significant amounts of water also undergo extreme surface oxidation.

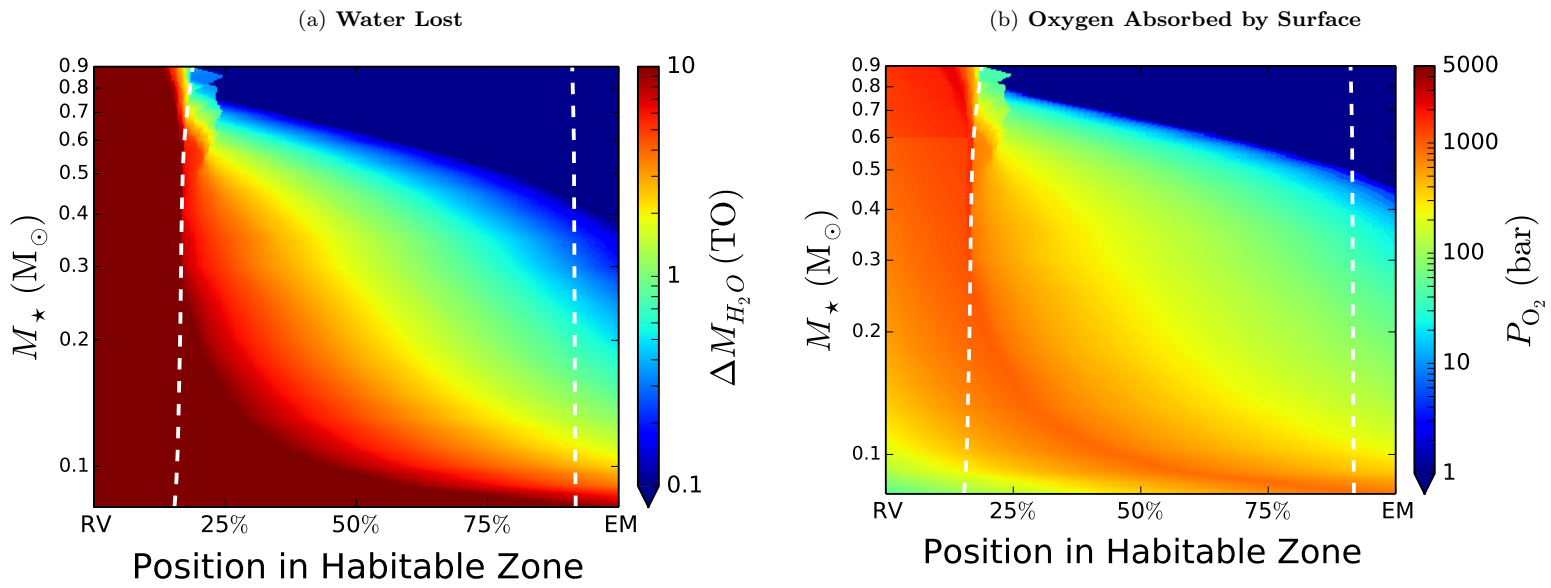

Figure 8: Same as Figure 7 (energy-limited escape), but for an initial water content of 10 TO. Note the change in the colorbar scales. Planets throughout most of the HZ of M dwarfs now lose at least 1 TO of water; those close to the RG limit and around low mass M dwarfs lose close to 10 TO. Most planets now retain several hundred to $\sim$ 1000 bars of O_2.

regarding the initial water content precludes an accurate determination of the amount of O_2 retained by the planet. Lammer et al. (2011b) argue that most of the O_2 that does not escape hydrodynamically could be removed by oxygen blow-off, but the exospheric temperatures may not be high enough for this to occur (see, e.g., Tian, 2009, and §2.6.4). In any event, the negligible O_2 content of the Venusian atmosphere today implies that the O_2 must have been removed either by surface sinks or other escape mechanisms. Given terrestrial rates of O_2 removal of a few hundred bars per Gyr (§2.2.2), my values are consistent with a maximum of a few TO of water on early Venus and therefore agree with the studies mentioned above.

2.5.2 Fast Oxygen Removal: Energy-Limited Escape

I first consider the limiting case where the rate of oxygen absorption by surface sinks is much larger than the rate at which it is photolytically produced. This could happen in the case of planets with vigorous resurfacing processes or convecting magma oceans (see §2.2.2.3 and §2.3 for a discussion). The oxygen content of the atmospheres of these planets is always low, the upper atmosphere is rich in water vapor, and hydrogen and oxygen escape at the energy-limited rate (§1.3.2.1).

In Figure 6 I plot the results for a $1M_\oplus$ planet with 1 TO of initial surface water, formed at 10 Myr, assuming a runaway greenhouse occurs interior to the RG limit. The axes are the same as in Figure 4, but the colors now indicate the total amount of water lost (left panel) and the equivalent pressure of O_2 absorbed by the surface in bars (right panel). While Figure 6 shows my results on a traditional HZ plot, I re-scale the x-axis to be the relative position in the (empirical) 5 Gyr HZ in Figure 7. This corresponds to zooming in on the HZ at each stellar mass in Figure 6 (note that the data plotted here are exactly the same). For reference, the dashed vertical lines once again indicate the locations of the RG (left) and MG (right) limits.

Planets throughout most of the HZ of M dwarfs are completely desiccated (left panel). Above $\sim 0.2M_\odot$, planets close to the outer edge of the HZ retain some of their water due to the shorter runaway phase; however, even planets in the center of the HZ of high mass M dwarfs ($M_\star \gtrsim 0.4M_\odot$) are completely desiccated. Planets around K dwarfs ($M_\star \gtrsim 0.6M_\odot$), on the other hand, lose significant amounts of water only close to the RG limit. Note, importantly, that planets that form *in situ* in the HZs of K dwarfs likely take longer than 10 Myr to form (Raymond et al., 2007); therefore, above $\sim 0.6M_\odot$, this figure applies only to planets that form quickly beyond the snow

line and migrate into the HZ. Finally, although the figures encompass only M and K dwarfs, I note that planets in the theoretical HZ of solar-mass stars are never in an insolation-induced runaway greenhouse and therefore lose no water via this mechanism.

Planets that lose significant water also retain on the order of 100 bars of O_2 (right panel), which by assumption is quickly absorbed by the surface. Such a high oxidative power could have strong implications for planetary evolution and habitability, which I discuss in §2.6. Perhaps surprisingly, the maximum oxygen pressure actually occurs in the *center* of the HZ of mid- to high mass M dwarfs and close to the *outer* edge for low mass M dwarfs. This is because close to the inner edge of the HZ, the ocean is lost quickly and the planet is desiccated early on, when the XUV flux is high and oxygen escape is efficient. Close to the outer edge, the RG phase is shorter, resulting in weaker ocean loss and similarly low O_2 amounts. Therefore, for each stellar mass, there exists a certain distance at which the O_2 pressure peaks. I discuss this trend further in §2.5.5.

Next, in Figure 8 I repeat the previous calculations for an initial water content of 10 TO. Complete desiccation now occurs only around low mass M dwarfs, particularly close to the RG limit. However, most M dwarf planets lose > 1 TO throughout the HZ and $\gtrsim 5$ TO close to the RG limit. Moreover, the equivalent pressure of oxygen absorbed by the surface is now on the order of several hundred to ~ 1000 bars, peaking in the outer HZ of the lowest mass M dwarfs and close to the RG limit of $0.3 M_\odot$ M dwarfs.

Finally, I note that there is an interesting feature near the RG limit of K dwarfs in these figures. Consider Figure 7, for instance: above $0.6 M_\odot$, both the total water loss and the O_2 amount change discontinuously as the stellar mass increases, leading to a jagged pattern near the RG limit and high O_2 buildup. This behavior is rooted in the non-monotonic luminosity evolution of K dwarfs prior to ~ 100 Myr. In Figure 9 I plot the evolution of the bolometric flux received by a planet near the inner edge of the HZ (defined at 5 Gyr), in units of the critical runaway greenhouse flux; values above 1 (red shading) correspond to a runaway greenhouse state. The thick curves indicate the flux evolution for different stellar masses (blue: a late M dwarf; green: an early M dwarf; red: a K dwarf; cyan: a solar-type G dwarf). Note that while M dwarfs dim monotonically during their PMS contraction phase, stars of type K and earlier display a bump in their luminosity prior to 100 Myr. This is due to the fact that stars more massive than about 0.6 $M_\odot$ switch from convective to radiative energy transport towards the end of their contraction phase, during which time their

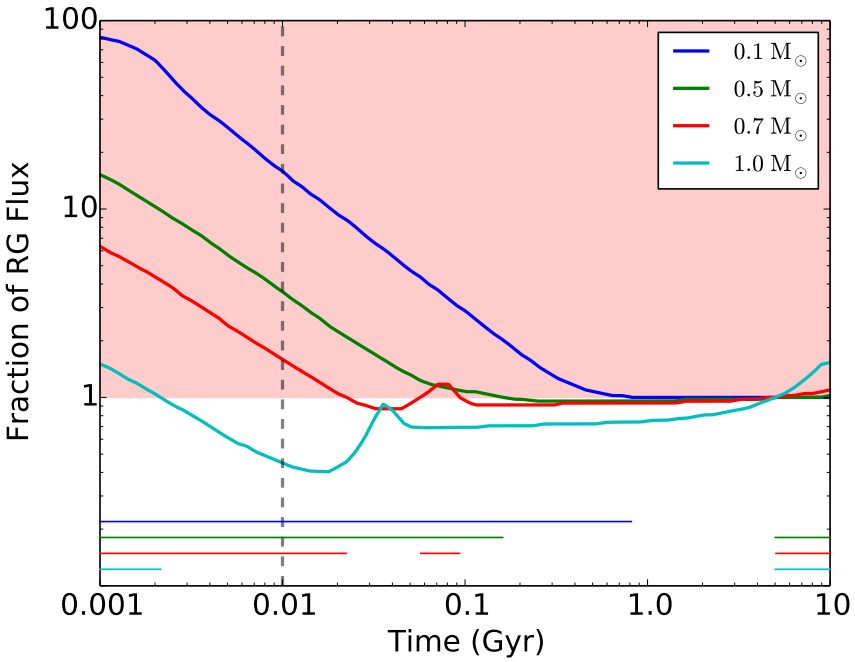

Figure 9: Evolution of the flux received by planets that are at the inner edge of the theoretical HZ at 5 Gyr for different stellar masses (blue: $0.1M_{\odot}$, green: $0.5M_{\odot}$, red: $0.7M_{\odot}$, cyan: $1.0M_{\odot}$). The vertical axis is the bolometric flux normalized to the runaway greenhouse flux; the region shaded in pink corresponds to planets that are interior to the RG limit. The dashed vertical line corresponds to the 10 Myr formation time I assume in my model. The horizontal lines at the bottom of the plot indicate the times during which planets are in a runaway greenhouse. See text for a discussion.

effective temperatures rise, leading to a temporary increase in L just before reaching the MS (see, e.g., Reid and Hawley, 2005). M dwarfs, which remain mostly convective even once they reach the MS, do not display such a bump.

For a star like the Sun (cyan curve), the bump in the luminosity does not significantly affect planets in the HZ, since the flux on a planet near the inner edge does not exceed the critical flux during that time. For K dwarfs, however, this is not true; the flux on a planet at the inner edge of the HZ around a 0.7 $M_{\odot}$ star exceeds the critical flux for $\sim$ 30 Myr, leading to two distinct runaway greenhouse episodes prior to 5 Gyr.

For reference, the thin lines at the bottom of the plot indicate when planets are in the runaway regime. While M dwarf planets experience the longest runaway greenhouse episodes, planets that form prior to about 20 Myr in the inner HZ of K dwarfs undergo two runaways, the second of

which occurs late enough that the XUV flux has largely tapered off, leading to efficient oxygen buildup.

2.5.3 Inefficient Oxygen Sinks: Diffusion-Limited Escape

In the previous section I assumed that the rate of oxygen removal at the surface was much higher than the rate at which oxygen was produced. I now repeat all calculations in the opposite limit, assuming the production rate is much higher than the absorption rate. As discussed in §2.2.2, this could be the case for water worlds or for planets with a pre-oxidized surface/interior, inefficient outgassing of reducing compounds, inefficient resurfacing processes, etc. In these runs, the oxygen remains in the atmosphere and hydrogen must diffuse through it in order to escape. I therefore calculate loss rates in the diffusion limit as described in §1.3.2.3.

My results are plotted in Figures 10 and 11 for initial surface water contents of 1 and 10 TO, respectively. Water loss amounts are generally slightly lower (compare to Figures 7 and 8), since the diffusion limit is slower than the energy-limited escape rate at early times. Nevertheless, planets throughout a large portion of the HZ of M dwarfs are still desiccated in the 1 TO case. Planets with larger water inventories lose $>$ 1 TO around low mass M dwarfs and close to the inner edge of high mass M dwarfs.

In the panels on the right, it is clear that oxygen buildup is larger than in the energy-limited case; recall that these panels now represent the equivalent pressure of oxygen *in the atmosphere* at the end of the runaway phase. Despite the subdued water loss, the fact that no oxygen escapes leads to the buildup of $\sim$ 240 bars of O_2 (the equivalent pressure of O_2 in 1 TO of water) throughout a large portion of the HZ in Figure 10. In Figure 11, planets build up between $\sim$ 100 and $\sim$ 1000 bars of O_2 in their atmospheres throughout most of the HZ of M dwarfs. Note also that, unlike in the previous figures, the oxygen amount is a monotonic function of the position in the HZ; because no oxygen escapes, planets closer to the inner edge build up more O_2.

2.5.4 Higher Planet Mass

In Figure 12 I show the results for a 5 $M_\oplus$ super-Earth with an initial water inventory of 10 TO and efficient oxygen sinks. Compared to a 1 $M_\oplus$ planet (Figure 8), both water loss and oxygen amounts are significantly *higher*. Thousands of bars of O_2 are now retained throughout most of

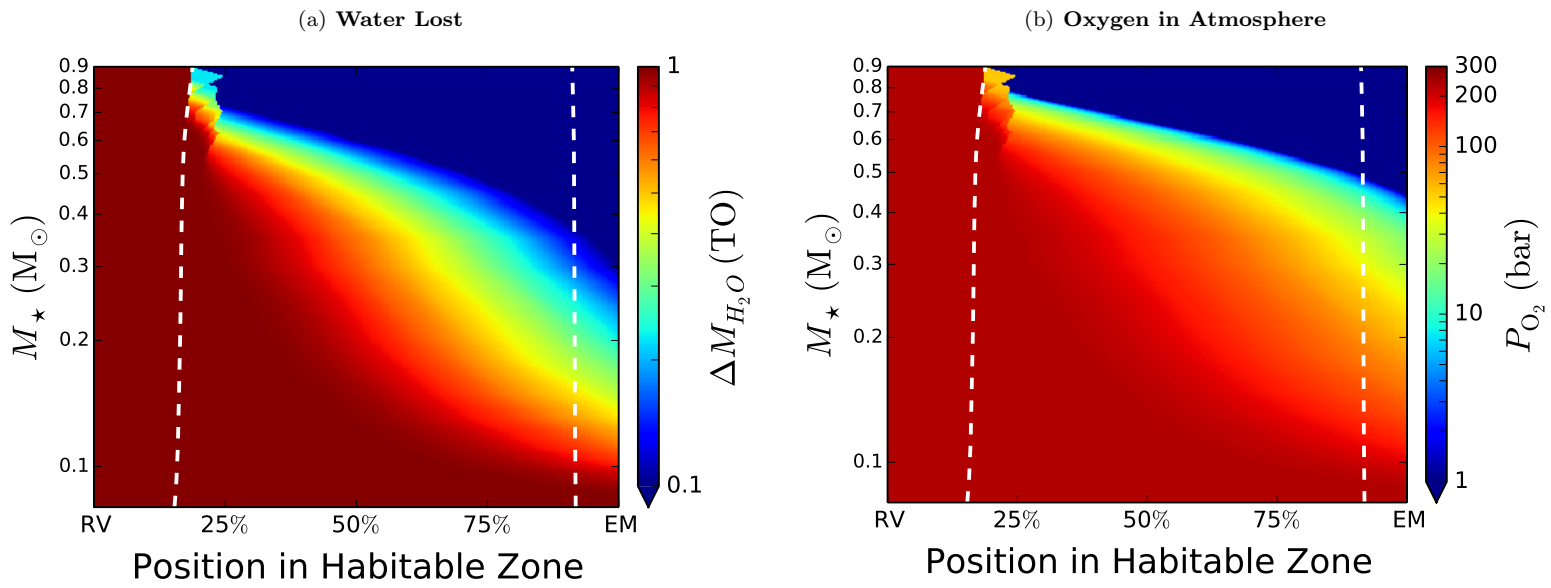

Figure 10: Similar to Figure 7 (1 TO), but assuming that the escape of hydrogen is diffusion-limited. This corresponds to planets with slow/ineffective oxygen sinks that retain all of the photolytically-produced O_2 in their atmospheres. While water loss amounts are generally lower than in the energy-limited case, planets throughout a large fraction of the HZ of M dwarfs are still desiccated. Moreover, the amount of oxygen that builds up is substantially greater than in Figure 7, since the oxygen cannot escape if the loss of hydrogen is diffusion-limited. Thus, planets that lose 1 TO of water build up $\frac{16}{18} \times 270 = 240$ bars of O_2 in their atmospheres.

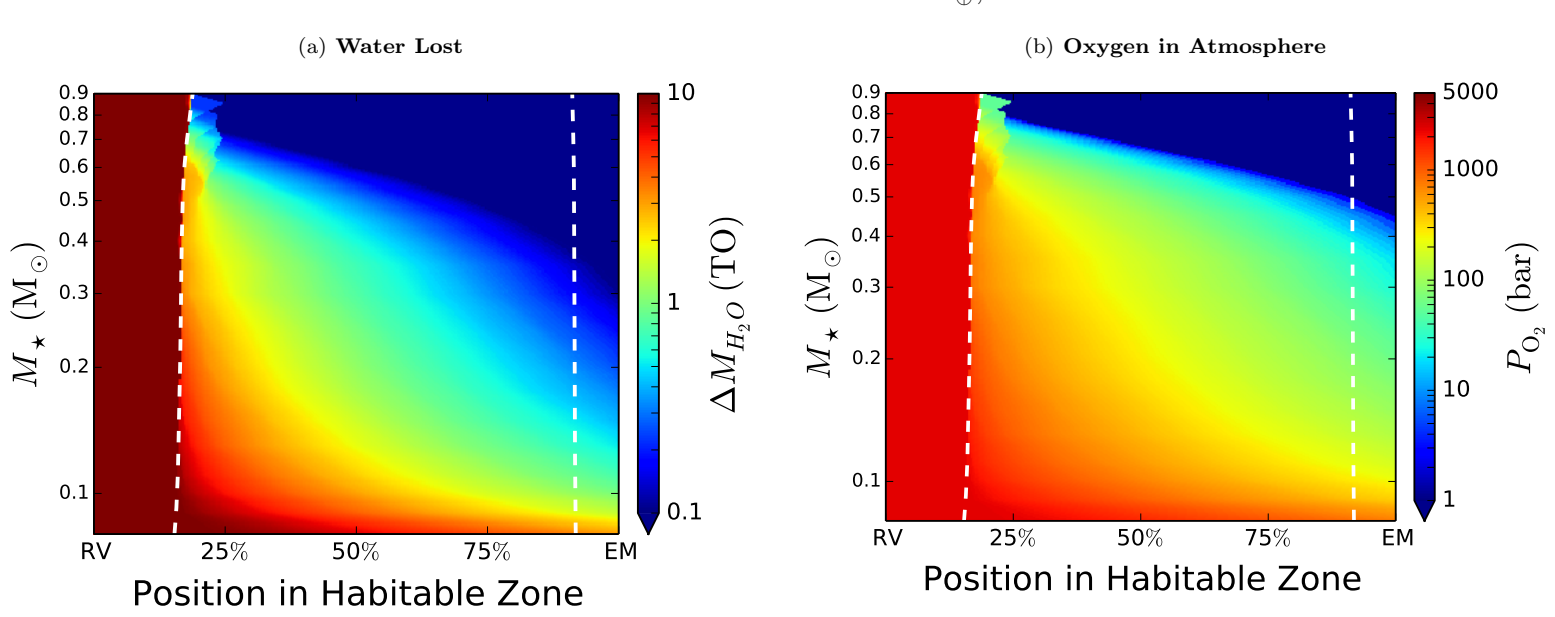

Figure 11: Same as the previous figure (diffusion-limited escape), but for an initial water content of 10 TO; compare to Figure 8, the corresponding energy-limited case. Once again, water loss amounts are smaller, but oxygen amounts greater than 2000 bars are now possible around the lowest mass M dwarfs.

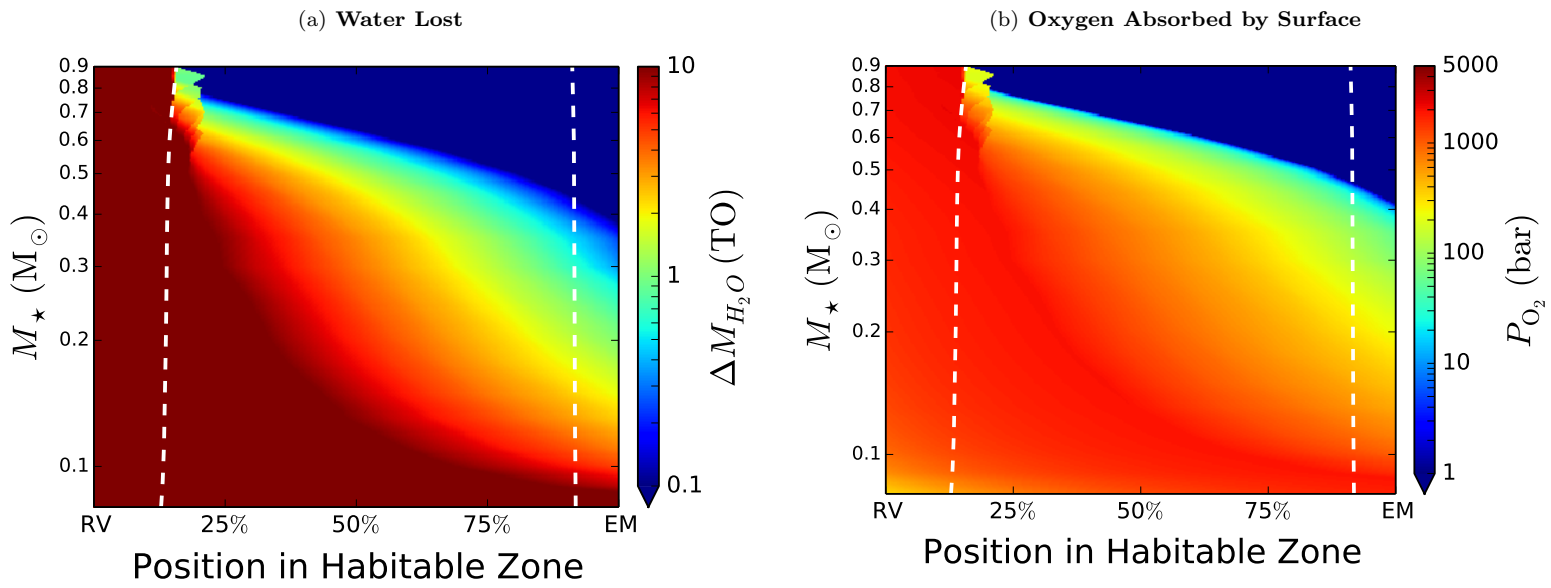

Figure 12: Similar to the previous figures, but for a super-Earth of mass $5M_\oplus$ with 10 TO, assuming energy-limited escape. Note the large fraction of the HZ in the lower left portion of the left panel where planets are completely desiccated. Elsewhere, super-Earths lose several TO of water. In the right panel, thousands of bars of O_2 are absorbed at the surface of planets throughout most of the HZ of M dwarfs.

Diffusion-Limited Escape: 5 $M_\oplus$, 10 TO

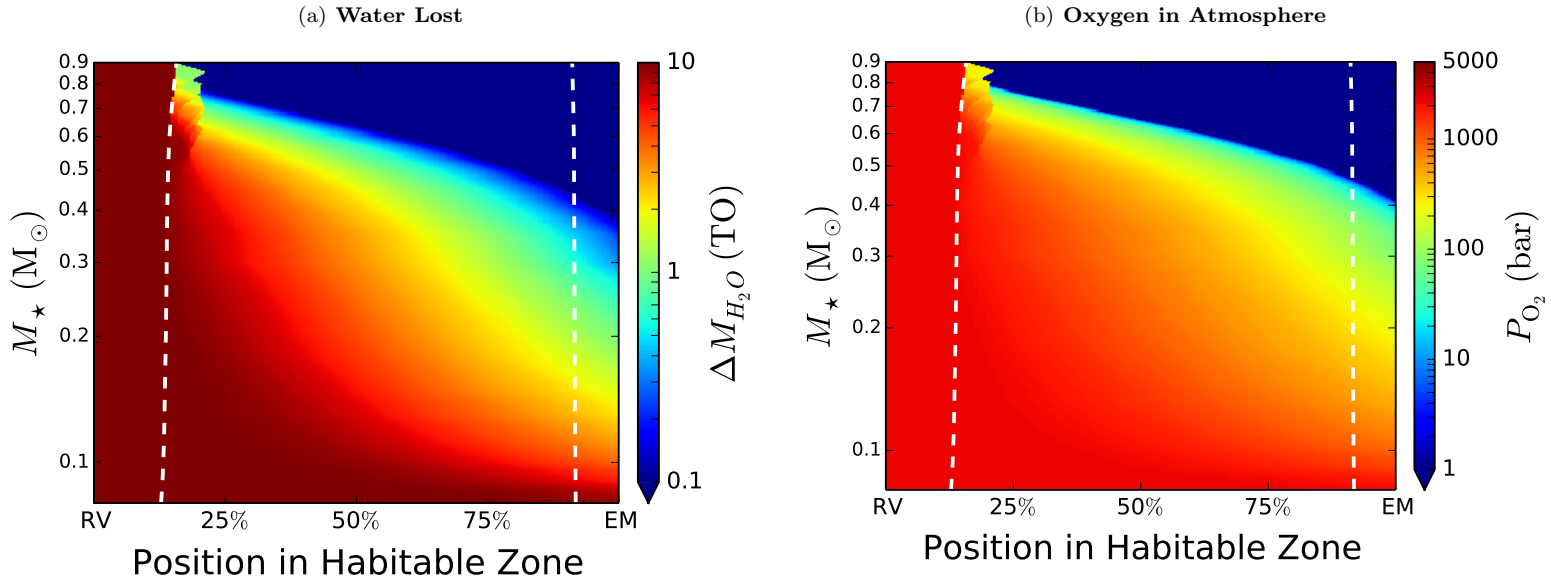

Figure 13: Similar to Figure 12 (a 5 $M_\oplus$ super-Earth with 10 TO of surface water), but assuming diffusion-limited escape. Despite a slight decrease in the total water loss amounts, the results are very similar to those in the energy-limited case. In general, these planets lose several TO of water and build up several hundred to a few thousand bars of O_2 in their atmospheres.

the HZ. Close to the inner edge of the lowest mass M dwarfs, planets with larger initial water inventories can lose several tens of TO.

This perhaps counter-intuitive result stems from the fact that the crossover mass (Equation 19) scales inversely with the surface gravity; on super-Earths, the escape of oxygen is greatly suppressed, leading to faster buildup at the surface. Moreover, when the oxygen does not escape, the loss of the ocean happens more quickly, since the escape of hydrogen is nine times more efficient at depleting the water content of the planet (at fixed $\mathcal{F}_{\rm XUV}$). Recall that in the limit $\eta \to 0$, Equation (32) gives $\dot{m}_{\rm ocean} = 9\dot{M}_{\rm EL}$.

In Figure 13, I perform the same calculation, but for diffusion-limited escape, assuming the O_2 remains in the atmosphere. Once again, both the amount of water lost and the oxygen pressure are substantially higher than in the 1 $M_\oplus$ case. This is a straightforward consequence of Equation (33); the diffusion limit flux scales inversely with the scale height of the background atmosphere, which is smaller on the super-Earth by a factor of ~ 2.2 due to the higher surface gravity.

Perhaps even more interestingly, Figures 12 and 13 are very similar; the desiccation process of super-Earths is relatively insensitive to the escape regime and to whether or not the oxygen remains in the atmosphere or is absorbed by the surface. This similarity is due to the fact that the energy-limited escape rate and the diffusion-limited escape rate are comparable for the range of XUV fluxes received by these planets. This applies even at early times, when the XUV flux is very high; the escape of even a small amount of oxygen in the energy-limited regime tends to slow down the rate of ocean loss, given that a large fraction of the XUV energy goes into driving the escape of the heavier species.

It is important to note that since I have been plotting the O_2 equivalent *pressure* rather than the actual amount, care must be taken when comparing it between planets of different masses. Under the plane-parallel approximation, the atmospheric pressure scales as $M_{\rm p}/R_{\rm p}^4$. Given that a $5M_\oplus$ planet with an Earth-like composition has $R_{\rm p} = 1.52R_\oplus$ (Fortney et al., 2007), the pressure exerted by a fixed amount of O_2 will actually be *smaller* (by about 6%) on the higher mass planet. Nevertheless, these figures show that the O_2 pressures are a factor of $2-3$ times *larger* on a super-Earth than on an Earth-mass planet, implying substantially more O_2 buildup.

For reference, in Figure 14 I overplot a handful of known planets on the oxygen pressure contours of Figure 13 (b): GJ 581g[7] (Vogt et al., 2010), Kepler 61b (Ballard et al., 2013), Kepler

[7]The existence of GJ 581g has recently been contested; see Robertson et al. (2014).

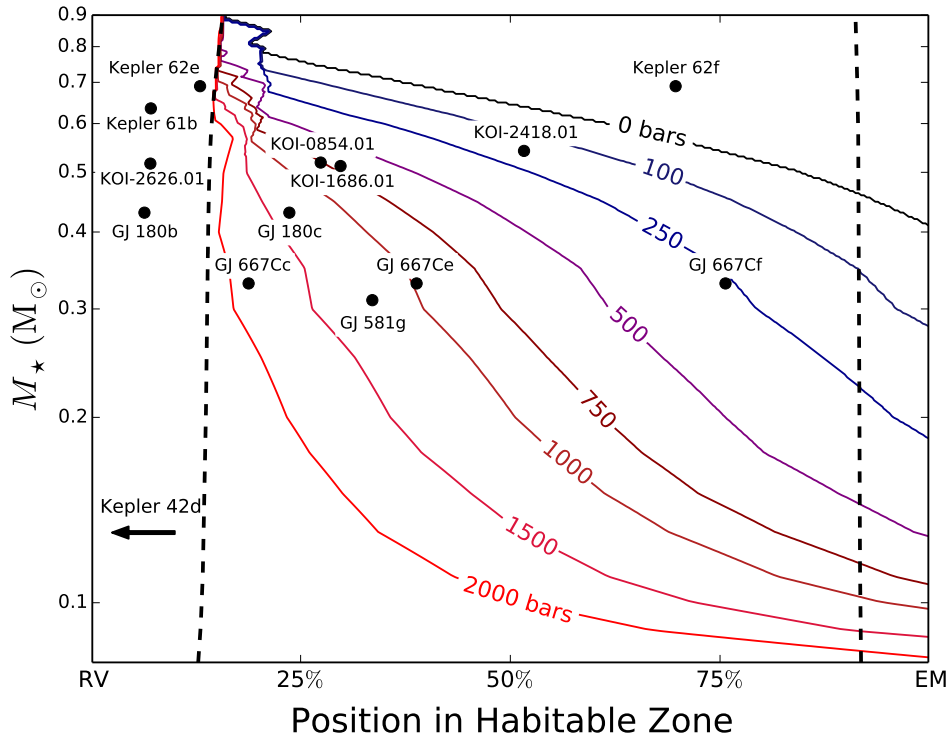

Figure 14: A selection of M and K dwarf super-Earths that could have detectable O_2 atmospheres if they formed with abundant surface water; as before, the dashed lines represent the RG and MG limits. Calculations were performed assuming a mass of 5 $M_\oplus$, an initial water content of 10 TO, and diffusion-limited escape. Contours correspond to the equivalent O_2 atmospheric pressure in bars at the end of the runaway phase, assuming all of the O_2 remains in the atmosphere. Of the planets shown here, only Kepler 62f does not build up any oxygen.

62e & f (Borucki et al., 2013), GJ 180c (Tuomi et al., 2014), the GJ 667 system (Anglada-Escudé et al., 2013), and four Kepler candidate planets (Dressing and Charbonneau, 2013); Kepler 42d (Muirhead et al., 2012), which is interior to the HZ, is indicated by an arrow. As in Figure 13, I assume that the planet mass is 5 $M_\oplus$, the oxygen remains in the atmosphere, and the escape is diffusion-limited. Provided they formed with abundant water, many of the currently known super-Earths could have built up hundreds to thousands of bars of O_2. In particular, this could be the case for GJ 667Cc, which could have lost as many as 10 TO early on, accumulating close to 2000 bars of O_2; as a result, it may not be habitable today.

Whether this oxygen remains in these planets' atmospheres past the early runaway phase depends on the efficiency of their surface sinks. Given poor constraints on exoplanet tectonics, the

physics of oxygen absorption by a magma ocean, and other aspects of exoplanet atmospheres, it is reasonable to expect that some super-Earths may be unable to remove all the photolytically-produced oxygen within the ages of their systems.

2.5.5 The Rate of Oxygen Buildup

While the final O_2 equivalent pressure is a complex function of the stellar/planetary mass and the semi-major axis, my results in the energy-limited regime can be understood in fairly simple terms by considering the mass loss rates (Equations 29–31) derived in the previous chapter. In particular, the rate of oxygen accumulation (Equation 31) is a function of both η and $\dot{M}_{\rm EL}$, which are themselves functions of the XUV flux. By expressing $\dot{M}_{\rm EL}$ as a function of η, it is possible to show that the rate at which oxygen accumulates in the atmosphere/at the surface is *completely independent* of the XUV flux above the critical value given in Equation (20). To demonstrate this, consider Equation (31), the rate of oxygen buildup in the atmosphere in the energy-limited regime. Since both η and $\dot{M}_{\rm EL}$ are functions of $\mathcal{F}_{\rm XUV}$, we will rearrange this expression in order to make the dependence more explicit. First, we solve Equation (22) for the XUV flux, making use of Equations (23) and (16):

$$\mathcal{F}_{\rm XUV} = \left(\frac{40 G^2 m_{\rm H}^2 b M_{\rm p}^2 K_{\rm tide}}{kT \epsilon_{\rm XUV} R_{\rm p}^3} \right) \frac{1+8\eta}{1-\eta}, \tag{60}$$

The energy-limited mass escape rate (Equation 5) may then be written

$$\dot{M}_{\rm EL} = \left(\frac{40 \pi G m_{\rm H}^2 b M_{\rm p}}{kT} \right) \frac{1+8\eta}{1-\eta}. \tag{61}$$

Plugging this into Equation (31), one obtains $\dot{m}_{\rm O}^{\rm atm}(\eta)$:

$$\dot{m}_{\rm O}^{\rm atm} = \left(\frac{320 \pi G m_{\rm H}^2 b M_{\rm p}}{kT} \right), \tag{62}$$

which is independent of η. This means that *the rate of oxygen buildup is constant in time* and does not vary with the XUV flux (provided $\mathcal{F}_{\rm XUV} > \mathcal{F}_{\rm crit}$). Instead, it depends only on the planet mass and the temperature of the flow. For XUV fluxes below $\mathcal{F}_{\rm crit}$, the oxygen escape rate is zero and the rate of buildup in the atmosphere scales linearly with the flux; one must calculate this

directly from Equation (31).

This result may seem very counter-intuitive, since it implies that the rate of oxygen buildup is independent of the XUV flux and relatively independent of the escape regime. In particular, one might expect that in the energy-limited regime, an increase in the XUV flux would lead to more oxygen escape and thus a slower rate of buildup. However, increasing $\mathcal{F}_{\rm XUV}$ also leads to a higher *hydrogen* escape rate and a faster net production of O atoms. While the ratio of the oxygen to hydrogen mass escape rates in Equation (28) approaches a maximum value of 8 for $\eta \to 1$, the *difference* between the oxygen production and escape rates remains constant.

The constant buildup rate of O_2 may be understood more easily if we consider the flux of oxygen atoms into the atmosphere. If we divide Equation (62) by $4\pi R_{\rm p}^2 m_{\rm H}$, we obtain the rate at which oxygen atoms build up in the atmosphere:

$$F_{\rm O}^{\rm atm} = 5bgm_H/kT, \tag{63}$$

which is exactly one-half the diffusion limit for hydrogen when $X_{\rm O}/X_{\rm H} = 1/2$. If we momentarily consider the hydrodynamic flow from the frame of the escaping hydrogen particles, we see that in this case it is the *oxygen* that is diffusing through a static hydrogen background, albeit downwards. The diffusion limit is, from Equation (33),

$$\phi_{\rm O} = \frac{bg(m_{\rm H} - m_{\rm O})}{kT(1 + X_{\rm H}/X_{\rm O})} \tag{64}$$

$$= -5bgm_{\rm H}/kT, \tag{65}$$

where the negative sign indicates a downward flux. This result is no coincidence; it implies that oxygen is retained *at its diffusion limit*. In other words, in order for the oxygen to accumulate in the atmosphere, it must diffuse out of the hydrodynamic flow that is attempting to carry it away, and the rate at which it can do so is capped at the diffusion limit.

When the escape of hydrogen is diffusion-limited instead of energy-limited, the rate at which oxygen builds up in the atmosphere is again one-half the diffusion limit (since one oxygen atom is left behind for every two escaping hydrogen atoms). Initially, therefore, the rate of buildup is the same as in the energy-limited case. However, as $X_{\rm O}/X_{\rm H}$ begins to increase, both the H escape rate

and the O_2 buildup rate decrease. Nonetheless, only once oxygen becomes the dominant species in the atmosphere does the rate at which it is produced begin to taper off.

When the O_2 buildup rate is constant, the expressions for the escape rates of hydrogen (Equation 29) and oxygen (Equation 30) and the expression for the rate of ocean loss (Equation 32) greatly simplify. It can be shown from Equations (24) and (26) that, provided $\mathcal{F}_{\rm XUV} > \mathcal{F}_{\rm crit}$,

$$\dot{m}_{\rm H}^{\uparrow} = \frac{1}{9}\dot{M}_{\rm EL} + C \tag{66}$$

$$\dot{m}_{\rm O}^{\uparrow} = \frac{8}{9}\dot{M}_{\rm EL} - C \tag{67}$$

$$\dot{m}_{\rm ocean} = \dot{M}_{\rm EL} + C \tag{68}$$

where

$$C = \frac{\dot{m}_{\rm O}^{\rm atm}}{9} = \frac{320\pi G m_{\rm H}^2 b M_{\rm p}}{9kT}. \tag{69}$$

We can also derive the rate of change of the equivalent oxygen pressure with time, assuming $X_{\rm O}/X_{\rm H} = 1/2$. Noting that

$$\dot{P}_{\rm O_2} \approx \frac{GM_{\rm p}}{4\pi R_{\rm p}^4}\dot{m}_{\rm O}^{\rm atm}, \tag{70}$$

we obtain the following from Equations (62) and (31):

$$\dot{P}_{\rm O_2} = \begin{cases} 5.35\left(\dfrac{M}{\rm M_\oplus}\right)^2\left(\dfrac{R}{\rm R_\oplus}\right)^{-4} \text{ bars Myr}^{-1} & \text{if } \mathcal{F}_{\rm XUV} \geq \mathcal{F}_{\rm crit} \\ 0.138\left(\dfrac{\mathcal{F}_{\rm XUV}}{\mathcal{F}_\oplus}\right)\left(\dfrac{R}{\rm R_\oplus}\right)^{-1}\left(\dfrac{\epsilon_{\rm XUV}}{0.30}\right) \text{ bars Myr}^{-1} & \text{if } \mathcal{F}_{\rm XUV} < \mathcal{F}_{\rm crit}. \end{cases} \tag{71}$$

In Figure 15 I plot η, $\dot{P}_{\rm O_2}$, and $\dot{m}_{\rm ocean}$ as a function of $\mathcal{F}_{\rm XUV}$ for two planet masses and two XUV absorption efficiencies. Note the constant value of $\dot{P}_{\rm O_2}$ and the linear behavior of $\dot{m}_{\rm ocean}$ above $\mathcal{F}_{\rm crit}$.

Since this rate is constant in the energy-limited regime and declines slowly in the diffusion-limited regime, the primary factor controlling the final oxygen amount retained by the planet is the duration of the hydrodynamic escape phase. In general, an Earth-mass planet will build up oxygen at a constant rate of about 5 bars/Myr until (a) the planet leaves the runaway regime

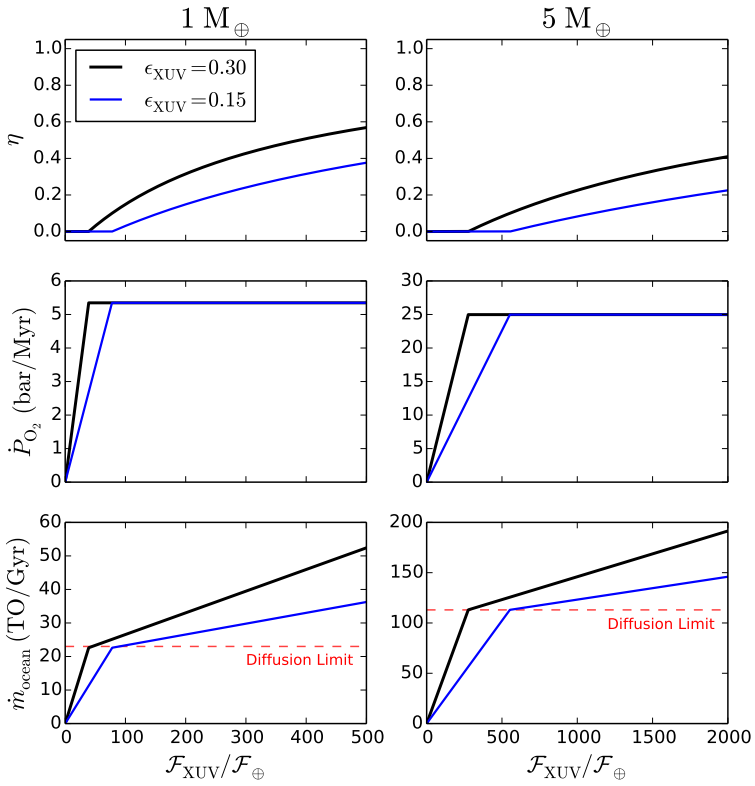

Figure 15: Dependence of the oxygen escape parameter η (top), the rate of oxygen buildup $\dot{P}_{O_2}$ (center), and the ocean loss rate $\dot{m}_{ocean}$ (bottom) on the XUV flux for a 1 $M_\oplus$ Earth (left) and a 5 $M_\oplus$ super-Earth (right) in the energy-limited regime. Results for two different XUV efficiencies are plotted: 0.30 (black) and 0.15 (blue). For $\eta = 0$, corresponding to $\mathcal{F}_{XUV} < \mathcal{F}_{crit}$, the rates of oxygen buildup and ocean loss are linear in $\mathcal{F}_{XUV}$. For $\eta > 0$, the oxygen buildup rate is constant at $\sim$ 5 bar/Myr for the Earth and $\sim$ 25 bar/Myr for the super-Earth. The rate of ocean loss is still linear in $\mathcal{F}_{XUV}$, but increases more slowly. For reference, in the bottom panel I plot the diffusion-limited ocean loss rate as a dashed red line.

(enters the HZ); (b) enough O_2 builds up in the atmosphere to slow the H escape rate; (c) the planet loses all surface water; or (d) the XUV flux drops below $\mathcal{F}_{crit}$.

This helps shed light on the behavior seen in several of the figures in the previous sections. Consider, for instance, Figure 7, where the final oxygen pressure is not a monotonic function of either $M_\star$ or the position in the HZ, peaking close to the outer edge of the HZ in some cases. This trend is shown more clearly in Figure 16, where I plot a cross-section along $M_\star = 0.4 M_\odot$ in Figures 7 (a) and (b). The blue curve corresponds to the amount of water lost and the red curve corresponds to the final oxygen pressure.

Since $M_p = 1 M_\oplus$ and the escape regime is energy-limited, all planets build up O_2 at the same

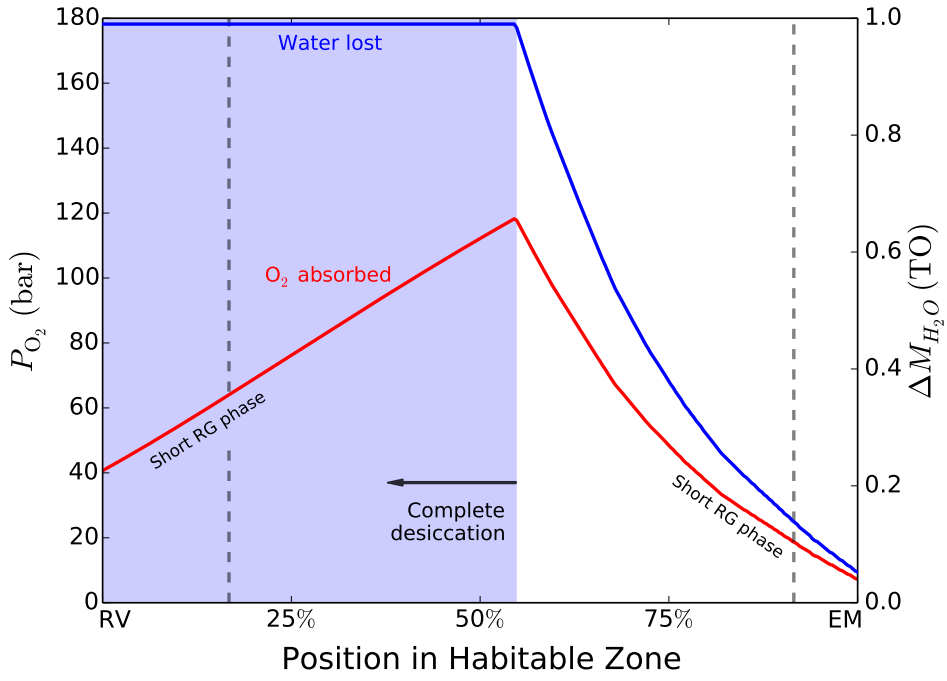

Figure 16: Cross-section along $M_\star = 0.4 M_\odot$ in Figures 7 (a) and (b), showing the oxygen absorbed by the surface (red) and total amount of water lost (blue) as a function of the position in the HZ. Water loss scales with the time spent in the runaway phase; interior to a critical distance, complete desiccation occurs (shaded region). From right to left, the O_2 amount initially increases due to the increase in the amount of water lost. However, interior to the critical distance, complete desiccation occurs at progressively earlier times. The higher XUV flux early on results in more oxygen escape and less buildup. See text for a discussion.

rate of 5.35 bars/Myr. The difference in the final value of P_{O_2} is solely due to the duration of the escape. Towards the outer edge of the HZ, planets experience short-lived runaway greenhouse phases; see Figure 4 (b). As one moves from the outer HZ to the inner HZ, the increasing duration of the RG phase leads to higher O_2 pressures.

However, at a certain critical position in the HZ ($\sim 55\%$ in Figure 16), the final O_2 pressure begins to decrease. This is because planets in the shaded region are completely desiccated, which terminates the escape phase and sets an upper limit to the O_2 pressure. Interior to this distance, faster water loss (blue line) results in an earlier desiccation time, leading to lower O_2 amounts.

I note, finally, that the rate of O_2 buildup is constant only when the XUV flux exceeds $\mathcal{F}_{\rm crit}$. Below this flux, the escape of hydrogen is energy-limited, the oxygen escape rate is zero and its rate of buildup is linearly proportional to the H escape flux. However, nearly all Earth-mass planets in my runs enter the HZ before the XUV flux drops below the critical value. Compare Figures 2 and

9, noting that planets around M dwarfs remain in a runaway for $0.1 - 1$ Gyr, during which time $\mathcal{F}_{\rm XUV} \gtrsim 300\mathcal{F}_{\oplus}$ for a saturation time of 1 Gyr. For a shorter saturation time (typical of K and G dwarfs) and/or a higher planet mass (for which the critical flux is higher), planets may leave the oxygen escape regime prior to the end of the runaway greenhouse.

2.6 Discussion

2.6.1 Prolonged Runaway Greenhouse

I showed in the first part of this chapter that the pre-main sequence evolution of M dwarfs causes their habitable zones to migrate inwards by nearly an order of magnitude in semi-major axis during the first few hundred Myr. This means that many of the planets currently in the HZs of M dwarfs were not in the HZs when they formed. The fact that these planets spend several tens to several hundreds of Myr in insolation-induced runaway greenhouses establishes initial conditions in stark contrast to those on Earth, the only known habitable planet to date. This could significantly compromise their habitability.

Because of the extended runaway greenhouse, the atmospheric and thermal evolution of many M dwarf planets could follow very different paths than Earth's. Chemical equilibrium between the atmosphere and a magma ocean during a runaway could lead to the buildup of hundreds of bars of CO_2 early on (Elkins-Tanton, 2008). High surface temperatures could also inhibit plate tectonics by promoting rapid lithospheric healing and grain growth, increasing the viscosity and erasing weak zones where plate subduction can occur (e.g., Driscoll and Bercovici, 2013). This may prevent the onset of a carbonate-silicate cycle on these planets, making them unable to remove atmospheric CO_2 and maintaining permanently high surface temperatures.

Currently, the closest analog to the planets I consider here is Venus, which as a result of a prolonged runaway greenhouse phase underwent complete desiccation, irreversible crustal oxidation, termination of plate tectonics, and the buildup of a dense CO_2 atmosphere which maintains the surface temperature over 700 K. Whether this is the fate of all planets that undergo prolonged runaway greenhouses is unclear. Earth itself is thought to have been in an impact-induced runaway greenhouse for a few Myr (Zahnle et al., 1988; Hamano et al., 2013), yet it is habitable today. Detailed models coupling the atmosphere to the interior evolution of planets around M dwarfs should be performed in the future in order to address the long-term consequences of the early RG

phase.

2.6.2 The Fate of the Oxygen

In the second part of the results, I showed that the prolonged runaway greenhouse can lead to the loss of several TO of water and the buildup of hundreds to thousands of bars of O_2 on M dwarf planets that end up in the HZ. I considered two limiting cases regarding the escape: energy-limited and diffusion-limited. I argued that the former applies to planets with efficient surface sinks for the O_2, as the sinks prevent oxygen from becoming a major atmospheric gas and thus bottlenecking the escape of hydrogen via diffusion. The latter applies to the opposite case, where the O_2 builds up quickly and caps the hydrogen escape rate at the diffusion limit.

In reality, a planet with an Earth-like redox state is likely to start out in the energy-limited regime and transition to the diffusion-limited regime as its surface sinks get overwhelmed; my results for water loss and O_2 buildup should therefore bracket these processes on many M dwarf planets. In fact, any Earth-mass planet that removes oxygen at a rate slower than ~ 5 bars/Myr (~ 25 bars/Myr for a super-Earth) will build up O_2 in its atmosphere during the runaway phase and likely transition to diffusion-limited escape. Moreover, *any* major atmospheric gas will tend to cap the escape at the diffusion limit. If species such as CO_2 or N_2 are abundant relative to H_2O during the runaway, these would also enforce diffusion-limited escape.

Nevertheless, I have shown that even at the diffusion limit, extreme water loss and atmospheric O_2 buildup may occur on these planets. Given the ability of surface processes to remove tens to hundreds of bars of O_2 per Gyr (§2.2.2), one might expect that some planets that build up oxygen atmospheres in excess of ~ 100 bar could retain detectable amounts of O_2 today. Such elevated quantities of O_2 are possible throughout the HZ of all M dwarfs, except near the outer edge of those more massive than about 0.5 $M_\odot$, where planets are in runaway greenhouses for only a few Myr.

It is important to note that, because of the potential volatile deficiency of planets that form *in situ* in the HZs of M dwarfs (Raymond et al., 2007; Lissauer, 2007), many of the planets in the HZ with abundant surface water early on could be different from Earth in composition, given that they may have formed beyond the snow line and migrated in. In particular, they could have large inventories of highly reducing compounds, since hydrogen-rich species such as methane and

ammonia condense at large distances from the host star and are thus easily accreted during planet formation. Not only might these planets have highly reducing surfaces, but outgassing of reduced compounds could also lead to the quick removal of atmospheric O_2. This is especially the case for planets with nonzero eccentricity around low mass M dwarfs, as strong tidal forces in the HZ can drive vigorous volcanism (Jackson et al., 2008; Barnes et al., 2013). On the other hand, whether these planets develop a magma ocean capable of absorbing most of the O_2 is unclear, as they may not have rocky surfaces, but instead a thick layer of water and ice extending down to great depths.

Another issue is the possibility that planets forming beyond the snow line could develop substantial hydrogen/helium envelopes before they migrate into the HZ. While the hydrogen may act as a direct sink for the oxygen (via the formation of H_2O), the envelope could also inhibit the escape of water in the first place. A dense enough envelope could effectively shield the surface and prevent the dissociation of H_2O, provided the water settles below the base of the hydrodynamic wind. However, planetary cores that accrete dense H/He envelopes may be unable to completely shed them, in which case the planet will never be terrestrial (and probably not habitable). In order for the envelope to prevent water loss and O_2 buildup and yet not preclude the planet's eventual habitability, it must bear the brunt of the XUV flux early on and dissipate before the flux drops too quickly and atmospheric escape becomes negligible. I explore this possibility in detail in the next chapter, noting that such evaporated cores could potentially be habitable.

As for planets that happen to form *in situ* with abundant water, one of the major factors controlling whether or not a long-term O_2 atmosphere will develop could be the interaction with a potential magma ocean (§2.2.2.3). Gillmann et al. (2009) and Hamano et al. (2013) argue that this process could be responsible for the lack of oxygen in the Venusian atmosphere today. It is entirely plausible that a deep magma ocean could remove most of the atmospheric O_2 on M dwarf planets, provided solidification of the mantle occurs only at the end of the runaway. If, on the other hand, solidification occurs after $\sim$ 10 Myr as it likely did on Venus (Lebrun et al., 2013), planets that spend longer than this amount of time in a runaway may accumulate a large fraction of the O_2 in their atmospheres.

Finally, I caution that I have not fully accounted for the role of photochemistry in moderating the production and escape of hydrogen atoms. I assumed that the escape rate is limited either by the available energy (§2.5.2) or by the rate of hydrogen diffusion (§2.5.3), but in reality the

escape may also be limited by the availability of free hydrogen (and oxygen) atoms at the base of the flow, which is dictated by the rate of photolysis and various chemical reactions among hydrogen and oxygen-bearing species. If, for instance, oxygen recombines quickly with hydrogen following photolysis, the hydrodynamic flow may be suppressed. Alternatively, the production of ozone (O_3) could shield lower atmospheric layers from UV photons necessary for water photolysis. Photochemistry could also act to cause hydrogen and oxygen to escape in stoichiometric proportions (Zahnle and Catling, 2017). However, a detailed hydrodynamic/photochemical treatment is necessary to address these issues and accurately track the amount of oxygen that remains in the planetary atmosphere.

2.6.3 Implications for Habitability

I have shown that during the early high luminosity phase of M dwarfs, terrestrial planets can lose several Earth oceans of water. Over 10 oceans may be lost close to the inner edge of the HZ, particularly for super-Earths, whose hydrogen diffusion limit is higher, and for planets with efficient O_2 absorption processes, as these could prevent atmospheric O_2 from building up, resulting in an escape rate that approaches the energy-limited rate. Given the fundamental importance of water to life as we know it, complete planetary desiccation could severely hamper the ability of life to originate and evolve on planets around M dwarfs. Moreover, water is an essential ingredient for both plate tectonics (Mackwell et al., 1998; Moresi and Solomatov, 1998) and the carbonate-silicate cycle, which together regulate CO_2 in the Earth's atmosphere. Without a mechanism to remove atmospheric CO_2, desiccated planets may build up dense CO_2 atmospheres and maintain high surface temperatures even after the end of the greenhouse phase, much like Venus. In this case, even a late delivery of water by comets or asteroids may be unable to restore the planet's habitability.

I have also shown that planets that lose significant amounts of water retain tens, hundreds, or even thousands of bars of photolytically-produced O_2, which in certain cases could remain in the atmosphere. It is widely accepted that prebiotic chemistry happened in a reducing environment (Oparin, 1924; Haldane, 1929); moreover, life on Earth evolved in the absence of atmospheric oxygen for at least 1 Gyr (e.g., Schopf et al., 2007; Anbar et al., 2007). Early organisms relied on the free energy available in redox reactions involving a variety of hydrogen compounds; on an

O_2-rich planet, organisms would have to compete with the oxygen for this free energy. Abiogenesis in the presence of massive amounts of atmospheric oxygen could therefore be difficult, though these issues should be investigated further.

2.6.4 Other Remarks

I assumed that terrestrial planets form with abundant surface water. Planets that form dry and receive water at a late stage (such as by cometary impacts) are naturally more robust against water escape and oxygen buildup, since the runaway greenhouse will not last as long. The inner edge of the HZ for dry planets may also be significantly closer to the star (Abe et al., 2011). Moreover, I only considered planets that migrate via disk interactions. Planet-planet scattering processes, for instance, need not occur early; if a terrestrial planet scatters into the HZ after the star settles onto the MS, it will be safe from the early runaway period.

However, I only considered water loss and O_2 buildup during the runaway greenhouse state. Water loss may occur during a moist greenhouse state (Kopparapu et al., 2013) at the end of the star's contraction phase. Moreover, diffusion-limited escape of hydrogen can occur even on a planet that is not in a runaway state. Wordsworth and Pierrehumbert (2014) show that an N_2-poor Earth can lose up to 0.3 TO and produce 66 bars of O_2 due to diffusion-limited escape. A planet whose crust and mantle have been highly oxidized during the runaway regime may be unable to remove this oxygen, leading to an O_2-rich atmosphere even if all the oxygen produced during the runaway was absorbed into the magma.

I neglected the possibility of cold-trapping of the water at the end of the runaway phase. On Earth, water vapor is strongly inhibited from reaching the stratosphere by condensation in the troposphere. However, due to the high surface temperature on a runaway planet, the thermal structure follows a dry adiabat throughout most of the troposphere (e.g., Kasting, 1988), along which the vapor pressure is lower than the saturation vapor pressure and thus water cannot condense. Nevertheless, as a planet gets desiccated and H_2O becomes progressively less abundant relative to O_2, a cold trap could eventually be established. In principle, this could prevent planets from becoming completely desiccated, allowing a small fraction of the initial water content to remain after the end of the runaway.

I also ignored atmospheric loss due to flares and stellar winds, which is likely significant around

M dwarfs (Scalo et al., 2007; Kislyakova et al., 2013) and can lead to the loss of substantially more water than I calculate here. Flares could also remove some of the oxygen, lowering the amount that builds up in the atmosphere or at the surface, particularly around the lowest mass M dwarfs. On the other hand, stars less massive than about $0.1 M_\odot$ may be in a supersaturation regime early on, saturating at XUV fluxes one or even two orders of magnitude below those of higher mass M dwarfs (Cook et al., 2014). This could reduce water loss rates from planets around these lowest mass stars.

I further assume blow-off conditions for oxygen are not met. Tian (2009) showed that super-Earths with CO_2 atmospheres receiving up to 1000 $F_\oplus$ are stable to carbon escape; oxygen escape may be similarly inhibited. However, as Lammer et al. (2011b) point out, this may not be the case on Earth-mass planets that lack IR-cooling species. On such planets, significantly less oxygen could build up in the atmosphere.

I only presented figures where the runaway greenhouse occurs interior to the RG limit. I also considered the case of planets that are in a runaway only while they are interior to the RV limit (the inner edge of the empirical HZ). I find that total water loss amounts and O_2 pressures are similar, but all contours shift to the left (following the shift in the HZ boundary). This is due to the fact that these planets remain in a runaway for less time.

All plots shown in the chapter assume an XUV absorption efficiency $\epsilon_{\rm XUV} = 0.30$. This is likely an overestimate, since radiative cooling by oxygen can greatly reduce the amount of energy available to drive the escape. Recently, Bolmont et al. (2017) showed that the escape efficiency for water-rich terrestrial planets is likely in the range 0.01–0.1, which would decrease the total amount of water lost. Recall, however, that the rate of oxygen buildup is independent of the XUV flux (and by extension $\epsilon_{\rm XUV}$); nevertheless, the delayed desiccation of planets with lower $\epsilon_{\rm XUV}$ can lead to *more* oxygen buildup in the long run.

Finally, I neglected tidal heating and orbital evolution due to tides. For late M dwarfs, tidal evolution can be quite strong in the HZ (see Barnes et al., 2008, and the next chapter), meaning that planets on circular orbits in the HZ today may have started outside of the HZ on eccentric orbits. While this may help reduce the insolation on such planets early on, tidal heating could provide sufficient surface heat flux to trigger a runaway (Barnes et al., 2013). A tidal runaway could lead to a longer period of water loss and oxygen buildup.

2.7 Conclusions

I have shown that the extended pre-main sequence contraction phase of M dwarfs causes the habitable zones of these stars to move inwards by up to an order of magnitude in semi-major axis over the course of the first several hundred Myr. Since terrestrial M dwarf planets probably form within $\sim$ 10 Myr after the formation of the parent star, many planets currently in the HZs of M dwarfs were not in the HZs when they formed, in agreement with Lissauer (2007). If these planets formed with water, they may have experienced prolonged runaway greenhouses, lasting between $\sim$ 10 Myr for high mass M dwarfs and $\sim$ 1 Gyr for the lowest mass M dwarfs. Such prolonged runaways could lead to planetary evolution fundamentally different from Earth's, potentially compromising their habitability in the long run.

During a runaway greenhouse, photolysis of water vapor in the stratosphere followed by the hydrodynamic escape of the upper atmosphere can lead to the rapid loss of a planet's surface water. Because hydrogen escapes preferentially over oxygen, large quantities of O_2 may also build up. I have shown that planets currently in the HZs of M dwarfs may have lost up to several tens of terrestrial oceans (TO) of water during the early runaway phase, accumulating O_2 at a constant rate that is set by diffusion: about 5 bars/Myr for Earth-mass planets and 25 bars/Myr for super-Earths. At the end of the runaway phase, this leads to the buildup of hundreds to thousands of bars of O_2, which may or may not remain in the atmosphere. I considered two limiting cases regarding the oxygen: (i) highly efficient surface sinks, resulting in an atmospheric O_2 content that is always low; and (ii) inefficient O_2 sinks, leading to quick atmospheric buildup.

In the first case, I assume that water vapor is the dominant atmospheric species and that atmospheric escape occurs in the energy-limited regime. Both hydrogen and oxygen escape in proportions controlled by the stellar XUV flux. I find that for $M_\star \lesssim 0.3\mathrm{M}_\odot$, nearly all Earth-mass planets in the HZ lose at least 1 TO, though tens of TO are typically lost for $M_\star \lesssim 0.15\mathrm{M}_\odot$. For $0.3\mathrm{M}_\odot \lesssim M_\star \lesssim 0.6\mathrm{M}_\odot$, several TO are lost in the center of the HZ and close to the inner edge. The surfaces of these planets undergo extreme oxidation, absorbing the equivalent of hundreds to thousands of bars of O_2.

In the second case, I assume that O_2 is produced faster than surface sinks can remove it, resulting in an oxygen-rich atmosphere. I thus calculate escape rates according to the diffusion limit of hydrogen; in this case, the escape flux is insufficient to drag any of the oxygen off to

space. Because of the lower escape flux, water loss rates are slightly lower. However, several TO are still lost, particularly around low mass M dwarfs and close to the inner edge of the HZ. Oxygen amounts, on the other hand, are slightly *higher*; planets around all M dwarfs can develop atmospheres with hundreds to thousands of bars of O_2.

Perhaps counter-intuitively, I find that both the amount of water lost and the final oxygen pressure scale with planet mass; super-Earths tend to lose substantially more water and develop more massive O_2 atmospheres than Earth-mass planets. Despite their higher surface gravity, which reduces the total energy-limited escape rate, super-Earths lose water primarily via the escape of hydrogen, which causes a faster net loss of the oceans compared to the case in which both hydrogen and oxygen escape (which primarily occurs on Earth-mass planets). This is also the case for loss at the diffusion limit, since the escape flux scales with the surface gravity of the planet. I showed that as a result of this enhanced escape, some recently discovered super-Earths in the HZs of M dwarfs such as GJ 667Cc could have lost on the order of 10 TO and built up $\sim$ 2000 bars of O_2.

Given the variety of possible planetary compositions and processes capable of removing O_2 from the atmosphere, my two cases should roughly bracket the evolution of many exoplanets orbiting M dwarfs. However, the values reported here should be viewed as strict upper limits to the thermal escape rate of hydrogen and the rate of atmospheric oxygen buildup, as I assume a high XUV absorption efficiency ($\epsilon_{\rm XUV} = 0.3$) and neglect photochemical processes that could negatively feed back on the water photolysis and hydrogen escape rates (see, e.g., Zahnle and Catling, 2017). The production of ozone, for instance, could shield lower parts of the atmosphere from photolysis by UV radiation, bottlenecking the escape. Detailed photochemical modeling is therefore required to assess the actual amounts of O_2 these planets build up in their atmospheres.

Nevertheless, it is possible that many planets in the HZs of M dwarfs, in particular super-Earths, could retain enough atmospheric O_2 to be spectroscopically detectable by future missions such as the James Webb Space Telescope and the WSO-UV space observatory (Fossati et al., 2014). My work thus strengthens the results of Wordsworth and Pierrehumbert (2014), Tian et al. (2014), and Domagal-Goldman et al. (2014) that O_2 in a planetary atmosphere is not a reliable biosignature; in fact, such elevated quantities of atmospheric oxygen could potentially be an anti-biosignature. The habitability of many planets around M dwarfs should thus be questioned. Many such planets could be no more than "mirage" Earths.

3 Habitable Evaporated Cores

Portions of this chapter were originally published in collaboration with R. Barnes, E. Lopez, J. Fortney, B. Jackson, and V. S. Meadows in the journal Astrobiology (Luger et al., 2015, Astrobiology, ©2015 Mary Ann Liebert, Inc.), and are reproduced below with the permission of Mary Ann Liebert, Inc.

3.1 Introduction

As I discussed in the previous chapter, terrestrial planets in the habitable zones of M dwarfs face many challenges to their habitability. Due to the extended pre-main sequence phase of their host stars and their prolonged and vigorous magnetic activity, these planets may undergo extreme water loss following their formation. While the water loss rates we calculated in the previous chapter—in excess of 10 Earth oceans for planets around the lowest mass stars—are likely upper limits to the thermal escape rates, we neglected nonthermal processes such as stripping by flares or ion pickup (e.g.; Kislyakova et al., 2013, 2014) and shut off all escape once the planets entered the HZ. We thus concluded that water loss may pose the single biggest threat to the habitability of these planets.

However, even under the most pessimistic assumptions for atmospheric escape, planets in the HZs of M dwarfs need not be desiccated today. Late water delivery by comets, late outgassing from a water-rich mantle, or a large initial inventory of surface water could all guard a planet against complete desiccation. In this chapter, I introduce a new potential pathway to the formation of water-rich planets in the HZs of M dwarfs, drawing from recent studies that show migration may be quite common in planetary systems. Both disk-driven migration and planet-planet interactions can lead to substantial orbital changes, potentially bringing planets from outside the snow line (the region of the disk beyond which water and other volatiles condense into ices, facilitating the formation of massive planetary cores) to within the HZ (Hayashi et al., 1985; Ida and Lin, 2008a,b; Ogihara and Ida, 2009). Disk-driven migration relies on the exchange of angular momentum between a planet and the surrounding gaseous disk, which induces rapid inward migration for planets in the terrestrial mass range (Ward, 1997); this occurs on timescales shorter than 10–20 Myr for M dwarfs (Carpenter et al., 2006; Pascucci et al., 2009; Ribas et al., 2014). Planet-planet scattering, on the other hand, is by nature a stochastic process and may take place at any point

during a system's lifetime, although it is more likely in the first few tens of Myr (Ford and Rasio, 2008).

Given the abundance of ices beyond the snow line, planets that migrate into the HZ from the outer regions of the disk should have abundant water—potentially enough to safeguard against complete desiccation. Late planet-planet scattering events could also protect the planet from water loss during the most active phase of the star. But perhaps most importantly, planets that form beyond the snow line likely accrete large quantities of hydrogen and helium from the protoplanetary disk, forming extended gaseous envelopes. Because of the large scaleheight of this envelope, the hydrogen/helium should escape first under the XUV irradiation of the host star. Provided a planet forms with just enough H/He, the envelope could shield the surface during the most active phase of the star and eventually dissipate, revealing a water-rich world as the planet enters the HZ at the end of the pre-main sequence phase of the star. Investigating whether these so-called "super-Earths" and "mini-Neptunes" can lose their envelopes and form planets with rocky or liquid water surfaces is therefore critical to understanding the habitability of planets around low mass stars.

In this chapter I focus on planets with initial masses in the range $1\,\mathrm{M}_\oplus \leq M_p \leq 10\,\mathrm{M}_\oplus$ with up to 50% hydrogen/helium by mass that have migrated into the HZs of mid- to late M dwarf stars. I investigate whether it is possible for atmospheric escape processes to remove the thick H/He envelopes of these planets within the HZ, potentially turning them into volatile-rich Earths and super-Earths, which I refer to as "habitable evaporated cores" (HECs). I consider two atmospheric loss processes: XUV-driven escape, in which stellar X-ray/extreme ultraviolet (XUV) photons heat the atmosphere and drive a hydrodynamic wind away from the planet, and a simple model of Roche lobe overflow (RLO), in which the atmosphere is so extended that part of it lies exterior to the planet's Roche lobe; that gas is therefore no longer gravitationally bound to the planet. I further couple the effects of atmospheric mass loss to the thermal and tidal evolution of these planets. Planets cool as they age, undergoing changes of up to an order of magnitude in radius, which can greatly affect the mass loss rate. Tidal forces arising from the differential strength of gravity distort both the planet and the star away from sphericity, introducing torques that lead to the evolution of the orbital and spin parameters of both bodies.

While many studies have considered the separate effects of atmospheric escape (Erkaev et al., 2007; Murray-Clay et al., 2009; Tian, 2009; Owen and Jackson, 2012; Lammer et al., 2013), Roche

lobe overflow (Trilling et al., 1998; Gu et al., 2003), thermal evolution (Lopez et al., 2012), and tidal evolution (Jackson et al., 2008a; Ferraz-Mello et al., 2008; Correia and Laskar, 2011) on exoplanets, none has considered the coupling of these effects in the HZ. For some systems, in particular those that may harbor HECs, modeling the coupling of these processes is essential to accurately determine the evolution, since several feedbacks can ensue. Tidal forces in the HZ typically act to decrease a planet's semi-major axis, leading to higher stellar fluxes and faster mass loss. The mass loss, in turn, affects the rate of tidal evolution primarily via the changing planet radius, which is also governed by the cooling rate of the envelope. Changes to the star's radius and luminosity lead to further couplings that need to be treated with care.

Jackson et al. (2010) considered the effect of the coupling between mass loss and tidal evolution on the hot super-Earth CoRoT-7 b, finding that the two effects are strongly linked and must be considered simultaneously to obtain an accurate understanding of the planet's evolution. However, an analogous study has not been performed in the HZ, in great part because both tidal effects and atmospheric mass loss are generally orders of magnitude weaker at such distances from the star. This is not necessarily true around M dwarfs, for two reasons: (a) their low luminosities result in a HZ that is much closer in, exposing planets to strong tidal effects and possible RLO; and (b) M dwarfs are extremely active early on, so that the XUV flux in the HZ can be orders of magnitude higher than that around a solar-type star (see, for instance, Scalo et al., 2007).

In this chapter I present the results of a model that couples tides, thermal evolution, and atmospheric mass loss in the habitable zone, showing that for certain systems the coupling is key in determining the long-term evolution of the planet. I further demonstrate that it is possible to turn gas-rich super-Earths into HECs within the habitable zone, providing an important pathway to the formation of potentially habitable, volatile-rich planets around M dwarfs.

In §3.2 I describe the model I use in detail, followed by the results in §3.3. I then discuss the main findings and the corresponding caveats in §3.4, followed by a summary in §3.5.

3.2 Model Description

The model I developed evolves planet-star systems forward in time in order to determine whether HECs can form from hydrogen-rich super-Earths that have migrated into the HZs of M dwarfs. I perform my calculations on a grid of varying planetary, orbital, and stellar properties in order to

Table 1: Parameters used in my code to model the evaporation of super-Earths and mini-Neptunes in the HZ of M dwarfs. The second column indicates the ranges assumed and the third column indicates the values adopted in the default cases.

Parameter	Range	Default	Notes
$M_\star (M_\odot)$	$0.08 - 0.4$	-	Late-mid MD
$M_p (M_\oplus)$	$1 - 10$	-	-
$R_{\rm XUV} (R_p)$	$1.0 - 1.2$	1.2	See §1.3.1.2
a	IHZ - OHZ	-	See §3.2.1
e	0.0 - 0.95	-	-
$P_{0,\star}$ (days)	$1.0 - 100$	30.0	Initial rot. per.
f_H	$10^{-6} - 0.5$	-	H mass fraction
$\epsilon_{\rm XUV}$	$0.1 - 0.4$	0.3	-
$\xi_{\rm min}$	$1 + 10^{-5} - 3$	3	See §3.2.3
Atmos. esc.	R/R-Lim / E-Lim	-	See §3.2.3
Tidal model	CPL/CTL	CTL	-
$Q_\star$	$10^5 - 10^6$	10^5	CPL only
Q_p	$10^1 - 10^5$	10^4	CPL only
$\tau_\star$ (s)	$10^{-2} - 10^{-1}$	10^{-1}	CTL only
τ_p (s)	$10^{-3} - 10^3$	10^{-1}	CTL only
β	$0.7 - 1.23$	1.23	See Eq. (1)
$t_{\rm sat}$ (Gyr)	$0.1 - 1.0$	1.0	XUV sat. time
t_0 (Myr)	$10.0 - 100.0$	10.0	Integration start
$t_{\rm stop}$ (Gyr)	$0.01 - 5.0$	5.0	Integration end

determine the types of systems that may harbor HECs. The complete list is provided in Table 1, where I indicate the ranges of values I consider as well as the default values adopted in the plots in §3.3 (unless otherwise indicated).

Integrations are performed from $t = t_0$ (the time at which the planet is assumed to have migrated into the HZ) to $t = t_{\rm stop}$ (the current age of the system) using the adaptive timestepping method described in Appendix E of Barnes et al. (2013).

3.2.1 Stellar Model

As in the previous chapter, I use the evolutionary tracks of Baraffe et al. (1998) for solar metallicity to calculate $L_{\rm bol}$ and $T_{\rm eff}$ as a function of time. I then use Equation (1) to calculate $L_{\rm XUV}$, given $(L_{\rm XUV}/L_{\rm bol})_{\rm sat} = 10^{-3}$ and values of $t_{\rm sat}$ and β given in Table 1.

Using $L_{\rm bol}$ and $T_{\rm eff}$, I calculate the location of the HZ from the equations given in Kopparapu et al. (2013), adding the eccentricity correction (Equation 2). Given the uncertainty in the actual HZ boundaries and their dependence on a host of properties of a planet's climate, I choose my

inner edge (IHZ) to be the average of the Recent Venus and the Runaway Greenhouse limits and my outer edge (OHZ) to be the average of the Maximum Greenhouse and the Early Mars limits. Throughout this chapter I will also refer to the center of the HZ (CHZ), which I take to be the average of the IHZ and OHZ. Since I am concerned with the formation of ultimately habitable planets, I take the locations of the IHZ, CHZ, and OHZ to be their values at 1 Gyr, at which point the stellar luminosity becomes roughly constant.

3.2.2 Planet Radius Model

To determine the planetary radius R_p as a function of the core mass M_c, the envelope mass fraction $f_H \equiv M_e/M_p$, and the planet age, I use the planet structure model described in Lopez et al. (2012) and Lopez and Fortney (2014), which is an extension of the model of Fortney et al. (2007) to low-mass low-density (LMLD) planets. These models perform full thermal evolution calculations of the interior as a function of time. In my runs, the core is taken to be Earth-like, with a mixture of 2/3 silicate rock and 1/3 iron, and the envelope is modeled as a H/He adiabat. A grid of values of R_p is then computed in the range $1\,\mathrm{M}_\oplus \leq M_c \leq 10\,\mathrm{M}_\oplus$, $10^{-6} \leq f_H \leq 0.5$, and 10^7 years $\leq t \leq 10^{10}$ years. For values between grid points, I perform a simple trilinear interpolation. For gas-rich planets, R_p is the 20 mbar radius; for gas-free planets, it corresponds to the surface radius. The evolution of R_p with age due solely to thermal contraction is plotted in Figure 17 for a few different core masses and values of f_H.

The models of Fortney et al. (2007) and Lopez et al. (2012) are in general agreement with those of Mordasini et al. (2012b,a) and, by extension, Rogers et al. (2011). Mordasini et al. (2012b) presented a validation of their model against that of Fortney et al. (2007), showing that for planets spanning 0.1 to 10 Jupiter masses, the two models predict the same radius to within a few percent. At the lower masses relevant to this study, the two models are also in agreement. To demonstrate this, in Figure 17 I shade the regions corresponding to the spread in radii at a given mass and age in Figure 9 of Mordasini et al. (2012a). Since those authors used a coupled formation/evolution code, at low planet mass the maximum envelope mass fraction f_H is small; for a total mass of $4\mathrm{M}_\oplus$, Mordasini et al. (2012a) find that all planets have $f_H < 0.2$. At $2\mathrm{M}_\oplus$, most planets have $f_H \lesssim 0.1$. It is clear from Figure 17 that at these values of f_H, the two models predict very similar radius evolution. Note that Mordasini et al. (2012a) did not consider planets

less massive than $2M_\oplus$.

The maximum envelope fraction merits further discussion. Since I do not model the formation of these planets, I do not place a priori constraints on the value of f_H at a given mass; instead, I allow it to vary in the range $10^{-6} \leq f_H \leq 0.5$ for all planet masses. At masses $\lesssim 5M_\oplus$, planets accumulate gas slowly and are typically unable to accrete more than $\sim$ 10-20% of their mass in H/He (Rogers et al., 2011; Bodenheimer and Lissauer, 2014); values of $f_H \approx 0.5$ are likely unphysical. However, as I argue in §3.4.1, the longer disk lifetimes around M dwarfs (Carpenter et al., 2006; Pascucci et al., 2009) allow more time for gas accretion, potentially increasing the maximum value of f_H. Nevertheless, and more importantly, if a planet with $f_H = 0.5$ loses its entire envelope via atmospheric escape, *any planet with the same core mass and* $f_H < 0.5$ *will also lose its envelope.* Below, where I present integrations with $f_H = 0.5$, my results are therefore conservative, as planets with $f_H \ll 0.5$ will in general evaporate more quickly.

While my treatment of the radius evolution is an improvement upon past tidal-atmospheric coupling papers (Jackson et al., 2010, for instance, calculate R_p for super-Earths by assuming a constant density as mass is lost), there are still issues with my approach: (1) I do not account for inflation of the radius due to high insolation. Instead, I calculate the radii from grids corresponding to a planet receiving the same flux as Earth. While at late times this is justified, since planets in the HZ by definition receive fluxes similar to Earth, at early times this is probably a poor approximation; recall that planets in the HZ around low mass M dwarfs are exposed to fluxes up to two orders of magnitude higher during the host star's pre-main sequence phase. The primary effect of a higher insolation is to act as a blanket, delaying the planet's cooling and causing it to maintain an inflated radius for longer. This will result in mass loss rates higher than what I calculate here. (2) Since I am determining the radii from pre-computed grids, I also do not model the effect of tidal dissipation on the thermal evolution of the planet. Planets undergoing fast tidal evolution can dissipate large amounts of energy in their interiors, which should lead to significant heating and inflation of their radii. (3) The radius is also likely to depend on the mass loss rate. Setting R_p equal to the tabulated value for a given mass, age, and composition is valid only as long as the timescale on which the planet is able to cool is significantly shorter than the mass loss timescale. Otherwise, the radius will not have enough time to adjust to the rapid loss of mass and the planet will remain somewhat inflated, leading to a regime of runaway mass loss (Lopez et al.,

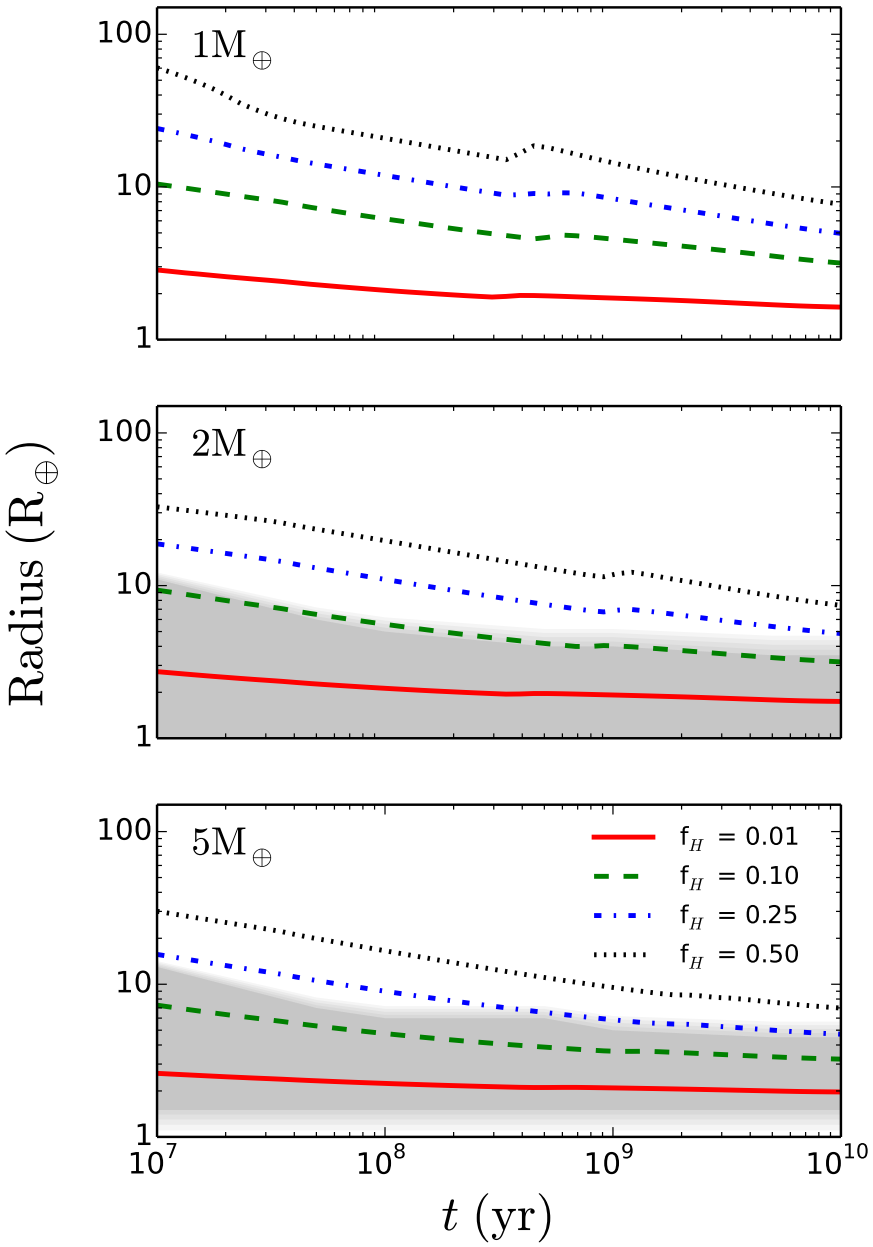

Figure 17: Evolution of the radius as a function of time due to thermal contraction of the envelope, in the absence of tidal effects and atmospheric mass loss. From top to bottom, the plots correspond to planets with initial total masses (core + envelope) of 1, 2, and 5 $M_\oplus$. Line styles correspond to different initial hydrogen mass fractions: 1% (red, solid), 10% (green, dashed), 25% (blue, dot-dashed), and 50% (black, dotted). For comparison, the grey shaded regions in the bottom two plots are the spread in radii calculated by Mordasini et al. (2012a) for $f_H \lesssim 0.20$. See text for a discussion.

2012). While the planets considered here are probably not in the runaway regime (Lopez et al. (2012) found that runaway mass loss occurred only for H/He mass fractions $\gtrsim 90\%$), I might still be significantly underestimating the radii during the early active phase of the parent star.

All points outlined above lead to an *underestimate* of the radius at a given time. Since the mass loss rate is proportional to R_p^3 (Equation 5) or $R_p^{3/2}$ (Equation 13), calculating the radius in this fashion leads to a *lower bound* on the amount of mass lost and on the strength of the coupling to tidal effects. Because my present goal is to determine whether it is possible to form habitable evaporated cores via this mechanism, this conservative approach is sufficient. Future work will incorporate a self-consistent thermal structure model to better address the radius evolution.

3.2.3 Atmospheric Escape Model

I assume that the escape of H/He from the planet atmosphere is hydrodynamic (blow-off) at all times, which is valid at the XUV fluxes I consider here (see Erkaev et al., 2013, and Figure 2). I run two separate sets of integrations: one in which I assume the flow is energy-limited (Equation 5) for all values of $F_{\rm XUV}$, and one in which I switch from energy-limited to radiation/recombination-limited (Equation 13) above the critical value of the flux (see §1.3.1.4). For planets whose orbits are eccentric enough that they switch between the two regimes over the course of one orbit, I make use of the expressions derived in §3.2.4.3. These two sets of integrations should roughly bracket the true mass loss rate.

I vary $\epsilon_{\rm XUV}$ and $R_{\rm XUV}$ in the ranges given in Table 1. I choose $\epsilon_{\rm XUV} = 0.30$ as my default case. While this is consistent with values cited in the literature (see §1.3.1.2), it could be an overestimate. I discuss the implications of this choice in §3.4.3.

Given the large planetary radii at early times, many of the planets I model here are not stable against Roche lobe overflow in the HZ. During RLO, the stellar gravity causes the upper layers of the atmosphere to become unbound from the planet; this occurs when $R_{\rm p} > R_{\rm Roche}$, where $R_{\rm Roche}$ is given by Equation (8). For a planet that forms and evolves *in situ*, RLO never occurs, since any gas that would be lifted from the planet in this fashion would have never accreted in the first place. However, an inflated gaseous planet that forms at a large distance from the star may initially be stable against overflow and enter RLO as it migrates inwards (since $R_{\rm Roche} \propto a$). This is particularly the case for planets in the HZs of M dwarfs, since a and consequently $R_{\rm Roche}$ are

small.

Ideally, the tidally-enhanced mass loss rate (Equation 5) should capture this process, but instead it predicts an infinite mass loss rate as $R_{\rm XUV} \to R_{\rm Roche}$ (or as $\xi \to 1$) and unphysically changes sign for $\xi < 1$. This is due to the fact that the energy-limited model implicitly assumes that the bulk of the atmosphere is located at $R_{\rm XUV}$ (the single-layer assumption). Realistically, one would expect the planet to quickly lose any mass above the Roche lobe and then return to the stable hydrodynamic escape regime. However, upon loss of the material above $R_{\rm Roche}$, the portion of the envelope below the new XUV absorption radius $R'_{\rm XUV}$ will not be in hydrostatic equilibrium; instead, an outward flow will attempt to redistribute mass to the evacuated region above, leading to further overflow.

Several models exist that allow one to calculate the mass loss rate due to RLO (e.g., Ritter, 1988; Trilling et al., 1998; Gu et al., 2003; Sepinsky et al., 2007). These often involve calculating the angular momentum exchange between the outflowing gas and the planet, which can lead to its outward migration, given by

$$\frac{1}{a}\frac{da}{dt} = -\frac{2}{M_p}\frac{dM_p}{dt}, \tag{72}$$

for a planet on a circular orbit (Gu et al., 2003; Chang et al., 2010). This leads to a corresponding increase in $R_{\rm Roche}$ until it reaches $R_{\rm XUV}$ and the overflow is halted. By differentiating the stability criterion $R_{\rm XUV}(M_p) = R_{\rm Roche}(M_p)$, one may then obtain an approximate expression for dM_p/dt in terms of the density profile $dM(< R)/dR$ of the envelope.

However, for super-Earths and mini-Neptunes that migrate into the HZ early on, RLO should occur *during* the initial migration process, which I do not model here. Instead, I begin my calculations by assuming that the planets are stable to RLO in the HZ. If a planet's radius initially exceeds the Roche lobe radius, I set its envelope mass equal to the *maximum* envelope mass for which it can be stable at its current orbit; the difference between the two envelope masses is the amount of H/He it must have lost prior to its arrival in the HZ. It is important to note that these planets will initially have $R_{\rm XUV} = R_{\rm Roche}$, which as I mentioned above, leads to an infinite mass loss rate in (5). An accurate determination of $\dot{M}_p$ in this case therefore requires hydrodynamic simulations. However, the mass loss rate can be approximated by imposing a minimum value $\xi_{\rm min}$ in (5). For $\xi < \xi_{\rm min}$, I set the mass loss rate equal to $\dot{M}_p(\xi = \xi_{\rm min})$. This is equivalent to imposing

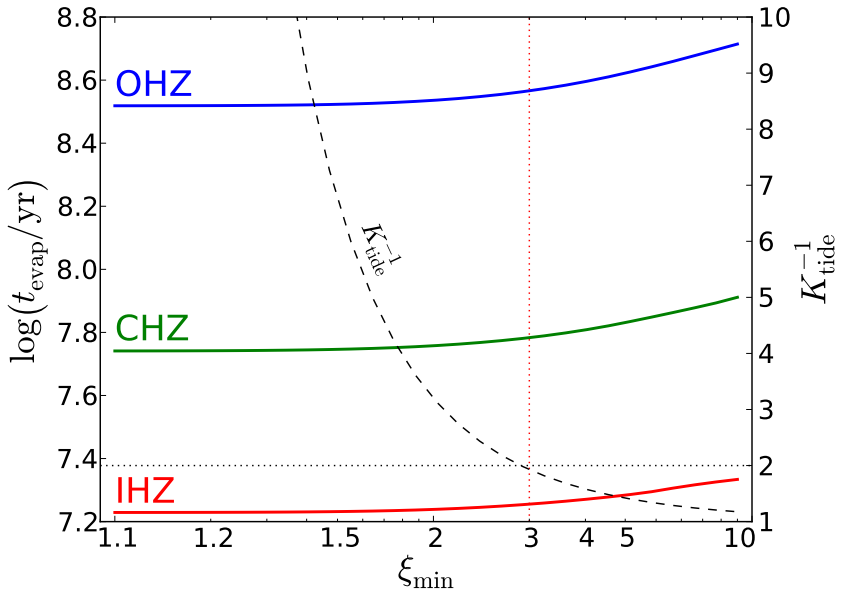

Figure 18: Complete evaporation time $t_{\rm evap}$ as a function of the cutoff value $\xi_{\rm min}$ for a $2\,{\rm M}_\oplus$ planet with $f_{H_0} = 0.5$ on a circular orbit around a $0.08\,{\rm M}_\odot$ star. The red, green, and blue lines correspond to planets in the IHZ, CHZ, and OHZ, respectively. Also plotted is the mass loss enhancement factor $1/K_{\rm tide}$ (black dashed line), which approaches infinity as $\xi \to 1$. Note that for $\xi_{\rm min} \lesssim 3$, the evaporation time is relatively insensitive to the exact cutoff value, despite the fact that $1/K_{\rm tide}$ blows up. I therefore choose $\xi_{\rm min} = 3$ as the default cutoff, corresponding to a maximum enhancement factor $1/K_{\rm tide} \approx 2$.

a maximum mass loss enhancement factor $1/K_{\rm tide}$, preventing the mass loss rate from reaching unphysical values as $R_{\rm XUV} \to R_{\rm Roche}$.

In Figure 18 I show how the time $t_{\rm evap}$ at which complete evaporation takes place scales with $\xi_{\rm min}$ for a typical gas-rich super-Earth/mini-Neptune in the IHZ (red), CHZ (green), and OHZ (blue). Also plotted is the mass loss enhancement factor $1/K_{\rm tide}$ (Equation 6, black dashed line) as a function of $\xi = \xi_{\rm min}$. Interestingly, despite the fact that the instantaneous mass loss rate ($\dot{M}_p \propto 1/K_{\rm tide}$) approaches infinity as $\xi \to 1$, the evaporation time is relatively constant for $\xi_{\rm min} \lesssim 3$. This is because for very large $\dot{M}_p$, the planet loses sufficient mass in an amount of time Δt to decrease R_p substantially and terminate the overflow. In other words, cases in which $\xi \approx 1$ are very unstable, and as mass is lost ξ will quickly increase beyond ~ 3. Both the net amount of mass lost and the evaporation time are therefore insensitive to the particular value of $\xi_{\rm min}$, provided it is less than about 3. I therefore choose $\xi_{\rm min} = 3$ as the default value for my runs, noting that this corresponds to a maximum mass loss enhancement of $1/K_{\rm tide} \approx 2$.

3.2.4 The Effect of Eccentricity on Atmospheric Escape

Most of the formalism that has been developed to analytically treat hydrodynamic escape only considers circular orbits. For planets on sufficiently eccentric orbits, neither the stellar flux nor the tidal effects may be treated as constant over the course of an orbit. Below I derive expressions analogous to Equation (5) that apply to planets on eccentric orbits.

3.2.4.1 The Quasi-Static Approximation

There are two separate effects that enhance the mass loss for planets on eccentric orbits. Most papers account for the first effect, which is the increase of the orbit-averaged stellar flux by a factor of $1/\sqrt{1-e^2}$ (see, for instance, Kopparapu et al., 2013). However, for $e \lesssim 0.3$, this effect is quite small. The second, more important effect is that the Roche lobe radius is no longer constant over the course of an orbit, and (8) is not valid. Instead, we must replace a with the instantaneous planet-star separation $r(t)$:

$$R_{\rm Roche}(t) = \left(\frac{M_p}{3M_\star}\right)^{1/3} r(t). \tag{73}$$

One might wonder whether this replacement is valid. Specifically, if $R_{\rm Roche}(t)$ changes faster than the atmosphere is able to respond to the changes in the gravitational potential, we would expect that the time dependence of the mass loss rate would be a complicated function of the tides generated in the atmosphere. On the other hand, if the orbital period is very large compared to the dynamical timescale of the planet, the atmosphere will have sufficient time to assume the equilibrium shape dictated by the new potential. This limit is known as the *quasi-static approximation* (Sepinsky et al., 2007). To show that this approximation is valid, I begin by introducing the parameter *alpha* (Sepinsky et al., 2007):

$$\alpha(e,f) = \frac{(1+e)^{1/2}}{(1-e)^{3/2}} \times |1-f|, \tag{74}$$

where

$$f = \frac{P_{\rm orb}}{P_{\rm rot}} \frac{(1-e)^{3/2}}{(1+e)^{1/2}} \tag{75}$$

is the ratio of the rotational angular velocity to the orbital angular velocity at periastron. Sepinsky et al. (2007) show that provided the condition

$$\frac{P_{\rm orb}}{\tau_{\rm dyn}} \gg \alpha(e, f) \tag{76}$$

holds, the system may be treated as quasi-static. The orbital period is given by Kepler's law,

$$P_{\rm orb} = 2\pi\sqrt{\frac{a^3}{GM_\star}}, \tag{77}$$

while $\tau_{\rm dyn}$, the dynamical timescale of the planet, is

$$\tau_{\rm dyn} \sim \frac{1}{\sqrt{GM_p/R_p^3}}. \tag{78}$$

The ratio of the two will be smallest for close-in, low-mass planets with large radii. The minimum value in my runs occurs for a super-inflated $2\,{\rm M}_\oplus$, $30\,{\rm R}_\oplus$ planet in the IHZ of a $0.08\,{\rm M}_\odot$ M dwarf, for which $P_{\rm orb}/\tau_{\rm dyn} \approx 2.4$.

We must now compare this to $\alpha(e, f)$. The equilibrium rotational period for a synchronously-rotating planet is obtained from Equations (41) and (55) in the CPL and CTL models, respectively, so that we may write $\alpha(e)$ as

$$\alpha(e, f) = \frac{(1+e)^{1/2}}{(1-e)^{3/2}} \times \left|1 - \frac{(1-e)^{3/2}}{(1+e)^{1/2}} \frac{\omega_{2,\rm eq}}{n}\right|. \tag{79}$$

This equation is plotted in Figure 19. Note that α only begins to approach $P_{\rm orb}/\tau_{\rm dyn}$ for $e \gtrsim 0.6$ in the CPL model and $e \gtrsim 0.4$ in the CTL model, and only for the most inflated planets in the IHZ. I therefore urge caution in interpreting results above $e \gtrsim 0.5$, where the quasi-static approximation may not hold for some planets.

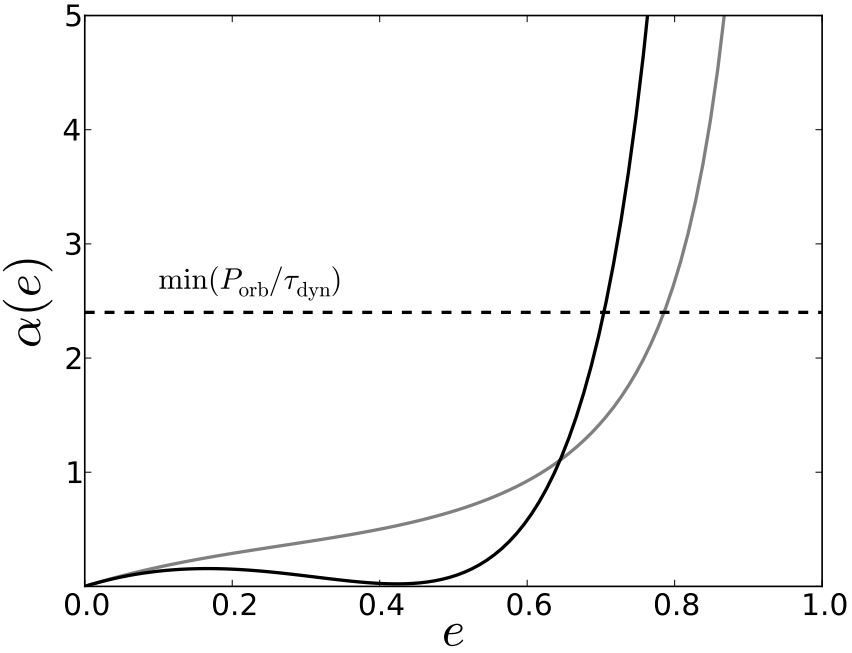

Figure 19: The Sepinsky et al. (2007) parameter α as a function of e for a synchronously-rotating planet in the CPL model (black) and in the CTL model (gray). The dashed line corresponds to the minimum value of the ratio $P_{\rm orb}/\tau_{\rm dyn}$ across all my runs. Note that for $e \lesssim 0.5$, the quasi-static approximation is probably valid.

3.2.4.2 The Mass Loss Enhancement Factor

Since $R_{\rm Roche} = R_{\rm Roche}(t)$, ξ, $K_{\rm tide}$, and dM/dt (Equations 5, 6 and 7) are now also functions of t, varying significantly over a single orbit. Equation (5) must thus be modified slightly:

$$\dot{M}(t) = \frac{R_{\rm XUV}^3 \epsilon_{\rm XUV} L_{\rm XUV}}{4GM_p} \left[r(t)^2 \left(1 - \frac{3}{2\xi(t)} + \frac{1}{2\xi(t)^3} \right) \right]^{-1}, \tag{80}$$

where $r(t)$ is the instantaneous separation between the centers of mass of the star and the planet and where I have plugged-in for $F_{\rm XUV}$ in terms of $L_{\rm XUV}$ and made use of Equation (6). The parameter ξ must also be modified, as the value of the Roche radius will also change as the

planet's distance from the star changes during one orbit:

$$\begin{aligned}\xi(t) &\equiv \frac{R_{\rm Roche}}{R_{\rm XUV}} \\ &= \left[\left(\frac{M_p}{3M_\star}\right)^{1/3}\frac{1}{R_{\rm XUV}}\right] r(t) = \frac{r(t)}{a}\xi. \end{aligned} \tag{81}$$

To avoid confusion, henceforth ξ will denote the original expression for circular orbits (Equations 7 and 8), while $\xi(t)$ is the time-dependent parameter given by the expression above. The time-average of $\dot{M}$ over one orbit is

$$\langle \dot{M} \rangle_t = \frac{1}{2\pi/n}\int_0^{2\pi/n} \dot{M}(t)dt \tag{82}$$

where n is the mean motion. Now, we know that the relationship between n and the eccentric anomaly E is

$$nt = E - e\sin E \tag{83}$$

and that

$$r(E) = a(1 - e\cos E), \tag{84}$$

so

$$\frac{dt}{dE} = \frac{1}{n}(1 - e\cos E) = \frac{r(E)}{an}. \tag{85}$$

Substituting into (82), we have

$$\begin{aligned}\langle \dot{M} \rangle_t &= \frac{1}{2\pi/n}\int_0^{2\pi} \dot{M}(E)\frac{dt}{dE}dE \\ &= \frac{1}{2\pi a}\int_0^{2\pi} \dot{M}(E)r(E)dE. \end{aligned} \tag{86}$$

Now, introducing the mass loss rate for $e = 0$ and $K_{\rm tide} = 1$ from Equation (5),

$$\dot{M}_0 \equiv \frac{R_{\rm XUV}^3 \epsilon_{\rm XUV} L_{\rm XUV}}{4GM_p a^2}, \tag{87}$$

it follows from Equations (80) and (81) that

$$\dot{M}(E) = \dot{M}_0 a^2 \left[r(E)^2 \left(1 - \frac{3}{2\xi} \left(\frac{a}{r(E)} \right) + \frac{1}{2\xi^3} \left(\frac{a}{r(E)} \right)^3 \right) \right]^{-1}. \tag{88}$$

Plugging this into Equation (86), and using Equation (84),

$$\begin{aligned} \langle \dot{M} \rangle_t &= \frac{\dot{M}_0}{2\pi} a \int_0^{2\pi} \left(r(E) - \frac{3}{2\xi} a + \frac{1}{2} \left(\frac{a}{\xi} \right)^3 r(E)^{-2} \right)^{-1} dE \\ &= \frac{\dot{M}_0}{2\pi} \int_0^{2\pi} \left[(1 - e\cos E) - \frac{3}{2\xi} + \frac{1}{2\xi^3 (1 - e\cos E)^2} \right]^{-1} dE \\ &= \frac{\dot{M}_0}{K_{\rm ecc}}, \end{aligned} \tag{89}$$

where I define the eccentric mass loss enhancement factor

$$1/K_{\rm ecc} \equiv \frac{1}{2\pi} \int_0^{2\pi} \left[(1 - e\cos E) - \frac{3}{2\xi} + \frac{1}{2\xi^3 (1 - e\cos E)^2} \right]^{-1} dE. \tag{90}$$

Note that for $e = 0$, $K_{\rm ecc}$ reduces to the circular version of $K_{\rm tide}$, Equation (6). Unfortunately, the integral in Equation (89) is not analytic. However, if the last term in the integral is small compared to the first two, an analytic solution is possible. In particular, if we write

$$\frac{1}{2\xi^3 (1 - e\cos E)^2} = \eta \left[(1 - e\cos E) - \frac{3}{2\xi} \right], \tag{91}$$

then we may neglect the last term provided $\eta \ll 1$. Since the term is largest at pericenter ($E = 0$), we may instead require that

$$0 < \frac{1}{2\xi^3 (1 - e)^2 (1 - \frac{3}{2\xi} - e)} \ll 1. \tag{92}$$

If this holds, Equation (89) simplifies to

$$\begin{aligned}\langle \dot{M} \rangle_t &\approx \frac{\dot{M}_0}{2\pi} \int_0^{2\pi} \left(1 - \frac{3}{2\xi} - e\cos E\right)^{-1} dE \\ &\approx \frac{\dot{M}_0}{2\pi(1-\frac{3}{2\xi})} \int_0^{2\pi} \left[1 - \left(\frac{e}{1-\frac{3}{2\xi}}\right)\cos E\right]^{-1} dE \\ &\approx \frac{\dot{M}_0}{1-\frac{3}{2\xi}} \frac{1}{\sqrt{1-\left(\frac{e}{1-\frac{3}{2\xi}}\right)^2}} \\ &\approx \frac{\dot{M}_0}{\left(1-\frac{3}{\xi}-\frac{9}{4\xi^2}-e^2\right)^{1/2}}. \end{aligned} \tag{93}$$

Thus, provided condition (92) holds (typically for $\xi \gtrsim 10$), we may write

$$K_{\rm ecc} \approx \sqrt{1-\frac{3}{\xi}-\frac{9}{4\xi^2}-e^2}. \tag{94}$$

Note that for $e = 0$ and keeping only first-order terms in ξ, the above expression agrees with that of Erkaev et al. (2007) for $\xi \gg 1$.

In Figure 20 I plot the loss rate enhancement factor of Erkaev et al. (2013), $1/K_{\rm tide}$ (blue lines), along with the eccentric version, $1/K_{\rm ecc}$ (exact[8] expression in black, approximate expression in red). Recall that there are two distinct effects contributing to the extra mass loss for eccentric orbits: the $1/\sqrt{1-e^2}$ flux enhancement and the smaller Roche lobe distance during part of the orbit. In order to better distinguish between these two, I display both $1/K_{\rm tide}$ (dotted blue lines) and $1/(K_{\rm tide}\sqrt{1-e^2})$ (dashed blue lines). Thus, any difference between $1/K_{\rm tide}$ and $1/K_{\rm ecc}$ in the figure is due *solely* to the Roche lobe effect.

Let us first compare the curves corresponding to the exact expressions (blue and black). For $e = 0$, $K_{\rm ecc} = K_{\rm tide}$, as expected. As the eccentricity increases to 0.25, the flux enhancement effect remains small (i.e., the dashed and dotted curves are similar), while the Roche lobe effect begins to become significant (the dashed black curve exceeds the blue dashed curve), particularly for low values of ξ. For $e \gtrsim 0.5$, the effect becomes even more important, leading to an enhancement of a factor of several for $\xi \lesssim 10$.

[8]While I refer to Equation (90) as the "exact" expression for the (inverse of the) mass loss enhancement factor, it is important to remember that it is the eccentric version of the third-order expression derived by Erkaev et al. (2007) and is, in this sense, still an approximation to the true enhancement.

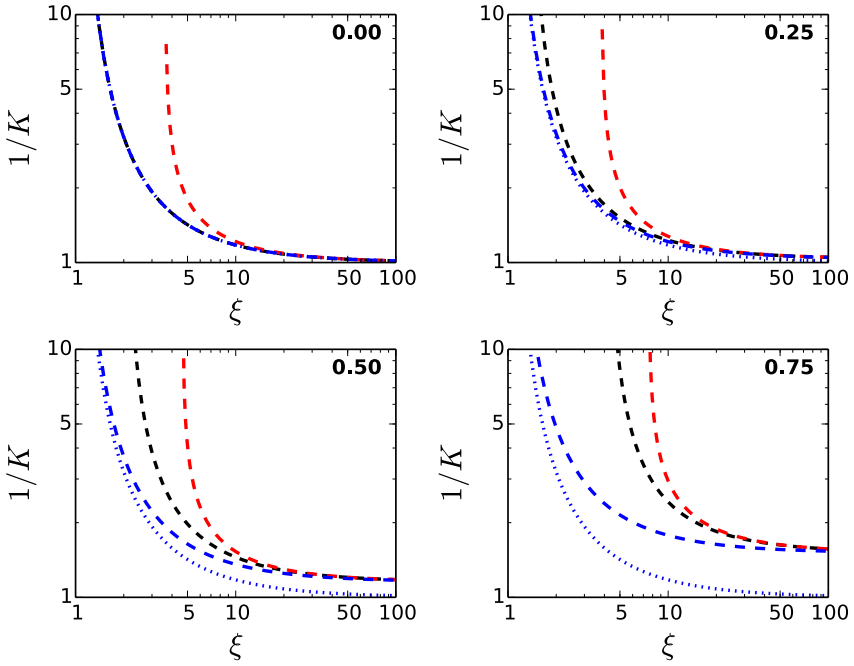

Figure 20: Loss rate enhancement factor $1/K$ as a function of the normalized Roche lobe distance $\xi = R_{\rm roche}/R_p$ for $e = 0,\ 0.25,\ 0.5$ and 0.75. The Erkaev et al. (2007) model ($K_{\rm tide}$) is plotted in blue, with the flux enhancement factor $1/\sqrt{1-e^2}$ (dashed) and without it (dotted). The exact expression $K_{\rm ecc}$ derived in this chapter (Equation 89) is plotted in black and the analytic approximation (Equation 93) is plotted in red. For low values of ξ, condition (92) is not satisfied and the approximation diverges significantly. However, at high eccentricity, as ξ increases, the analytic approximation converges to the exact value of $K_{\rm ecc}$ sooner than $K_{\rm tide}$. Finally, note that for sufficiently large ξ, all curves accounting for the flux enhancement converge to the same value.

A particularly important effect is that high eccentricities will cause the planet to undergo Roche lobe overflow at larger values of ξ. In the circular Erkaev et al. (2007) model, Roche lobe overflow occurs when $\xi = 1$ by definition. However, it is straightforward to show that for

$$\xi = \xi_{\rm crit} = \frac{1}{1-e}, \tag{95}$$

the expression in the integral of Equation (90) diverges, resulting in an infinite mass loss rate. This occurs because at pericenter, where $r = a(1-e)$, $\xi(r) = 1$. In other words, for $\xi \leq \xi_{\rm crit}$, the planet will overflow its Roche lobe during at least part of its orbit, leading to rapid mass loss.

It is important to note that the approximate expression (dashed red curve) converges only for $\xi \gtrsim 10$. For smaller values of ξ, it overestimates the mass loss enhancement significantly and

should therefore not be used. However, for high eccentricities (see the last panel, for instance), it converges to $1/K_{\rm ecc}$ quicker than $1/K_{\rm tide}$. Provided condition (92) is satisfied, the approximate analytic expression (Equation 93) does a good job at modeling the actual mass loss.

3.2.4.3 Orbits Crossing the Critical Radius

As I described in §1.3.1.4, the escape regime for an EUV-dominated flow may switch from energy-limited to radiation/recombination-limited above a certain critical value of the flux, $F_{\rm crit}$, given by

$$F_{\rm crit} = \left(\frac{B}{A}\right)^2, \tag{96}$$

where

$$A \equiv \frac{\pi \epsilon_{\rm XUV} R_{\rm XUV}^3}{G M_p K_{\rm tide}} \tag{97}$$

and

$$B \equiv 7.11 \times 10^7 \ {\rm g}^{\frac{1}{2}} {\rm s}^{\frac{1}{2}} \left(\frac{R_{\rm p}}{R_\oplus}\right)^{\frac{3}{2}} \tag{98}$$

are the coefficients multiplying the flux in the energy-limited and radiation/ recombination-limited equations, respectively (Equations 5 and 13). At fixed XUV luminosity, this corresponds to a certain critical orbital radius,

$$r_{\rm crit} = \sqrt{\frac{L_{\rm XUV}}{4\pi F_{\rm crit}}}. \tag{99}$$

Planets on low-eccentricity orbits that do not cross $r_{\rm crit}$ are therefore safely within either the energy-limited ($a \gg r_{\rm crit}$) or radiation/recombination-limited ($a \ll r_{\rm crit}$) regime. In this case, the orbit-averaged mass loss rate is determined simply by replacing F in Equations (5) and (13) by its orbit-averaged value, $\langle F \rangle$, given by Equation (2).

However, when an orbit is sufficiently eccentric that the atmospheric escape regime switches from energy-limited to radiation/recombination-limited over the course of a single orbit, the method outlined above is no longer rigorously correct. We must instead integrate the two mass

loss rate expressions over the portions of the orbit where they apply. Fortunately, as I demonstrate below, these integrals are analytic. This allows us to calculate the time-averaged value of the mass loss rate, $\langle \dot{M} \rangle_t$, much as I did in §3.2.4:

$$\langle \dot{M} \rangle_t = \frac{A}{P} \int_{\mathrm{EL}} F(t) dt + \frac{B}{P} \int_{\mathrm{RRL}} F^{\frac{1}{2}}(t) dt, \tag{100}$$

where P is the orbital period, $F(t)$ is the instantaneous XUV flux, and EL and RRL correspond to the energy-limited and radiation/recombination-limited portions of the orbit, respectively. By noting as in (85) that

$$dt = \frac{r(E)}{an} dE = \frac{r(E)P}{2\pi a} dE, \tag{101}$$

where E is the eccentric anomaly, we may write

$$\begin{aligned} \langle \dot{M} \rangle_t = & \frac{B}{2\pi a} \int_0^{E_{\mathrm{crit}}} F^{\frac{1}{2}}(E) r(E) dE + \\ & \frac{A}{2\pi a} \int_{E_{\mathrm{crit}}}^{2\pi - E_{\mathrm{crit}}} F(E) r(E) dE + \\ & \frac{B}{2\pi a} \int_{2\pi - E_{\mathrm{crit}}}^{2\pi} F^{\frac{1}{2}}(E) r(E) dE, \end{aligned} \tag{102}$$

where

$$E_{\mathrm{crit}} = \cos^{-1}\left(\frac{1}{e} - \frac{r_{\mathrm{crit}}}{ae}\right) \tag{103}$$

is the value of the eccentric anomaly when $r = r_{\mathrm{crit}}$. The three integrals above follow from the fact that starting from pericenter ($E = 0$), the planet is (by assumption) in the RRL regime up until $E = E_{\mathrm{crit}}$, switches to EL until $E = 2\pi - E_{\mathrm{crit}}$, and completes the orbit in the RRL regime. By symmetry of the orbit, the first and last integrals are identical, so we may simplify:

$$\begin{aligned} \langle \dot{M} \rangle_t = & \frac{A}{2\pi a} \int_{E_{\mathrm{crit}}}^{2\pi - E_{\mathrm{crit}}} F(E) r(E) dE + \\ & \frac{B}{\pi a} \int_0^{E_{\mathrm{crit}}} F^{\frac{1}{2}}(E) r(E) dE. \end{aligned} \tag{104}$$

Now, noting that $F(E) = L_{\rm XUV}/4\pi r^2(E)$ and $r(E) = a(1 - e\cos E)$, we have

$$\langle \dot{M} \rangle_t = \frac{AL_{\rm XUV}}{8\pi^2 a^2} \int_{E_{\rm crit}}^{2\pi - E_{\rm crit}} \frac{dE}{1 - e\cos E} + \frac{B}{\pi a}\sqrt{\frac{L_{\rm XUV}}{4\pi}} \int_0^{E_{\rm crit}} dE$$

$$= \frac{AL_{\rm XUV}}{8\pi^2 a^2} \int_{E_{\rm crit}}^{2\pi - E_{\rm crit}} \frac{dE}{1 - e\cos E} + \frac{BE_{\rm crit}}{\pi a}\sqrt{\frac{L_{\rm XUV}}{4\pi}}. \tag{105}$$

By evaluating the integral, we may finally write

$$\langle \dot{M} \rangle_t = A\bar{F}\left[1 - \frac{2}{\pi}\tan^{-1}\left(\frac{(1+e)\tan(E_{\rm crit}/2)}{\sqrt{1-e^2}}\right)\right] + B\bar{F}^{\frac{1}{2}}\left[\frac{(1-e^2)^{\frac{1}{4}}E_{\rm crit}}{\pi}\right]. \tag{106}$$

Note, importantly, that the expressions above are valid only for planets that cross the critical radius during their orbit; that is, planets for which the expression $a(1-e) < r_{\rm crit} < a(1+e)$ holds. The mass loss rate for planets outside this region must be calculated via the method described at the beginning of the section.

Finally, note that the formalism derived here makes use of $K_{\rm tide}$ rather than its eccentric version $K_{\rm ecc}$, which I derived in §3.2.4.2. It is in principle possible to account for the tidal enhancement during the energy-limited portion of the orbit, but the resulting integral would not be analytic. Moreover, the tidal enhancement due to the eccentricity is important primarily near pericenter, where the mass loss is radiation/recombination-limited and therefore independent of $K_{\rm tide}$. Thus, while the method outlined above may underpredict the mass loss rate in some cases, the effect will be small.

3.2.5 Tidal Model

I calculate the evolution of the semi-major axis, the eccentricity and the rotation rates from Equations (36)–(38) and (50)–(52) for the CPL and CTL models, respectively. For simplicity, I set the obliquities of all planets to zero. Since the tidal locking timescale is very short for close-in planets[9] (Gu et al., 2003), I assume that the planet's rotation rate is given by the equilibrium value (Equation 41 or 55).

I calculate the stellar spin by assuming different initial periods (see Table 1) and evolving

[9]Tidal locking refers to the state in which a body's rotation rate is fixed by tidal forces at an equilibrium value. While for circular orbits, this implies $\omega_i = n$, in the general case of an eccentric orbit in the CPL model, the planet's rotation rate assumes a slightly super-synchronous value. See Barnes et al. (2013) for a discussion.

it according to the tidal equations, while enforcing conservation of angular momentum as the star contracts during the pre-main sequence phase. I neglect the effects of rotational braking (Skumanich, 1972), whereby stars lose angular momentum to winds and spin down over time. This leads to an overestimate of the spin rate at later times, but tidal effects should only be important early on, when the radii and the eccentricity are higher. Bolmont et al. (2012) recently modeled the coupling between stellar spin and tides, following the evolution of a "slow rotator" star ($P_0 \approx 10$ days) and a "fast rotator" star ($P_0 \approx 1$ day). In both cases, the stars sped up during the first $\sim 300 - 500$ Myr, after which time angular momentum loss became significant. However, tidal evolution is orders of magnitude weaker at such late times, so rotational braking should have a minimal effect on the tidal evolution. Moreover, as I showed in §1.4, in general it is the tide raised on the planet that dominates the evolution; as this depends on the *planetary* rotation rate, and not on the stellar rotation rate, my results are relatively insensitive to the details of the spin evolution of the star.

In the CPL model, I adopt typical gas giant values $10^4 \leq Q_p \leq 10^5$ for gas-rich super-Earths/mini-Neptunes and typical terrestrial values $10 \leq Q_p \leq 500$ for planets that have lost their envelopes; I assume stellar values in the range $10^5 \leq Q_\star \leq 10^6$. In the CTL model, I consider time lags in the range $10^{-3}\,\mathrm{s} \leq \tau_p \leq 10^1\,\mathrm{s}$ for gas-rich planets and $10^{-1}\,\mathrm{s} \leq \tau_p \leq 10^3\,\mathrm{s}$ once they lose their envelopes. Following Leconte et al. (2010), I consider stellar time lags in the range $10^{-2}\,\mathrm{s} \leq \tau_\star \leq 10^{-1}\,\mathrm{s}$.

Once a planet loses all of its envelope, I artificially switch Q_p or τ_p to the terrestrial value adopted in that run. In reality, as the atmosphere is lost, the transition from high to low Q_p (or low to high τ_p) should be continuous. A detailed treatment of the dependence of Q_p and τ_p on the envelope mass fraction is deferred to future work.

Finally, I note that the second-order CPL model described above is valid only at low eccentricity. For $e \geq \sqrt{1/19} \approx 0.23$, the phase lag of the dominant tidal wave discontinuously changes from negative to positive, such that the model then predicts outward migration due to the planetary tide. This effect is unphysical, stemming from the fact that the CPL model considers only terms up to second order in the eccentricity (for a detailed discussion of this, see Leconte et al., 2010). I therefore restrict all my calculations in the CPL framework to values of the eccentricity $e \leq 0.2$. For higher values of e, I use the higher-order CTL model.

3.2.5.1 Rate of Change of the Flux

Since in my model I couple atmospheric escape to tidal evolution, it is instructive to briefly consider how the flux incident on a planet changes as the planet tidally evolves. Conservation of angular momentum requires that for tidal evolution, the rates of change of the eccentricity and the semi-major axis must be related through

$$\frac{1}{e}\frac{de}{dt} = \frac{1}{2a}\frac{da}{dt}. \tag{107}$$

The rate of change of the flux due to orbital changes is

$$\begin{aligned}
\frac{dF}{dt} &= \frac{L_{\rm bol}}{4\pi}\frac{d}{dt}\left(a^2\sqrt{1-e^2}\right)^{-1} \\
&= -\frac{L_{\rm bol}}{4\pi}\left(a^2\sqrt{1-e^2}\right)^{-2}\left[2a\frac{da}{dt}\sqrt{1-e^2} + \left(\frac{a^2}{2\sqrt{1-e^2}}\right)\left(-2e\frac{de}{dt}\right)\right] \\
&= \frac{L_{\rm bol}}{4\pi}\frac{1}{a^2}\left[-\frac{2}{a}\frac{da}{dt}\frac{1}{(1-e^2)^{1/2}} + \frac{e}{(1-e^2)^{3/2}}\frac{de}{dt}\right] \\
&= -F\frac{de}{dt}\left[\frac{4}{e\,(1-e^2)^{1/2}} - \frac{e}{(1-e^2)^{3/2}}\right] \\
&= -F\frac{de}{dt}\left[\frac{4-5e^2}{e(1-e^2)^{3/2}}\right].
\end{aligned} \tag{108}$$

For $e \lesssim 0.7$, this simplifies to

$$\frac{dF}{dt} = -4F\frac{1}{e}\frac{de}{dt}, \tag{109}$$

which is equivalent to the trivial result

$$\frac{1}{F}\frac{dF}{dt} = -\frac{2}{a}\frac{da}{dt}. \tag{110}$$

Thus at low eccentricities, the flux will always increase as the orbit shrinks. However, at very high eccentricities ($e \gtrsim 0.8$), Equation (108) predicts that dF/dt is *negative* when $da/dt < 0$: the decrease in the flux due to the circularization of the orbit overpowers the increase due to the shrinking orbit.

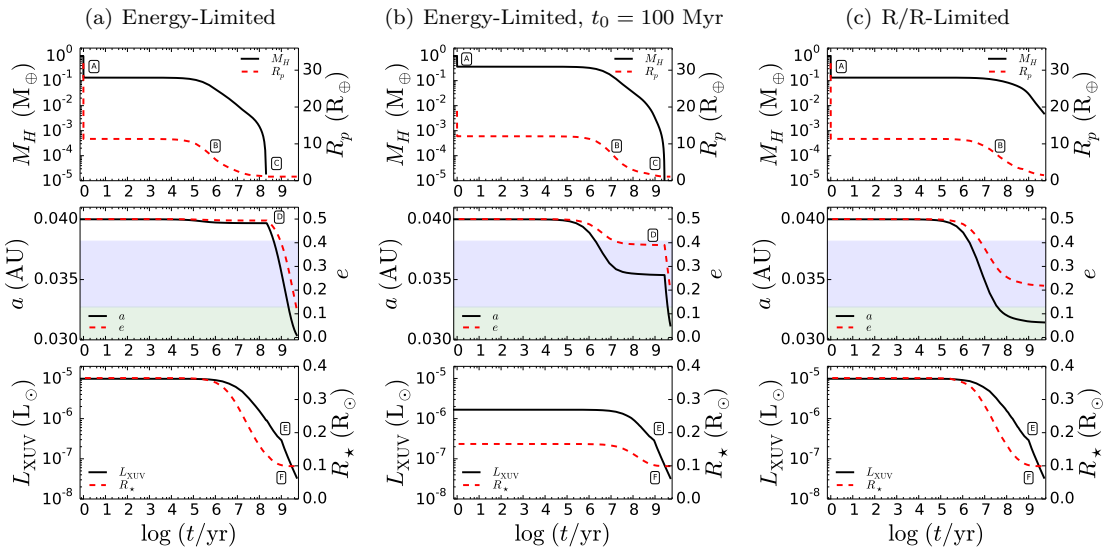

Figure 21: Three sample integrations of my code. The first row plots show the envelope mass (left axis) and planet radius (right axis) versus time since formation; the second row plots show the semi-major axis (left) and eccentricity (right) versus time; and the third row shows the stellar XUV luminosity (left) and stellar radius (right) versus time. The planet is initially a 1 $M_\oplus$ core with a 1 $M_\oplus$ envelope orbiting around a 0.08 $M_\odot$ M-dwarf with $e = 0.5$ at a semi-major axis of 0.04 AU, just outside the OHZ (light blue shading). Unless otherwise noted, all other parameters are set to their default values (Table 1). As it loses mass and tidally evolves, it migrates into the CHZ (light green shading). Note that the evolution of the HZ is not shown; the CHZ and OHZ are taken to be the long-term (> 1 Gyr) values. **(a):** In this run, I force the escape to be energy-limited, as in an X-ray dominated flow. The planet loses its entire envelope at $t \approx 100$ Myr. **(b):** Same as (a), except the calculation starts at $t_0 = 100$ Myr, corresponding to a planet that undergoes late migration. While the envelope still completely evaporates, this occurs at a much later time, $t \approx 2$ Gyr. **(c):** Same as (a), except that the escape is radiation/recombination limited above the critical flux. Here, the envelope does not fully evaporate and tidal migration is noticeably weaker. See the text for a discussion of the labels **A-F**.

3.3 Results

3.3.1 A Typical Run

In Figure 21 I show the time evolution of three gas-rich planets as a guide to understanding the results presented in the following sections. I plot planet mass and radius (top row), semi-major axis and eccentricity (center row), and stellar XUV luminosity and radius (bottom row), all as a function of time since the planet's initial migration, t_0, for 5 Gyr. In the center plots, the post-1 Gyr OHZ and CHZ are shaded blue and green, respectively. At $t = t_0$, the planet is a 1 $M_\oplus$ core with a 1 $M_\oplus$ envelope orbiting around a 0.08 $M_\odot$ M-dwarf with $e = 0.5$ at a semi-major axis of

0.04 AU. I set $\epsilon_{\rm XUV} = 0.30$; other parameters are equal to the default values listed in Table 1. Because of the high eccentricity, the tidal evolution is calculated according to the CTL model.

In column (a), I force the escape to be energy-limited (Equation 5), corresponding to an X-ray dominated flow. Prior to the first timestep, nearly 90 percent of the envelope mass is lost to RLO, indicated by the discontinuous drop marked **A** on the top plot. This is due primarily to the inflated radius shortly after formation, which reaches 30 $R_\oplus$ for a planet of age $t = t_0 = 10$ Myr. Once this mass is removed, the planet enters the energy-limited escape regime, which operates on a timescale of ~ 10 Myr (**B**). After ~ 100 Myr (**C**), the planet loses its entire envelope and becomes a HEC.

In the center plot, we see that the planet's orbit steadily decays as it circularizes, with a sharp discontinuity in the slope at $\sim$ 100 Myr (**D**), corresponding to when it transitions from a gaseous (low τ_p) to a rocky (high τ_p) body. The tidal evolution from that point forward is dramatically stronger, and e decreases to ~ 0.1 at 5 Gyr. The planet's semi-major axis decays by 25%, moving it well into the CHZ. As I noted earlier, the transition from low to high tidal time lags (or, alternatively, from high to low tidal quality factors) is likely to be gradual as the bulk of the energy shifts from being dissipated in the envelope to being dissipated in the core. In this case, the faster inward migration as τ_p increases is likely to accelerate the rate of mass loss, leading to slightly earlier evaporation times. However, given the large uncertainty in the values of τ_p and its dependence on planetary and orbital parameters, my current approach should suffice.

In the bottom plot, we see that the bulk of the mass loss occurs when the stellar XUV flux is high. After $t \approx 100$ Myr, the XUV luminosity is low enough that a planet with significant hydrogen left ($f_H \gtrsim 0.01$) is unlikely to completely evaporate. Here, the XUV saturation time is set to 1 Gyr, visible in the kink marked by the label **E**; prior to that time, the decrease in the XUV luminosity is simply a function of the rate of contraction of the star. After $t \sim 1$ Gyr (**F**), the stellar radius asymptotes to its main sequence value and the XUV flux decays as a simple power law.

In column (b) I repeat the integration but delay the start time, setting $t_0 = 100$ Myr. This corresponds to a planet that undergoes a late scattering event, bringing it to $a = 0.04$ AU when both the planet's radius and the XUV flux are significantly smaller. In this case, RLO is somewhat less effective, removing only 50% of the envelope initially (**A**). However, the planet still loses all of its hydrogen at $t \approx 2$ Gyr (**C**). Interestingly, because of its delayed evaporation, the planet's

eccentricity at 5 Gyr is significantly higher than in the previous run. This occurs because the transition from low to high τ_p (**D**) occurs much later. In this sense, a planet's current orbital properties can yield useful information about its atmospheric history. However, a more rigorous tidal model that accounts for the gradual change in Q_p and τ_p as f_H decreases is probably necessary to accurately infer the atmospheric evolution based on a planet's current eccentricity.

Finally, in column (c) I repeat the integration performed in (a), but this time set the flow to be radiation/recombination-limited above the critical flux (§1.3.1.4). Because of the significantly lower escape rate at early times, the envelope does not completely evaporate, and at 5 Gyr this is a super-Earth with slightly less than 1% hydrogen by mass. In order for a planet to completely lose its envelope in the radiation/recombination-limited regime, it must either migrate into an orbit closer to the star, have a larger eccentricity, have a smaller core, or be stripped by another process.

3.3.2 Dependence on M_c, M_H, $M_\star$, and t_0

In Figure 22 I plot the initial versus final envelope mass (M_H) for planets that end up in the IHZ (red lines), CHZ (green lines), and OHZ (blue lines). Line styles correspond to different values of t_0, the time at which the planet migrated into its initial close-in orbit (solid, 10 Myr; dashed, 50 Myr; dash-dotted, 100 Myr). The colored shadings are simply an aid to the eye, highlighting the spread due to different values of t_0. For reference, in dotted gray I plot the line corresponding to a planet that undergoes no evaporation. Note that in most plots, the curves approach this line as the initial M_H is increased: at constant core mass, it is in general more difficult to evaporate planets with more massive H/He envelopes. Finally, if a curve of a given color/line style is missing, the final hydrogen mass is zero for all values of the initial M_H.

The two columns correspond to runs in which the escape was forced to be energy-limited (left) and radiation/recombination-limited at high XUV flux and energy-limited at low XUV flux (right). Rows correspond to different planet properties, as labeled. In the top ("Default") row, the planet's core mass is set to $1\,M_\oplus$. The planet orbits an M dwarf with $M_\star = 0.08\,M_\odot$ in a circular orbit. All other parameters are set to their default values.

In the second row, I double the core mass. The third row is the same as the first, but for a $0.16\,M_\odot$ star; and the fourth row is the same as the first, but for an initial eccentricity $e = 0.4$. Since planets in this run undergo orbital evolution, the different color curves correspond to planets

that *end up* in the IHZ, CHZ, and OHZ, respectively (their initial semi-major axes are somewhat larger).

Let us first consider the general trends in the plots. For low enough M_H, all curves become nearly vertical. The *critical envelope mass* below which all mass is lost is marked with an X. Note that once planets lose sufficient mass such that $M_H \lesssim 10^{-4}$, the envelope is very unstable to complete erosion under the XUV fluxes considered here (corresponding to a near-vertical slope in this diagram). Many curves also display a flattening towards high initial envelope masses; some have prominent kinks beyond which the final envelope mass is constant. This is due to Roche lobe overflow, which causes any planets with radii larger than the Roche radius to lose mass prior to entering the HZ. Since increasing the envelope mass increases the planet radius, this results in a maximum envelope mass for some planets.

Planets that migrate in early (small t_0) lose significantly more mass than planets that migrate in late. This is a consequence of both the decay in the XUV luminosity of the star with time and the quick decrease of the planet radius as the planet cools. Another interesting trend is that the difference in the evaporated amount is much more pronounced in the energy-limited runs than in the radiation/recombination-limited runs. This is due to the $R_p^{3/2}$ scaling of the mass loss rate in the latter regime (versus the steeper R_p^3 scaling in the former). The difference in the initial radius across runs with different values of t_0 is less significant in the radiation/recombination-limited regime, resulting in more comparable mass loss rates. Note that the $F_{\rm XUV}^{1/2}$ scaling of the mass loss rate in this regime results in a weaker dependence on semi-major axis, as expected: the blue, green, and red curves (for a given t_0) are packed more closely together in the right column than in the left column.

Now let us focus on individual rows. For the default run (a $1\,M_\oplus$ core in a circular orbit around a $0.08\,M_\odot$ star), *all* planets in the IHZ lose all of their hydrogen and form HECs, regardless of migration time, envelope mass, or escape mechanism. In the CHZ, only planets that migrate within $\lesssim 50$ Myr and undergo energy-limited escape form HECs for all initial values of M_H. However, HECs still form from planets with $M_H \lesssim 0.5-0.9\,M_\oplus$ in the CHZ. In the OHZ, this is only possible for planets with less than about 1% H/He by mass.

At twice the core mass (second row), all curves shift up and to the left, approaching the zero evaporation line for planets in the OHZ. In the IHZ, HECs still form from planets with any initial

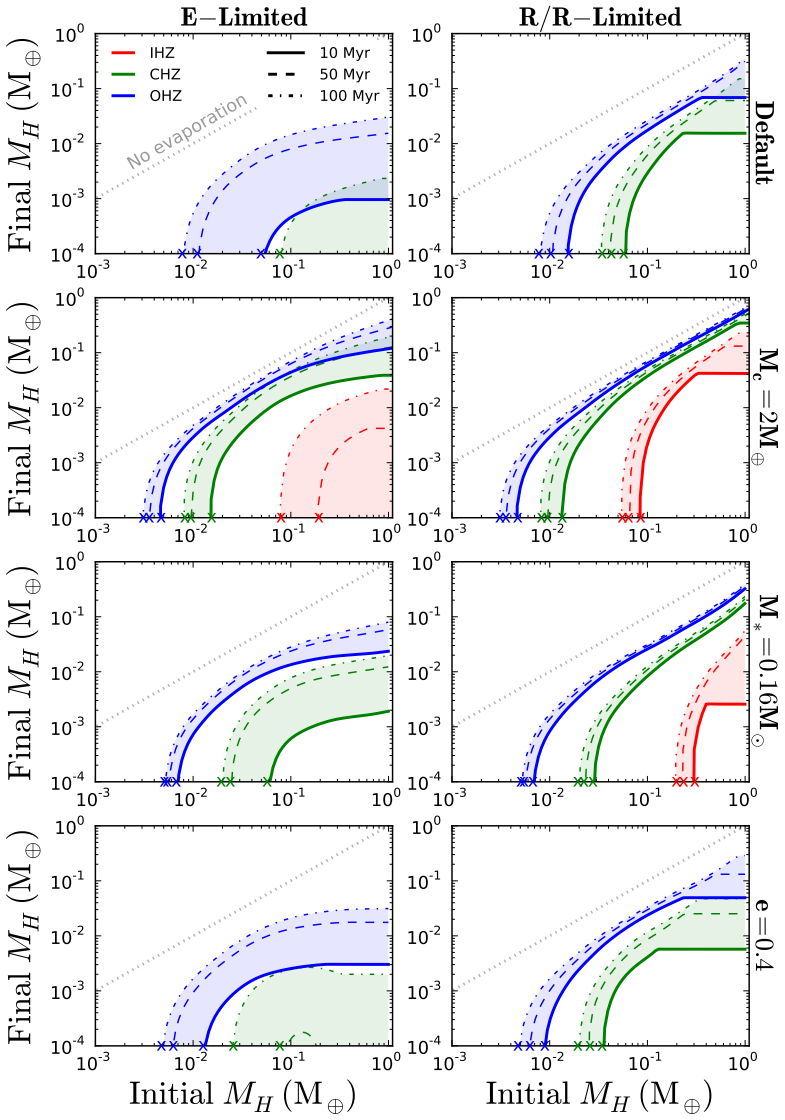

Figure 22: Initial versus final envelope mass (M_H) for planets that end up in the IHZ (red), CHZ (green), and OHZ (blue). Line styles correspond to different values of t_0 (solid, 10 Myr; dashed, 50 Myr; dash-dotted, 100 Myr). Columns correspond to runs in which the escape mechanism is energy-limited (left) and radiation/recombination-limited (right); rows vary certain parameters as labeled, with all others set to their default values. In the default run, the planet has a $1\,M_\oplus$ core and orbits an M dwarf with $M_\star = 0.08\,M_\odot$ in a circular orbit. The dotted gray line corresponds to a planet that undergoes no evaporation. An X marks the critical initial envelope mass below which full evaporation occurs within 5 Gyr. In some plots, curves of a given color/line style are missing; for those runs, the entire envelope was lost for all starting values of M_H.

hydrogen amount for $t_0 = 10$ Myr and in the energy-limited regime. In all other cases, HECs only form from planets with $M_H \lesssim 0.1\,\mathrm{M}_\oplus$. Of all the parameters I varied in my integrations, changing the core mass has the most dramatic effect on whether or not HECs can form. As I discuss below, habitable evaporated cores with masses greater than about $2\,\mathrm{M}_\oplus$ are unlikely in this model.

At higher stellar mass (third row), HECs are again more difficult to form, particularly in the radiation/ recombination-limited regime. Due to the more distant HZ, Roche lobe overflow is less effective in removing mass. The shorter super-luminous contraction phase of earlier M dwarfs also results in less total XUV energy deposition in the envelope. However, in the energy-limited regime, HECs still form from planets with up to 50% H/He envelopes in the IHZ.

Finally, the effect of a higher eccentricity (bottom row) is much more subtle. In general, these planets lose slightly less mass than in the default run, but the plots are qualitatively similar. At high eccentricity, the orbit-integrated mass loss rate is higher (see §3.2.4), but because of the orbital evolution, the planet must start out at larger a in order for it to end up in the HZ at 5 Gyr. These effects roughly cancel out: in general, habitable evaporated cores are just as likely on eccentric as on circular orbits. Note, also, that the green curves in the bottom right plot are the only ones that are non-monotonic. The effect is very small, but hints at an interesting coupling between tides and mass loss. At high initial M_H, the radius is large enough to drive fast inward migration (see below), exposing the planet to high XUV flux for slightly longer than a planet with smaller M_H (and therefore a smaller radius), resulting in a change in the slope of the curve at initial $M_H \approx 0.15\,\mathrm{M}_\oplus$.

3.3.3 The Role of Tides

Next, I consider in detail how tides affect the evolution of HECs. We saw in §1.4 that the net effect of tides is to induce inward migration and circularization of planet orbits in the HZ of M dwarfs, an effect that couples strongly to the atmospheric mass loss. For $e \lesssim 0.7$, the flux increases with time as planets tidally evolve, accelerating the rate of mass loss; at higher eccentricities, the flux actually decreases due to the circularization of the orbit (see §3.2.5.1). The changing mass and radius of the planet can then act back on the tidal evolution, either accelerating it (in cases where $|dM_p/dt| \gg |dR_p/dt|$) or decelerating it in a negative feedback loop (otherwise).

In Figure 23 I show the results of an integration of my code on a grid of tidal time lag τ_p versus

the XUV absorption efficiency $\epsilon_{\rm XUV}$. Colors correspond to the final hydrogen mass fraction, f'_H; evaporated cores occur in the white regions ($f'_H = 0$). Dark gray indicates planets that either migrated beyond the HZ or remained exterior to it and are therefore not habitable. These plots provide an intuitive sense of the relative importance of tidal evolution (y-axis) and energy-limited escape (x-axis) in determining whether or not a HEC is formed. Note that once a planet loses its gaseous envelope, I switch the time lag to $\tau'_p = \max(\tau_p,\ 64{\rm s})$, where the latter is a typical (gas-free) tidal time lag of a rocky planet (see, for instance, Barnes et al., 2013).

In the default run (a), I consider a planet with a core mass of $1\,{\rm M}_\oplus$ and an envelope mass of $1\,{\rm M}_\oplus$ ($f_H = 0.5$) orbiting at 0.07 AU in a highly eccentric ($e = 0.8$) orbit around a $0.08\,{\rm M}_\odot$ star. The mass loss mechanism is radiation/recombination-limited escape at high XUV flux and energy-limited at low XUV flux. At a given value of log τ_p, say 0, the final hydrogen fraction is a strong decreasing function of $\epsilon_{\rm XUV}$, as expected: the higher the evaporation efficiency, the smaller the final envelope. Interestingly, the dependence of f'_H on the tidal time lag can be nearly just as strong. At a given value of $\epsilon_{\rm XUV}$, say 0.35, f'_H depends strongly on τ: the final hydrogen envelope mass can range from zero to up to about 3% the planet mass. In this case, the tidal evolution controls whether or not a HEC forms. This is due to the fact that at large τ_p, tidal migration is fast, bringing the planet into the HZ while its radius is still inflated and the stellar XUV luminosity is higher. Planets with lower values of τ_p migrate in later and undergo slower mass loss. At very low τ_p, tidal migration is too weak to bring the planets fully into the HZ. The effect is stronger for higher $\epsilon_{\rm XUV}$: these are planets whose radii decrease very quickly (due to the fast mass loss), slowing down the rate of tidal evolution and keeping them outside of the HZ for longer.

In plot (b), I increase the core mass slightly to $1.5\,{\rm M}_\oplus$. The effect on the mass loss is significant, and HECs no longer form. At high $\epsilon_{\rm XUV}$ and high τ_p, the lowest value of f'_H is about 1%. This reinforces what I argued above: HECs are most likely for the lowest mass cores.

In plot (c), I instead decrease the eccentricity to 0.7. As in (b), HECs no longer form in these runs due to the decreased strength of the tidal evolution, and again the minimum f'_H is about 1%. Note that this does not mean that HECs are more likely at higher eccentricity in general—this is only the case here because I fix the initial semi-major axis at 0.07 AU. At fixed *final* semi-major axis (i.e., what we can readily observe in actual systems), planets on initially circular orbits will still lose more mass than planets on initially eccentric orbits (which must have formed farther out).

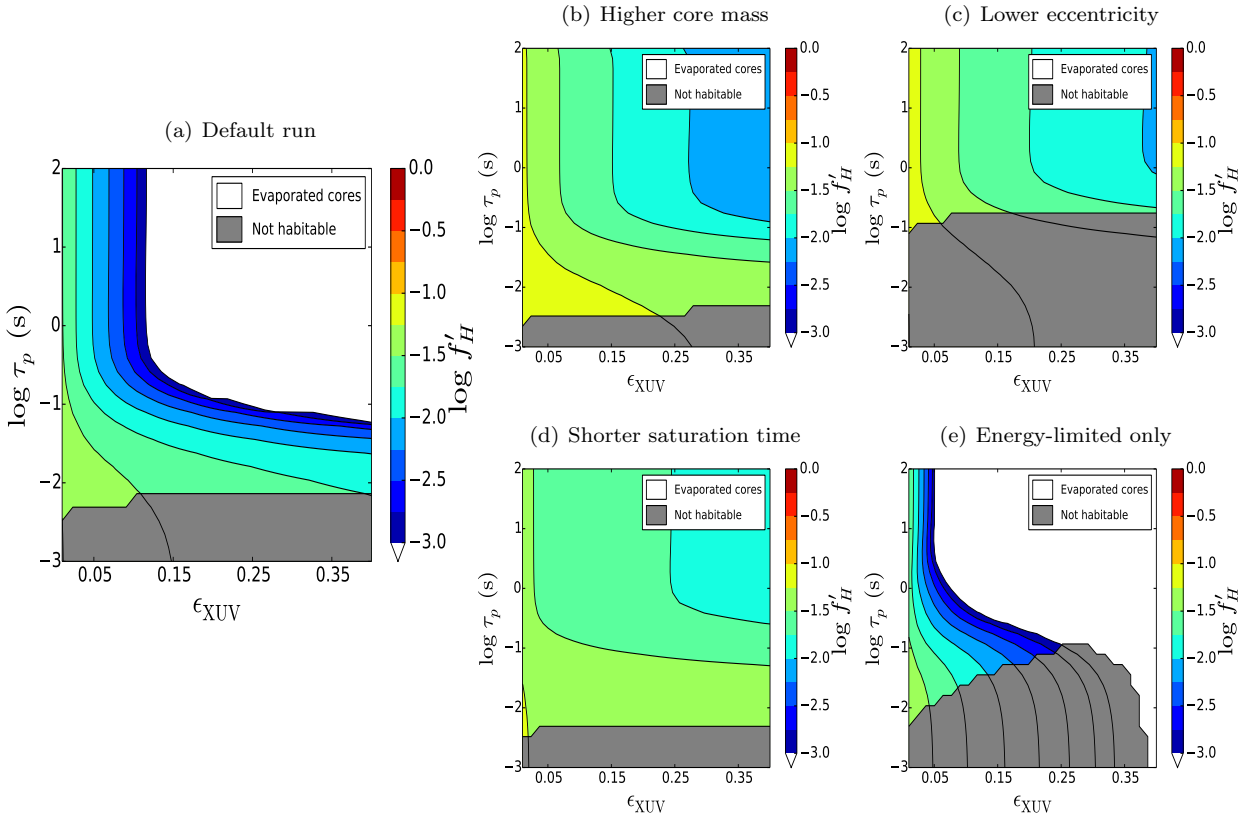

Figure 23: Contours of the log of the hydrogen mass fraction f'_H at 5 Gyr as a function of the tidal time lag τ_p and the XUV absorption efficiency $\epsilon_{\rm XUV}$ for five integrations of my code. White regions correspond to planets that completely lost their envelopes; dark gray regions indicate planets that are not in the HZ at 5 Gyr. **(a)** The default run. The planet has a core mass of 1 $M_\oplus$ with $f_{H_0} = 0.5$, orbiting around a 0.08 $M_\odot$ star in an initially highly eccentric orbit ($e = 0.8$) at $a = 0.07$ AU. All other parameters are set to their default values (Table 1), and the escape mechanism is radiation/recombination-limited at high XUV flux and energy-limited at low XUV flux. Note that the final envelope mass highly depends on *both* $\epsilon_{\rm XUV}$ and τ_p. **(b)** The same as (a), but for a higher core mass $M_c = 1.5\,M_\oplus$. No evaporated cores form in this scenario. **(c)** The same as (a), but for a lower eccentricity $e = 0.7$. Again, no habitable evaporated cores form, and the planet remains outside of the habitable zone for a larger range of τ_p. **(d)** The same as (a), but for a shorter XUV saturation time of the parent star, $t_{\rm sat} = 0.1$ Gyr. No evaporated cores form. Since the XUV flux drops off much more quickly, energy-limited escape is less effective in removing mass and thus f'_H is a weaker function of $\epsilon_{\rm XUV}$. **(e)** The same as (a), but for energy-limited escape only, which would be the case if the flow is X-ray dominated. Note that in this case whether or not the planet becomes an evaporated core is a much stronger function of both τ_p and $\epsilon_{\rm XUV}$.

In plot (d), I decrease the saturation time to $t_{sat} = 0.1$ Gyr, which is typical of early M/late K dwarfs (Jackson et al., 2012). No HECs form, and the final hydrogen fraction is less sensitive to both τ_p and $\epsilon_{\rm XUV}$: in general, mass loss is significantly suppressed.

Finally, in plot (e), I force the escape mechanism to be energy-limited only. Since the escape is now entirely controlled by $\epsilon_{\rm XUV}$, the dependence on this parameter is naturally much stronger, and complete evaporation occurs in this case for $\epsilon_{\rm XUV} \gtrsim 0.32$ at any τ_p. Because evaporation occurs more quickly in this case than in the other plots, at any given time the planet radii are smaller, resulting in less efficient migration at a given τ_p. More planets therefore do not make it into the HZ. Interestingly, for very large $\epsilon_{\rm XUV}$ ($\gtrsim 0.4$), all planets make it into the HZ. This is due to the fact that these planets transition from gaseous (low τ_p) to gas-free (high τ_p, equal to 64s in these runs) very early on. Despite their lower radii, they benefit from the stronger tidal dissipation of fully rocky bodies and are able to make it into the HZ after 5 Gyr.

In general, the coupling between tides and mass loss is quite complex. For planets with high initial eccentricities, this coupling can ultimately determine whether or not a HEC will form. I discuss this in more detail in §3.3.4.

3.3.4 Evaporated Cores in the Habitable Zone

Having shown that HECs are possible in certain cases, I now explore where in the habitable zone one may expect to find them. For a "terrestrial" (I use this term rather loosely[10]) planet detected in the HZ, it would be very informative to understand whether or not it may be the evaporated core of a gaseous planet, since its past atmospheric evolution may strongly affect its present ability to host life. To this end, I ran eight grids of 2.7×10^6 integrations each of my evolution code under different choices of parameters in Table 1. For initial semi-major axes in the range $0.01 \leq a_0 \leq 0.5$, initial eccentricities in the range $0 \leq e_0 \leq 0.95$, and stellar masses between $0.08\,{\rm M}_\odot \leq M_\star \leq 0.5\,{\rm M}_\odot$, I calculate the final (i.e., at $t_{\rm stop} = 5$ Gyr in the default run) values of a, e, and the envelope mass M_H. I then plot contours of M_H as a function of the stellar mass and the *final* semi-major axis and eccentricity (that is, the observable parameters). Since in principle, the mapping $(a_0, e_0, M_{H_0}) \rightarrow (a, e, M_H)$ is not necessarily a bijection (due to the nonlinear coupling

[10] As I explain below, it is quite possible that HECs are all but terrestrial, since it is likely that they have large ice mass fractions and are compositionally distinct from Earth. By "terrestrial," in this case, I mean planets that are similar in size and mass to Earth and are not gaseous; these may or may not have a rocky surface like Earth. My definition of "terrestrial" therefore encompasses water worlds.

between tides and mass loss and the finite resolution of my grid), the function $M_H(a, e)$ may be multi-valued at some points. In these cases, I take M_H at coordinates (a, e) to be the *minimum* of the set of all final values of M_H that are possible at (a, e). Because of this choice, the $M_H = 0$ contour in a versus $M_\star$ plots (Figures 24-26) separates currently terrestrial planets that *could be* the evaporated cores of gas-rich planets (left) from currently terrestrial planets that have *always* been terrestrial (right). In other words, I am showing where evaporated cores are ruled out.

In Figure 24 I perform the calculations described above for planets with $M_c = 1\,\mathrm{M}_\oplus$ and initial $f_H = 0.5$ (i.e., $M_H = 1\,\mathrm{M}_\oplus$). The habitable zone is plotted in the background for reference, where again the IHZ, CHZ, and OHZ are indicated by the red, green, and blue shading, respectively. Line styles correspond to different final values of the eccentricity (0, solid; 0.25, dashed; and 0.5, dash-dotted). Dark red lines correspond to the $M_H = 0$ contours in the radiation/recombination-limited escape model; dark blue lines correspond to the energy-limited model. Note that although I designate my two models as "energy-limited" and "radiation/recombination-limited", at low XUV fluxes the escape is energy-limited in *both* models, since below F_{crit} the energy-limited escape rate is always smaller than the radiation/recombination-limited escape rate. In this sense, the "radiation/recombination-limited model" is always conservative, as the mass loss rate is set to the minimum of Equations (5) and (13).

As an example of how to interpret this figure, consider a $1\,\mathrm{M}_\oplus$ rocky planet discovered orbiting a $0.2\,\mathrm{M}_\odot$ star at 0.1 AU, squarely within the CHZ. Since this planet lies to the left of all blue curves, under the assumption of energy-limited escape it could be a habitable evaporated core. On the other hand, if the atmospheric escape were radiation/recombination-limited, this planet must have always been terrestrial.

Next, consider a putative rocky planet discovered around the same $0.2\,\mathrm{M}_\odot$ star skirting the inner edge of the HZ (i.e., at 0.07 AU). Under the energy-limited assumption, this planet could be a HEC. For radiation/recombination-limited escape, however, whether or not it is a HEC depends on its present eccentricity. If the planet is currently on a circular orbit, we infer that it has always been terrestrial (since it lies to the right of the $e = 0$ contour). However, since the planet lies to the left of the higher eccentricity contours, if $e \gtrsim 0.25$, it could be a HEC.

One may conclude from Figure 24 that if energy-limited escape is the dominant mechanism around M dwarfs, planets with $M_p \sim 1\,\mathrm{M}_\oplus$ in the CHZ and IHZ of these stars could be habitable

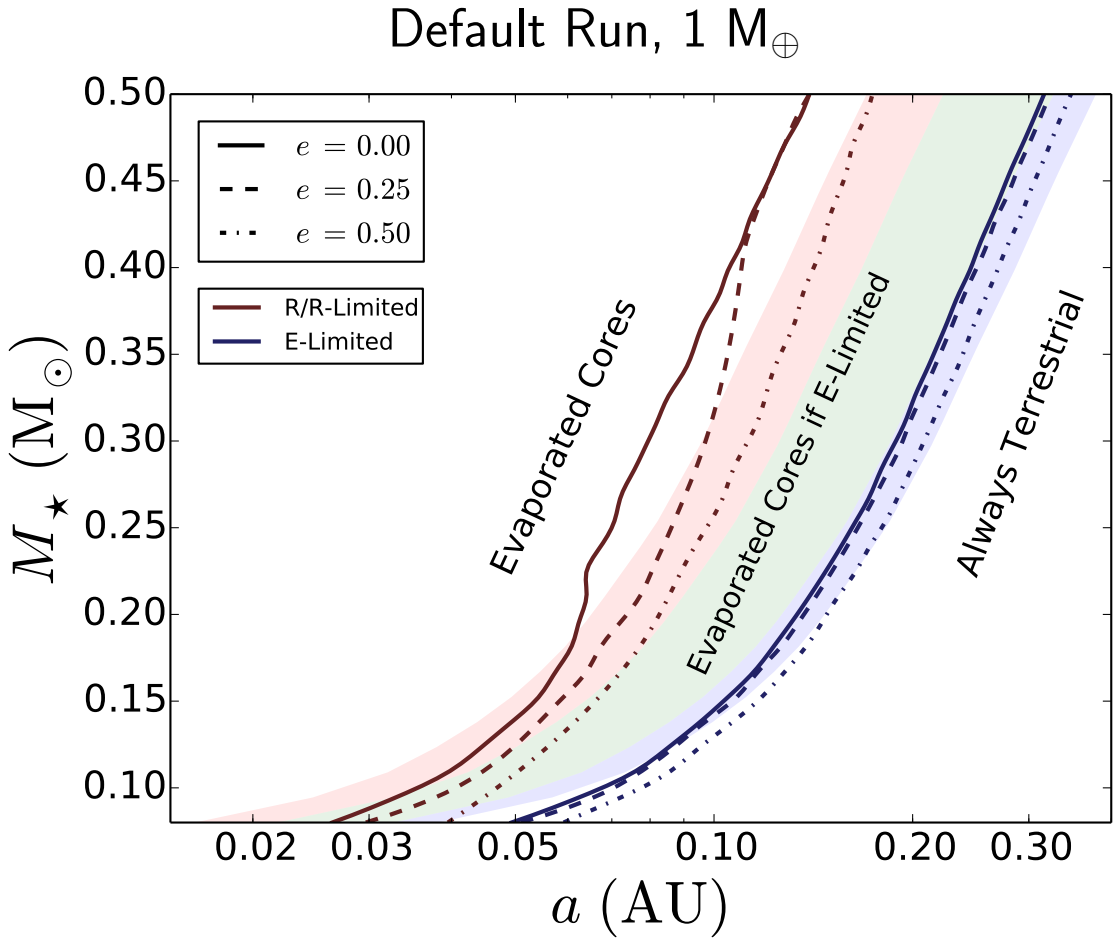

Figure 24: Regions of parameter space that may be populated by HECs, for $M_c = 1\,\mathrm{M}_\oplus$, $M_H \leq 1\,\mathrm{M}_\oplus$, and default values for all other parameters. Terrestrial planets detected today occupying the space to the left of each contour line could be the evaporated cores of gaseous planets with $f_H \leq 0.5$. Planets detected to the right of the contour lines have always been terrestrial/gaseous. Dark red lines correspond to the conservative mass loss scenario, in which mass loss is radiation/recombination-limited at high XUV flux and energy-limited at low XUV flux. Dark blue lines correspond to mass loss via the energy-limited mechanism only. Planets around stars with significant X-ray emission early on are likely to be in the latter regime. Different line styles correspond to different eccentricities today. Terrestrial planets detected at higher eccentricity (dashed and dash-dotted lines) could be evaporated cores at slightly larger orbital separations than planets detected on circular orbits (solid lines). Note that in the energy-limited regime, *all* $1\,\mathrm{M}_\oplus$ terrestrial planets in the HZ of low-mass M dwarfs could be habitable evaporated cores. At higher stellar mass, HECs are restricted to planets in the CHZ and IHZ. In the radiation/recombination-limited regime, the accessible region of parameter space is smaller, but around the lowest mass M dwarfs HECs are still possible in the CHZ.

evaporated cores. For $M_\star \lesssim 0.15\,\mathrm{M}_\odot$, HECs may exist throughout the entire HZ. If, on the other hand, these planets are shaped mostly by radiation/recombination-limited escape, HECs are only possible in the IHZ for $M_\star \lesssim 0.2\,\mathrm{M}_\odot$ and in the CHZ of M dwarfs near the hydrogen-burning limit.

The effect of the eccentricity is significant primarily in the radiation/recombination-limited regime, where a lower mass loss rate keeps the planet's radius large for longer than in the energy-limited regime, allowing it to migrate into the HZ faster, thereby enhancing the mass loss rate. In some cases, particularly near the inner edge of the HZ, the present-day eccentricity yields important information about a planet's evolutionary history: depending on the precise value of e, a given terrestrial planet may or may not be a HEC. However, I urge caution in interpreting the results in Figure 24 at nonzero eccentricity, given the large uncertainty in the tidal parameters of exoplanets.

On this note, it is important to bear in mind that the curves in Figure 24 are a function of my choice of parameters in Table 1. In order to assess the impact of my choice of "default" parameters on these contours, in Figure 25 I repeat the calculation for more conservative values of the two parameters that govern the mass loss mechanism: the XUV saturation time t_{sat} and the absorption efficiency ϵ_{XUV}. In this figure, I choose $t_{\mathrm{sat}} = 0.1$ Gyr, the nominal value for earlier K/G dwarf stars (§1.1), and $\epsilon_{\mathrm{XUV}} = 0.15$.

In this grid, all curves shift quite dramatically to lower a, and HECs are no longer possible under radiation/recombination-limited escape. For energy-limited escape, HECs are confined to the IHZ for $M_\star \lesssim 0.15\,\mathrm{M}_\odot$ and the CHZ for $M_\star \lesssim 0.1\,\mathrm{M}_\odot$ at all eccentricities considered here.

Given the large difference between the results of Figures 24 and 25, care must be taken in assessing whether a planet may be a HEC. Since it is likely that the XUV saturation time is much longer for M dwarfs than for K and G dwarfs, and since my goal at this point is to separate regions of parameter space where HECs can and cannot exist, Figure 24 is probably the more relevant of the two. I discuss this point further in §3.4.

In Figure 26 I raise the core mass to $M_c = 2\,\mathrm{M}_\oplus$. HECs are now possible only in the IHZ and only in the energy-limited regime. At the lowest $M_\star$, HECs may be possible in the CHZ (energy-limited) and at the very inner edge of the HZ (radiation/recombination-limited). At higher eccentricity, the parameter space accessible to HECs is slightly higher in the energy-limited regime, but it is still a small effect overall. For a conservative choice of escape parameters ($t_{\mathrm{sat}} = 0.1$ Myr and $\epsilon_{\mathrm{XUV}} = 0.15$) with $M_c = 2\,\mathrm{M}_\oplus$, no HECs form anywhere in the HZ.

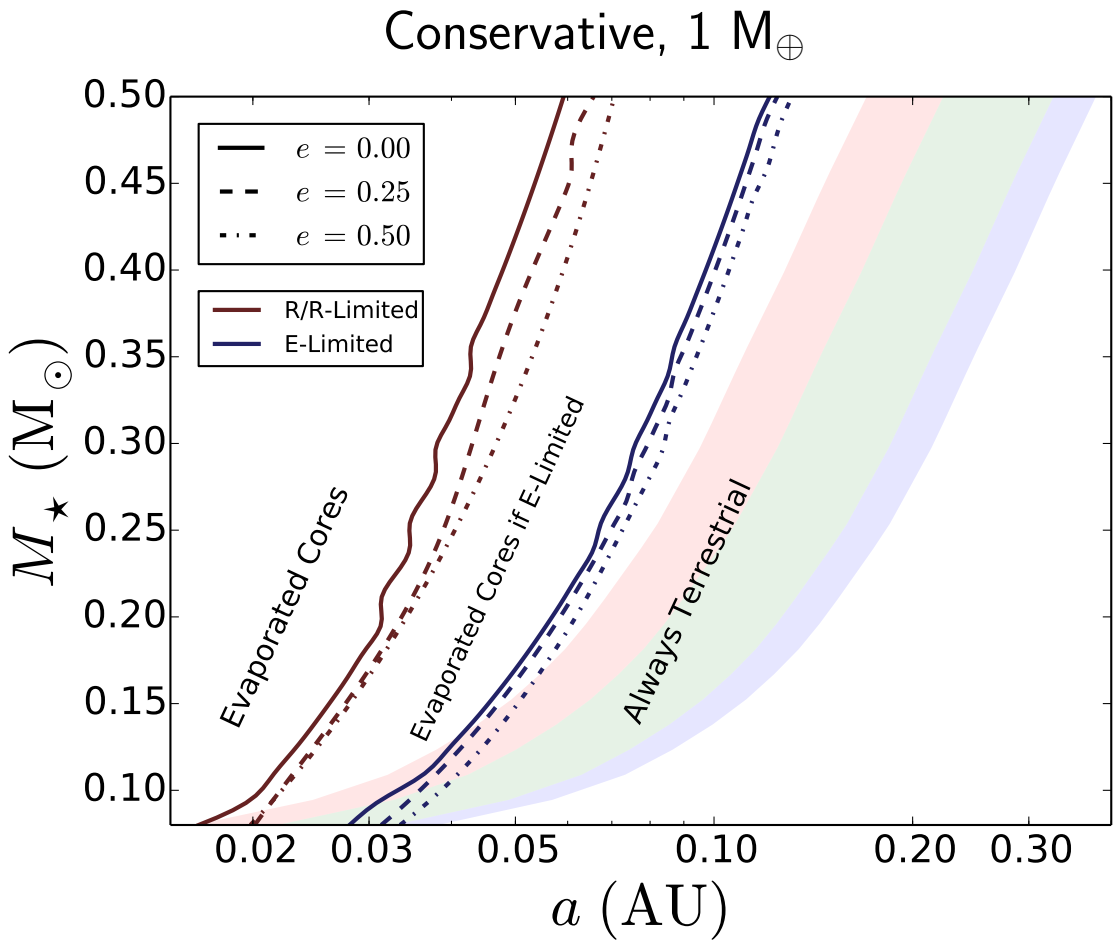

Figure 25: The same plot as Figure 24, but for a conservative choice of the parameters governing atmospheric escape: a short XUV saturation time $t_{\rm sat} = 0.1$ Gyr and a low XUV absorption efficiency $\epsilon_{\rm XUV} = 0.15$. In this case, HECs are no longer possible for radiation/recombination-limited escape. For energy-limited escape, HECs are only possible in the IHZ of low-mass M dwarfs and in the CHZ of M dwarfs at the hydrogen-burning limit. At high eccentricity ($e = 0.5$), HECs are only marginally more likely.

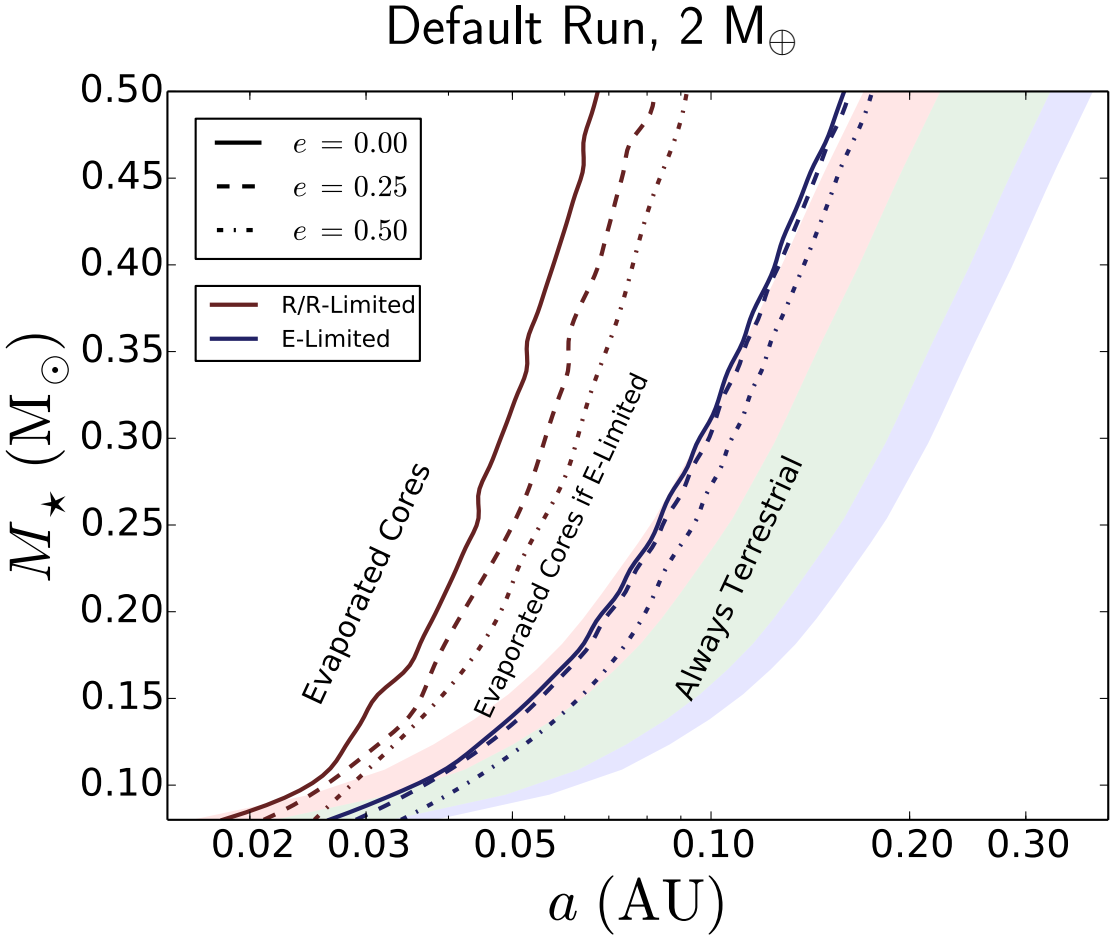

Figure 26: The same as Figure 24, but for a higher core mass $M_c = 2\,\mathrm{M}_\oplus$. HECs are now confined mostly to the IHZ of low-mass M dwarfs in the energy-limited regime. For planets undergoing radiation/recombination-limited escape, super-Earth HECs may not be possible.

3.4 Discussion

3.4.1 Initial Conditions

I have shown that HECs can form from small gas-rich super-Earths/mini-Neptunes ($M_p \lesssim 1\,\mathrm{M}_\oplus$) with H/He mass fractions as high as $f_H = 0.5$ around mid- to late M dwarfs. However, instrumental to my conclusions is the assumption that these planets formed beyond the snow line and migrated quickly into the HZ at $t = t_0$. Planets that form *in situ* in the HZ are unlikely to accumulate substantial H/He envelopes due to large disk temperatures and relatively long formation timescales. Significant gas accretion can only occur prior to the dissipation of the gas disk, which occurs on a timescale of a few to ~ 10 Myr; in particular, Lammer et al. (2014) showed that a 1 $\mathrm{M}_\oplus$ terrestrial planet generally accretes less than $\sim 3\%$ of its mass in H/He in the HZ of a solar-type star. Moreover, planets that form *in situ* do not undergo Roche lobe overflow, since accretion of any gas that would lead to $R_p > R_{\mathrm{Roche}}$ simply will not take place.

But can these gas-rich planets easily migrate into the HZ? The large number of recently discovered hot Neptunes and hot Super-Earths (e.g., Howard et al., 2012) suggests that inward disk-driven migration is an ubiquitous process in planetary systems, as it is highly unlikely that these systems formed *in situ* (Raymond and Cossou, 2014). Moreover, systems such as GJ 180, GJ 422, and GJ 667C each have at least one super-Earth/mini-Neptune in the HZ (Tuomi et al., 2014; Anglada-Escudé et al., 2013), some or all of which may have migrated to their current orbits. On the theoretical front, N-body simulations by Ogihara and Ida (2009) show that migration of protoplanets into the HZ of M dwarfs is efficient due to the proximity of the ice line and the fact that the inner edge of the disk lies close to the HZ. Super-Earths and mini-Neptunes that assemble early beyond the snow line could in principle also migrate into the HZ, but the likelihood that they migrate inwards and the distribution of these planets throughout the HZ needs to be investigated further in order to constrain the likelihood of formation of HECs.

A second issue is whether or not Earth-mass cores can accrete large H/He envelopes in the first place. One of my key results is that Earth-mass HECs can form from planets with initial envelope mass fractions of up to 50%; this would require a 1 $\mathrm{M}_\oplus$ planetary embryo to accrete an equivalent amount of gas from the disk. While such a planet would be stable against RLO beyond the snow line, whether or not it could have formed is unclear. Gas accretion takes place in two different

regimes: (i) a slow accretion regime, in which the envelope remains in hydrostatic equilibrium and gains mass only as it cools and contracts, evacuating a region that is then filled by nebular gas; and (ii) a runaway accretion regime, in which the rapidly increasing mass of the envelope leads to an increase in the size of the Hill sphere and increasingly faster gas accretion (e.g., Pollack et al., 1996). Since the critical core mass for runaway accretion is thought to be somewhere in the vicinity of 10–20 $M_\oplus$ (Pollack et al., 1996; Rafikov, 2006), the progenitors of HECs must accrete their gas slowly. As I mentioned above, Lammer et al. (2014) find that Earth-mass planets typically form with $f_H \lesssim 0.03$ in the HZ of solar-type stars. This is consistent with recent analyses of Kepler planet data by Rogers (2015) and Wolfgang and Lopez (2014), who find that most planets with radii less than about 1.5 $R_\oplus$ (corresponding to masses less than $\sim 5M_\oplus$) are rocky, with typical H/He mass fractions of about 1%.

Formation beyond the snow line can increase f_H: Rogers et al. (2011) show that a planet with core mass 2.65 $M_\oplus$ accretes only 0.54 $M_\oplus$ of gas (corresponding to $f_H \approx 17\%$). Bodenheimer and Lissauer (2014), on the other hand, show that planets with cores in the range 2.2–2.5 $M_\oplus$ generally accrete less than 10% of their mass in H/He. However, these authors terminated gas accretion at 2 Myr. For a longer disk lifetime of 4 Myr, Bodenheimer and Lissauer (2014) show that a planet with core and envelope masses of 2.8 $M_\oplus$ and ~ 1.2 $M_\oplus$, respectively, can form ($f_H \approx 30\%$). This is particularly relevant to planet formation around M dwarfs, as these stars may have disk lifetimes in excess of 5 Myr (Carpenter et al., 2006; Pascucci et al., 2009), which could allow for larger initial H/He envelope fractions.

Nevertheless, while a 1 $M_\oplus$ core with a 1 $M_\oplus$ envelope is probably an unlikely initial condition, it is important to keep in mind that my results apply to planets that form with smaller H/He fractions as well. Figures 24–26 show where planets with *up to* 50% H/He can form HECs; planets with lower H/He mass fractions also form HECs at the same values of a and $M_\star$. See §3.4.6 for a more detailed discussion.

3.4.2 Are HECs Habitable?

Under the core accretion mechanism, mini-Neptunes and gas-rich super-Earths are likely to form close to or beyond the snow line, where disk densities are higher due to the condensation of various types of ices; these planets should therefore have large quantities of volatiles, including water,

ammonia, and CO_2 ices. Assuming a disk composition similar to that around the young Sun, the bulk composition of their cores would probably be similar to that of comets: roughly equal parts ice and silicate rock, similar to what studies predict for the composition of water worlds (Léger et al., 2004; Selsis et al., 2007). Once these planets have migrated into the HZ and lost their envelopes, it is quite likely that they would be water worlds and therefore not "terrestrial" in the strictest sense of the word. Whether or not such planets are habitable may depend on their ability to sustain active geochemical cycles, which are crucial for life on Earth today.

One concern is a possible interruption of the carbon cycle by a high pressure ice layer at the bottom of the ocean (Léger et al., 2004). Alibert (2014) calculated the critical ocean mass for high pressure ice formation, finding that it lies between 0.02 and 0.03 $M_\oplus$ for an Earth-mass planet. If HECs are in fact comet-like in composition, a deep ice layer would separate their oceans from their mantles, which could inhibit the recycling of carbon and other bioessential elements between the two reservoirs, a process that is critical to life on Earth. However, whether this is the case is far from settled, as processes involving solid state ice convection could mediate the cycling of these elements. In particular, Levi et al. (2013) and Kaltenegger et al. (2013) presented a mechanism that could recycle CH_4 and CO_2 between the interior and the atmosphere of water worlds, invoking the ability of these molecules to form clathrates that could be convectively transported through the ice to the surface. Other mechanisms could also exist, and without further modeling and laboratory work, it is unclear whether a high pressure ice layer poses a threat to habitability.

The probable difference in composition between HECs and Earth is likely to have other geophysical implications. Hydrogen-rich compounds such as methane and ammonia could make up a non-negligible fraction of a HEC's mass, leading to extremely reducing conditions at the surface. These could also end up in a secondary atmosphere along with large quantities of CO_2, which could lead to strong greenhouse heating, although it is likely that most of the CO_2 and NH_3 would be sequestered in the ice mantle (Léger et al., 2004). The compositional difference of HECs would also likely lead to mantle convection and tectonic activity different from Earth's, as well as potential differences in magnetic field generation by a possible core dynamo. Since both an active tectonic cycle and a magnetic field may be necessary for habitability, these issues need to be investigated further.

Also critical to the habitability of a HEC is its ability to outgas a secondary atmosphere once

its primordial envelope is lost. In particular, the secondary atmosphere must be stable against erosion. While the power-law decay in XUV emission after $t_{\rm sat}$ could allow for such an atmosphere to form, low mass M dwarfs may remain active for $t \gg 1$ Gyr. Strong planetary magnetic fields could potentially shield these atmospheres; Segura et al. (2010) showed that flares on the extremely active M dwarf AD Leo would have a small effect on the atmosphere of an Earth-like planet in the HZ provided it has an Earth-like magnetic field. However, interactions with the stellar wind could still pose serious problems to these and other planets in the HZs of M dwarfs. In particular, the high XUV/EUV fluxes of active M dwarfs can lead to significant expansion of the upper atmosphere, potentially causing the radius of the exobase to exceed the distance to the stellar magnetopause (Lichtenegger et al., 2010; Lammer et al., 2011b). For a pure N_2 atmosphere, Lichtenegger et al. (2010) showed that nitrogen ionized by EUV radiation above the exobase is subject to ion pickup by the solar wind, leading to the complete erosion of a 1 bar atmosphere in as short a time as 10 Myr. A stronger planetary magnetic field or large quantities of a heavier background gas such as CO_2 may be necessary to suppress ion pickup and preserve secondary atmospheres on HECs.

Given the complex processes governing the habitability of HECs, a detailed investigation of these issues is left to future work. As I discuss below, stronger constraints on the details of the X-ray/EUV evolution of M dwarfs of all masses are essential to understanding the fate and ultimately the habitability of planets around these stars.

3.4.3 The Need For Better Constraints

Figure 24 shows that the formation of HECs depends strongly on the atmospheric escape mechanism. As I mentioned earlier, whether a flow is closer to radiation/recombination-limited or energy-limited will depend on the ratio of the X-ray to EUV luminosity of the parent star. Low-mass M dwarfs may have XUV luminosities as high as $\sim 4 \times 10^{29}$ erg/s early on (§1.1). If X-rays contribute more than a few percent of this luminosity, low-mass low-density planets in the HZs of these stars may undergo energy-limited X-ray-driven escape (Figure 11 in Owen and Jackson, 2012). Unfortunately, knowledge of the exact age-luminosity relation in the X-ray and EUV bands for M dwarfs is still very poor, in great part because of the large uncertainties on these stars' ages. However, recent studies suggest that X-rays contribute a significant fraction of this luminosity (for a review, see Scalo et al., 2007). In particular, Stelzer et al. (2013) report high ($L_X \gtrsim 10^{29}$ erg/s)

X-ray luminosities for a sample of early active M dwarfs in the solar neighborhood, with a steeper age dependence than in FUV and NUV bands, which dominate the emission for $t \gtrsim 1$ Gyr. This is consistent with Owen and Jackson (2012), who argue that close-in mini-Neptunes may undergo a transition from X-ray-driven escape at early times to EUV-driven escape at later times.

Moreover, atmospheric X-ray heating and cooling is primarily done by metals. As Owen and Jackson (2012) point out, atmospheric composition is also likely to play a role in determining whether hydrodynamic flows are EUV- or X-ray-driven. Additionally, the presence of dust in the envelopes of these planets could greatly affect their ability to cool via Lyman α radiation. Absorption of recombination radiation by dust particles lifted high into the envelope by vigorous convection could convert a significant fraction of this energy into heating, which could effectively increase the absorption efficiency $\epsilon_{\rm XUV}$ and bring the flow closer to the energy-limited regime. Unfortunately, my present parametric escape model is unable to address the effect of dust and metal abundances on the escape rate—this issue needs to be revisited in the future with more sophisticated hydrodynamic models.

The difference between Figure 24 and Figure 25 is just as significant. At lower $t_{\rm sat}$ and $\epsilon_{\rm XUV}$, the HEC parameter space is greatly reduced. As I discuss in §1.1, the lower value $t_{\rm sat} = 0.1$ Gyr is more representative of K and G dwarfs, while M dwarfs may remain saturated for $t_{\rm sat} > 1$ Gyr. In fact, the energy-limited contours in Figure 24 closely trace the CHZ/OHZ boundary as $M_\star$ increases, predicting that HECs may even be possible around solar type stars (not shown). This is of course not the case, since solar-type stars are known to leave the saturation phase around $t_{\rm sat} = 0.1$ Gyr (Ribas et al., 2005). For some stellar mass near the M/K Dwarf transition ($0.5\,{\rm M}_\odot \lesssim M_{\star,crit} \lesssim 0.7\,{\rm M}_\odot$), Figure 25 becomes a better description for HECs, but for low-mass M dwarfs Figure 24 is probably more appropriate.

Naturally, the exact value of $\epsilon_{\rm XUV}$ will also affect the possible distribution of HECs within the HZ. Recently, Shematovich et al. (2014) performed numerical simulations to solve the kinetic Boltzmann equation for XUV-irradiated hydrogen atmospheres, finding an upper limit to the heating efficiency of $\epsilon_{\rm XUV} \approx 0.20$. My nominal value $\epsilon_{\rm XUV} = 0.3$ may therefore be an overestimate. However, given that my goal in this chapter is to explore where in the HZ HECs are possible and to map regions where the transition from gaseous to terrestrial is not possible, my present approach should suffice. Nevertheless, further studies constraining $\epsilon_{\rm XUV}$ are crucial to improving

our understanding of this parameter space.

3.4.4 Eccentricity Effects

From Figures 24–26, we see that as we consider planets with higher current eccentricity, the region where evaporated cores are possible overlaps increasingly more of the HZ, particularly in the radiation/ recombination-limited regime. This does not necessarily mean that HECs are more likely at high eccentricity—but rather that planets in circular orbits at certain a cannot be HECs, while planets at higher eccentricity can. However, that being said, there is an interesting negative feedback between tides and mass loss that may enhance the probability of HECs on eccentric orbits. A planet with initially high eccentricity will undergo fast tidal evolution (due to the strong β dependence in Equation 51), particularly if its radius is large (since $de/dt \propto R_p^5$). If a planet loses mass quickly (as is usually the case for gaseous planets with large R_p and at large eccentricity), its radius will shrink and de/dt will decrease. The earlier a planet sheds its envelope, the more likely it is to maintain a nonzero eccentricity in the long run (i.e., after ~ 5 Gyr). Since HECs typically form from such quickly evaporating planets, they are more likely to end up in eccentric orbits than their gaseous counterparts that did not lose their envelopes. There is, of course, a trade-off here in the sense that gas-free planets should have higher τ_p (and therefore faster de/dt) than gaseous planets, but the dependence on τ_p is linear and can thus be weaker than the dependence on the radius, even when τ_p varies by a few orders of magnitude. What this means is that there is an interesting link between present eccentricity and mass loss history. Translating a planet's present eccentricity into a probability that it is an evaporated core is no easy task, however, and likely requires a detailed understanding of its initial orbital state and the migration mechanism(s) it underwent. On the other hand, statistical surveys of planets found to the left of the contours in Figure 24 could uncover interesting trends.

Note that Figures 24–26 show eccentricity contours only up to $e = 0.5$. At final eccentricities higher than 0.5 today, the contour lines should move farther to the right, increasing the parameter space populated by HECs. However, since high eccentricities today in general require extremely high eccentricities in the past, I do not calculate contours above this level. I note, finally, that even though I run simulations with e_0 as high as 0.95, Figures 24–26 remain unchanged for a lower cutoff $e_0 \lesssim 0.7$.

3.4.5 Orbital Effects Due To Roche Lobe Overflow

In the present work I neglect any orbital effects due to Roche lobe overflow, which could in some cases lead to significant outward migration of the planet; see (72). I argued in section §3.2.3 that since RLO should occur during the initial stage of migration (i.e., from beyond the snow line into the HZ), modeling it was outside the scope of this work. However, a careful investigation of this initial migration process could uncover interesting couplings. For instance, gas-rich planets that initially overshoot the HZ would experience strong RLO, which could in principle cause them to migrate *back* into the HZ. These planets could have lost significantly more mass than the ones I considered here, since both RLO and XUV-driven escape would be stronger in their closer-in orbits.

A second important point concerning RLO-induced migration is that for nonzero eccentricity, (72) does not apply. In this case, it can be shown that angular momentum exchange between the gas and the planet at pericenter (where overflow is strongest) leads to a net *increase* in the eccentricity. In the limiting case that both bodies may be treated as point masses and the mass transfer rate may be approximated as a delta function at pericenter (which is appropriate for high e), Sepinsky et al. (2009) show that

$$\frac{da}{dt} = -\frac{a}{\pi}\frac{\dot{M}_p}{M_p}(1-e^2)^{1/2}\left(1-\frac{M_p}{M_\star}\right) \tag{111}$$

$$\frac{de}{dt} = -\frac{1}{\pi}\frac{\dot{M}_p}{M_p}(1-e^2)^{1/2}(1-e)\left(1-\frac{M_p}{M_\star}\right), \tag{112}$$

predicting that both a and e will tend to increase with time (recall that $\dot{M}_p$ is negative). Their relative rates of change are

$$\frac{de}{dt} = (1-e)\frac{1}{a}\frac{da}{dt}. \tag{113}$$

In other words, at intermediate values of e, the fractional rates of change in the semi-major axis and the eccentricity are comparable, and the eccentricity will increase proportionally to the semi-major axis.

As e increases, so does the atmospheric mass loss rate (via $K_{\rm ecc}$), the Roche lobe overflow rate (modulo the increase in a), and the rate of tidal evolution, leading to potentially faster mass loss

and rich couplings between these different processes.

3.4.6 Different f_H

Figures 24-26 correspond to planets with initial hydrogen fraction $f_H = 0.5$. Planets with lower initial f_H are naturally more unstable to complete loss of their envelopes—this would move all contours to the right, expanding the region where HECs are possible. However, for a different choice of initial f_H in the range $0.1 \lesssim f_H \lesssim 0.5$, the figures change very little: at constant core mass, it is only marginally harder to fully evaporate a $f_H = 0.5$ envelope than it is to evaporate a $f_H = 0.1$ envelope (see discussion below). This is due to both Roche lobe overflow, which dramatically reduces the envelope mass early on, and the fast atmospheric escape for extremely inflated planets. In Figure 21(a), for instance, the envelope fraction decreases by a factor of 10 within the first ~ 10 Myr due to energy-limited escape. The escape process is in general very fast for large f_H, and bottlenecks for $f_H \lesssim 0.1$; thus the HEC boundary is relatively insensitive to the exact choice of the initial f_H.

Along the same lines, one can ask how the choice of *final* f_H affects my results. In this chapter, I defined an evaporated core as a planet with $f'_H = 0$ and no atmosphere. However, planets with substantial gaseous envelopes that evaporate down to $f'_H \sim 0.01$ become fundamentally solid planets. At such low f'_H, additional non-thermal mass loss processes, including flaring events and interactions with the stellar wind, may significantly erode the remaining envelope. Investigating the habitability of planets with nonzero f'_H is beyond the scope of this chapter, but it is worthwhile to consider how an alternative definition of a "habitable evaporated core" affects my results. To this end, in Figure 27 I plot contour lines corresponding to three different choices of the critical f'_H below which I consider a planet to be a HEC. Solid lines correspond to $f'_H = 0$ (the default choice); dashed lines correspond to $f'_H = 0.001$; dotted lines, to $f'_H = 0.01$. Note that in the radiation/recombination-limited regime, this choice does not significantly affect evaporated cores in the HZ. However, for $f'_H = 0.01$, *all* planets undergoing energy-limited escape in the HZ could be HECs.

Earths and super-Earths with up to a few percent of their mass in H/He may still harbor liquid water oceans, and so at this point their habitability can not be ruled out. At higher f'_H, however, the surface pressures in excess of ~ 1 GPa (e.g., Choukroun and Grasset, 2007) result in

the formation of high-pressure ices, at which point the planet will likely no longer be habitable.

3.4.7 Other Atmospheric Escape Mechanisms

I modeled the XUV emission of M dwarfs as a smooth power law, but active M dwarfs are seldom so well behaved. Frequent flares and coronal mass ejections punctuate the background XUV emission and will lead to erosion of planet atmospheres beyond what I model here. During flares, the ratio of the X-ray to bolometric luminosity ($L_{\rm X}/L_{\rm BOL}$) can increase by up to two orders of magnitude (Scalo et al., 2007). Moreover, interactions with the stellar wind and nonthermal escape mechanisms such as ion pick-up and charge-exchange can lead to a few Earth-ocean equivalents of hydrogen from these planets (Kislyakova et al., 2013, 2014). However, as Kislyakova et al. (2013) demonstrate, the escape rate due to these processes generally amounts to only a few percent of the hydrodynamic escape rate. Including these non-thermal mechanisms would thus have a minor effect on my results.

In this chapter I also neglect the effects of magnetic fields. The presence of a strong planetary magnetic field may inhibit escape during flares and possibly decelerate the mass loss rate if the planetary wind is ionized. Since a strong magnetic field is likely a requirement of surface habitability around an M dwarf (Scalo et al., 2007), this point must be addressed in future work.

Finally, stars less massive than about $0.1{\rm M}_{\odot}$ may be in a "supersaturation" regime early on, saturating at XUV fluxes one or even two orders of magnitude below higher mass M dwarfs (Cook et al., 2014). This could reduce escape rates from planets around M dwarfs close to the hydrogen burning limit.

3.4.8 Multi-Planet Systems

In this chapter I restricted my calculations to single-planet systems. However, many of the super-Earths and mini-Neptunes recently detected by Kepler reside in multi-planet systems, such as Kepler-32 (Fabrycky et al., 2012), Kepler-62 (Borucki et al., 2013), and Kepler-186 (Quintana et al., 2014). Moreover, Swift et al. (2013) demonstrated how the five-planet system Kepler-32 could be representative of the entire ensemble of Kepler M dwarf systems, arguing that multiplicity could be the rule rather than the exception for these stars.

While tidal processes still operate in systems with multiple planets, the orbital evolution of

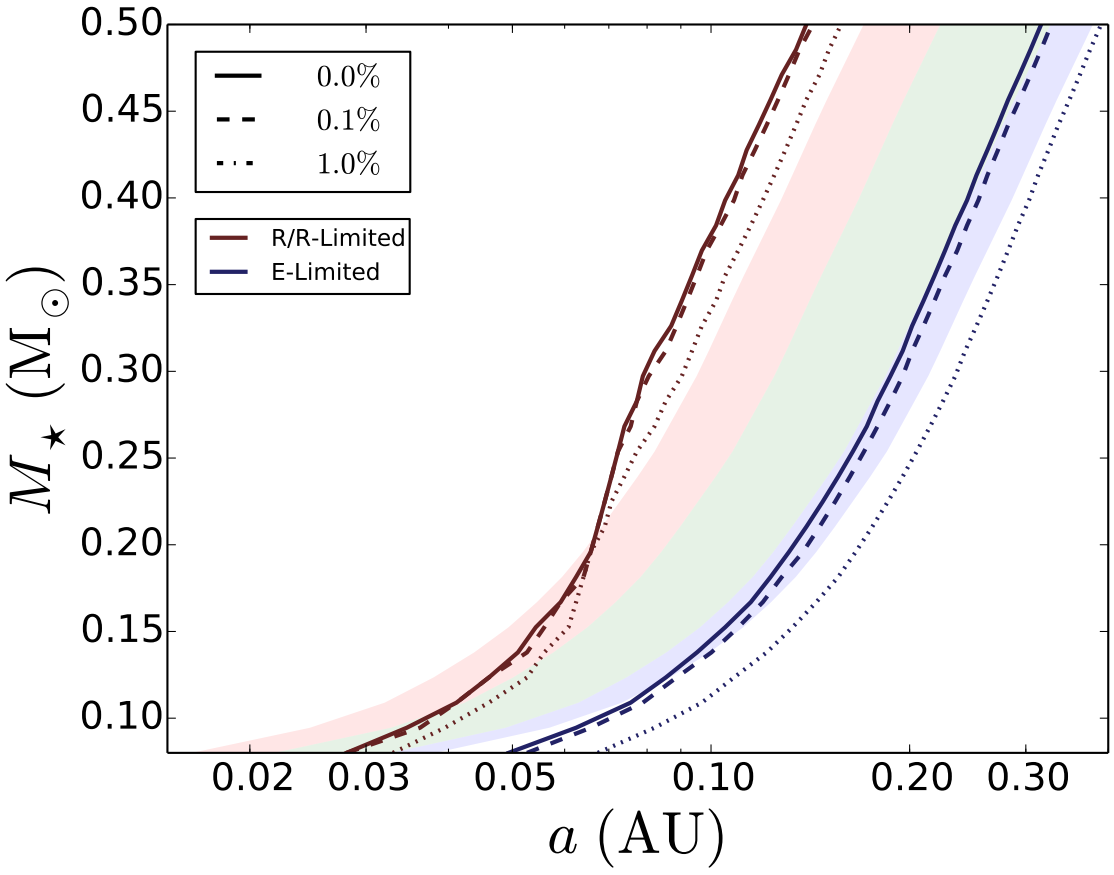

Figure 27: Similar to Figure 24, but here contours correspond to different choices of final f_H for a $1\,M_\oplus$ core on a circular orbit. The solid lines correspond to $f'_H = 0$ and are the same as the $e = 0$ contours in Figure 24. Dashed lines correspond to the transition between planets that have less than (left) and more than (right) 0.1% H/He ($f'_H = 0.001$) at 5 Gyr. Dotted lines correspond to the 1% H/He ($f'_H = 0.01$) transition. While the $f'_H = 0.001$ contours are barely distinguishable from the $f'_H = 0$ contours, at final $f'_H = 0.01$ the HEC parameter space is significantly larger. However, it is unclear whether planets with $f'_H = 0.01$ could be habitable.

individual planets will be much more complex. In general, though tidal dissipation may still lead to a decrease in the semi-major axis, the eccentricity evolution will be governed by secular interactions between the planets. The coupling between mass loss and tidal evolution I investigated in §3.3.3 would likely be weaker, particularly for closely-packed coplanar systems where the eccentricities are necessarily low. However, planet-planet interactions could allow for even richer couplings to the atmospheric evolution of planets in the HZ.

Consider, for instance, the case of a terrestrial planet in the HZ and a mini-Neptune interior to the HZ. Both mass loss and tidal dissipation scale inversely with the semi-major axis, so both effects could strongly shape the fate of the inner planet, which could in turn affect both the orbital and atmospheric evolution of the HZ planet in complex ways. The same could happen for a planet overflowing its Roche lobe interior to the HZ, which as I argued in §3.4.5 could experience large changes in a and e.

Similarly, consider a scenario in which a planet in the HZ loses mass at the same time as it is perturbed by a more massive companion exterior to the HZ. As the planet's mass decreases, it will undergo larger swings in eccentricity, which in turn lead to faster mass loss and stronger orbital evolution.

Another interesting scenario involves planets close to mean motion resonances. For a system of two planets in resonance, Lithwick and Wu (2012) and Batygin and Morbidelli (2013) showed that the presence of a dissipative mechanism such as tidal evolution naturally leads to the repulsion of the two planets' orbits: the inner planet's orbit decays, while the outer planet gets pushed outward. In the hypothetical case of two mini-Neptunes in resonant orbits interior to the HZ, this mechanism could result in the migration of the outer planet *into* the HZ, having lost significantly more mass (due to its initially smaller orbit) than if it had originated exterior to the HZ and migrated inwards. A chain of resonant planets, which is a potential outcome of disk-driven migration (Terquem and Papaloizou, 2007; Ogihara and Ida, 2009), could have similarly interesting consequences on planets in the HZ.

All of these scenarios, which involve coupling between atmospheric mass loss, tidal evolution, and secular planet-planet interactions, probably occur even for planets outside the HZ. Modeling this coupling could be critical to understanding systems like Kepler-36 (Carter et al., 2012), where two planets with semi-major axes differing by $\sim 10\%$ have a density ratio close to 8, which Lopez

and Fortney (2013) show could be the result of starkly different mass loss histories. Future work should investigate how mass loss modifies the orbital interactions in multiplanet systems.

3.4.9 Other Caveats

3.4.9.1 The Habitable Zone

Recently, Yang et al. (2013) argued that cloud feedback on tidally locked planets can greatly increase the planetary albedo and move the IHZ in by a substantial amount. Similarly, Abe et al. (2011) showed that planets with limited surface water are stable against runaway greenhouses at greater insolation than Earth, which also decreases the IHZ distance. Since evaporated cores are more likely at smaller a, a closer-in inner edge to the HZ could greatly increase the parameter space available to these planets. However, as I argued in §3.4.2, HECs are likely to have abundant surface water.

In the previous chapter, I showed that terrestrial planets in the HZ of M dwarfs may experience long runaway greenhouses during their host stars' extended pre-main sequence phases, during which time water loss to space can lead to their complete desiccation, rendering them uninhabitable. HECs are naturally more robust against severe water loss, given that their initially dense H/He envelopes can shield the surface and lower atmosphere from XUV radiation. Once the envelope is lost, water loss from the surface could occur, but by that point the stellar XUV flux will be much lower. Future work will address the fate of the water on HECs in detail.

3.4.9.2 Thermal Evolution

As mentioned before, I do not consider how the mass loss rate affects the planet radius, but instead assume the radius instantaneously returns to the "equilibrium" value for the given age, mass, and hydrogen fraction. In reality, the radius is likely to remain inflated for some time, particularly for fast mass loss (Lopez et al., 2012). Because the radii decrease too quickly, the mass loss feedback on the tidal evolution is always negative: mass loss always acts to dampen the tidal evolution. However, accurate modeling of the radius could enable a positive feedback, in which the effect of the decreasing mass (higher da/dt) overpowers the effect of the decreasing radius (lower da/dt) in (36) and (50), leading to faster orbital decay.

Currently, the planet radii in my model are also independent of both the insolation and the

degree of tidal heating. Under the extremely high fluxes and strong tidal forces at early times, planets in the HZs of M dwarfs are likely to be significantly more inflated than modeled here, resulting in faster atmospheric mass loss and likely a higher probability of complete envelope loss. I also do not model radiogenic heating, which could further add to the inflation of the planet. In this sense, my results are conservative, and HECs may in fact be possible at larger distances from their parent stars. A self-consistent thermal evolution model must be developed to accurately address this point.

I also emphasize that the quasi-static approximation (§3.2.4.1) may not be valid for the most inflated planets on eccentric orbits above $e \approx 0.4$. Given the large uncertainty concerning the orbital migration due to RLO for planets on eccentric orbits, I urge caution in interpreting my results quantitatively for highly eccentric planets.

3.4.9.3 Tidal Evolution

Given that the CPL tidal model is accurate only to second order in eccentricity, most of my integrations were performed under the CTL framework. At low e, however, these models predict qualitatively similar results for both the orbital migration and the coupling to the mass loss processes. At high e, on the other hand, tidal models lack observational validation, given the relatively low eccentricities of the major bodies in our solar system. I again urge caution in interpreting quantitative results for $e \gtrsim 0.5$.

Recall that my values of Q_p and τ_p are fixed throughout the integrations and change discontinuously to the more dissipative rocky values as soon as the envelope is lost; this is certainly an oversimplification. Although I defined Q via its relationship to the phase lag in Equation (34), it may also be defined as the *specific dissipation function* (Goldreich and Soter, 1966), defined as

$$Q^{-1} \equiv \frac{1}{2\pi E_0} \int \left(-\frac{dE}{dt} \right) dt, \tag{114}$$

where E_0 is the maximum tidal energy stored in the system and the integral is the energy dissipated over one complete cycle. Q is therefore inversely proportional to the dissipation per unit energy stored. Planets with large values of f_H, where the dissipation throughout most of the body is small, should therefore have large values of Q_p; rocky planets with small f_H, conversely, have large specific dissipation rates and correspondingly low values of Q_p, in agreement with the remarks in

§1.4.1. Therefore, one would expect that as planets lose mass, Q_p should decrease, corresponding to an increase in τ_p. Modeling $Q_p(f_H)$ and $\tau_p(f_H)$, however, is outside the scope of this chapter, and will be considered in future work. I simply note that as planets lose mass, the decreasing value of Q_p/increasing value of τ_p should offset the decreasing radii, leading to stronger tidal evolution and stronger couplings than reported here.

One must likewise be careful in choosing the planet radius that goes into calculating the tidal evolution. This radius should be the effective radius of the dissipating material, and is thus dependent on the chosen value of Q_p and τ_p; for high Q_p/low τ_p, the radius should probably be that of the envelope. I took this to be R_p, the 20 mbar radius shown in the tracks in Figure 17; other choices of the tidal radius will lead to different evolutions.

I only consider planets with zero initial obliquity. For nonzero obliquity, the tidal evolution equations must be modified and will lead to differences in the evolution; see Heller et al. (2011) and Barnes et al. (2013). I also ignore the spin-up of the planet as it thermally contracts over time, since it should return to the equilibrium spin almost instantaneously. The excess angular momentum would most likely be absorbed into the orbit, but it is a small enough fraction of the orbital angular momentum that it can be safely neglected.

Finally, I note that throughout this chapter I set $t_{\rm stop} = 5$ Gyr, based on the age of the solar system. For systems older than 5 Gyr, my results should change little, since both the tidal and atmospheric evolution will have strongly tapered off by this time. Younger systems, however, may still be in the throes of tidal decay and mass loss processes, and this must be kept in mind when searching for HECs.

3.5 Conclusions

I have shown that under certain conditions, gas-rich super-Earths and mini-Neptunes that migrate early ($t \lesssim 10$ Myr) into the HZs of M dwarfs can naturally shed their gaseous envelopes and form gas-free Earth-mass planets. Together, Roche lobe overflow and hydrodynamic escape can remove up to a few Earth masses of hydrogen and helium from these planets, potentially turning them into "habitable evaporated cores" (HECs). This process is most likely for planets with solid cores with masses lower than $1\,M_\oplus$ and up to about 50% H/He by mass, and can occur around all M dwarfs, particularly close to the inner edge fo the HZ. HECs are less likely to form around K and G

dwarfs because of these stars' shorter super-luminous pre-main sequence phases and shorter XUV saturation timescales. Furthermore, I find that HECs cannot form from planets with core masses greater than about $1\,M_\oplus$ and more than a few percent H/He by mass; thus, massive terrestrial super-Earths currently in the HZs of M dwarfs have probably always been terrestrial. My results are thus similar to those of Lammer et al. (2014), who showed that planets more massive than $\sim 1.5\ M_\oplus$ typically cannot lose their accreted nebular gas in the HZs of solar-type stars.

Whether or not a given gas-rich planet forms a HEC is highly dependent on the early XUV evolution of the host star. In particular, a long XUV saturation timescale ($t_{\rm sat} \gtrsim 1$ Gyr) is needed to fully evaporate the envelopes of gas-rich planets in the HZs of early and mid M dwarfs. While a large $t_{\rm sat}$ is consistent with the long activity timescales (West et al., 2008) and long spin-down times (Pizzolato et al., 2000) of M dwarfs, more observations are needed to pin down $t_{\rm sat}$ for these stars. Moreover, the relative strength of the X-ray and EUV luminosity early on also affects whether or not HECs can form, as this determines whether the atmospheric escape is energy-limited or radiation/recombination-limited. In the energy-limited regime, which occurs for an X-ray dominated flow, the escape is fast and HECs can form throughout most of the HZ of all M dwarfs. In the radiation/recombination-limited regime, which applies to an EUV dominated flow, HECs only form close to the inner edge of low mass M dwarfs.

I further find that HECs can form from planets on circular as well as eccentric orbits, though they are marginally more likely to have higher e in the long run. While there exist feedbacks between atmospheric mass loss and tidal evolution, I find that these are only significant at $e \gtrsim 0.5$; in these cases, whether or not a HEC forms can depend just as strongly on the tidal properties of the planet as on the efficiency of the atmospheric escape.

Many of the Earth or sub-Earth-mass terrestrial planets detected in the HZs of M dwarfs in the coming years could be habitable evaporated cores. These planets should have abundant surface water and are likely to be water worlds, whose potential for habitability should be investigated further. Given the extreme water loss that terrestrial planets undergo in the HZs of M dwarfs discussed in the previous chapter, the evaporation of gaseous envelopes from migrated planets could be a promising pathway to the formation of habitable planets around these stars.

Part II:

The Detection of Planets in the Habitable Zones of M Dwarfs via the Transit Method

4 Overview

Launched in 2009, the *Kepler* spacecraft (Borucki et al., 2010) has to date led to the discovery of nearly 5,000 extrasolar planet candidates and to a revolution in several fields of astronomy including but not limited to exoplanet science, eclipsing binary characterization, asteroseismology, stellar variability and supernovae studies. The most precise photometer ever launched, *Kepler* stared at $\sim$ 150,000 stars in a small patch of the sky in the constellation Cygnus over the course of four years, collecting the flux of each one at either a 1-minute or 30-minute cadence. *Kepler* thus provided light curves with unprecendented photometric precision and time baseline, enabling the detection of the sub-part per thousand signals of exoplanets transiting their host stars.

Exoplanet transits occur for exoplanets whose orbital planes lie along or close to the line of sight to Earth. As the planet crosses the face of its star (as seen from Earth), it blocks an amount of light roughly equal to the ratio of its projected area to that of the star's. While a transiting Jupiter-sized planet blocks close to 1% of its star's light, the signal of a transiting Earth-sized planet can be less than one-hundredth that value, or 100 parts per million (ppm). Nonetheless, the exquisite precision of *Kepler* permits the detection of these events, seen as dips in the light curve at periodic or quasi-periodic intervals.

Only a small fraction of exoplanets can ever been seen to transit, as the random orientation on the sky of any given planetary system guarantees that most planets never cross the face of their host star as seen from Earth (Borucki and Summers, 1984). While the geometric probability of transit depends on the size of the star and the period of orbit, on average, about one percent of all exoplanets can be expected to transit. Nevertheless, the power in numbers ensured *Kepler* would detect many exoplanets, provided these were in fact common. One of the most surprising results of the mission, in fact, was the vast number of exoplanets and exoplanet candidates discovered. Given a one percent average transit probability, the nearly 5,000 planets detected from surveying 150,000 stars suggest, at face value, that *every star in the sky hosts at least one planet.* While our simple analysis ignores the dependence of the transit probability on orbital period (shorter-period planets are easier to detect) and the large number of multi-planet systems in *Kepler*, the large number of detections are definitive evidence that exoplanets are extremely common around Sun-like stars.

While exoplanets of all kinds are interesting in their own right, one of the primary goals of

the *Kepler* mission was to constrain the quantity known as $\eta_\oplus$ ("eta-sub-Earth"), loosely defined as the frequency of Earth-size planets around Sun-like stars. Although many studies have been able to constrain this frequency (Petigura et al., 2013, e.g.,), uncertainties are still extremely large (Foreman-Mackey et al., 2014) and values vary between a couple percent and twenty percent. In particular, very few small planets have been detected in the habitable zones of solar-type stars, *none* of which are rigorously "Earth-size" (the vast majority are super-Earths and/or mini-Neptunes). As of October 2016, there are only 20 candidates with radii less than $2R_\oplus$ within the conservative HZ (Kane et al., 2016); see Figure 28. However, this relative paucity does not imply infrequency: planets in the HZ of solar-type stars are less likely to transit, and if they do, they transit less frequently. A true Earth analogue would have been seen to transit only four times (or fewer) in the *Kepler* dataset; with so few transits, stacking to increase the signal-to-noise of the detection is not as effective. Moreover, Earth-size planets are small and Sun-like stars can be quite noisy at the sub-100 ppm level, making the detection of these planets quite challenging. Estimates of $\eta_\oplus$ therefore rely on noisy histograms and extrapolation from shorter periods and larger radii.

Many of the difficulties in detecting and studying planets in the habitable zone can be overcome by turning to lower mass stars. As I discussed in previous chapters, M dwarfs offer the best opportunities for the detection of habitable zone planets via the transit method due to the smaller planet/star radius ratio, which increases the signal of a transiting planet and the shorter periods of planets in the HZ, which allows more transits to be observed in a given time window. With this in mind, many studies (e.g., Kopparapu, 2013; Dressing and Charbonneau, 2015) used the *Kepler* M dwarf sample to constrain the frequency of Earth-size planets in the HZs of those stars. These studies found much higher occurence rates for these planets relative to the solar-type sample, ranging from 20 to upwards of 50%, suggesting nearly half of all M dwarfs host terrestrial-size planets in the HZ.

This realization came at a fortuitous point in time. At the end of its primary four-year mission, *Kepler* suffered a large setback with the failure of the second of its four reaction wheels in May 2013. These wheels are what allow the spacecraft to maintain its fine sub-milliarcsecond pointing and stare at a fixed portion of the sky; with only two of them operational, motion along the roll axis of the spacecraft could no longer be controlled at the level required for precise photometry, and the primary mission came to an end. Fortunately, the *Kepler* team engineers realized that by

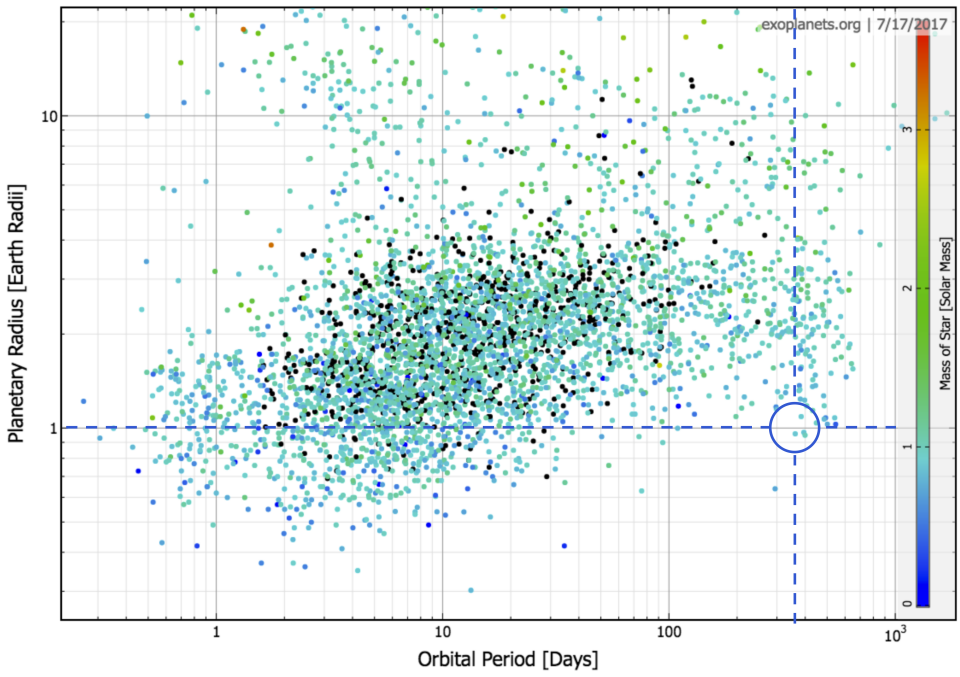

Figure 28: All *Kepler* planets and planet candidates discovered to date, plotted according to their radii and orbital periods, using the tools at `exoplanets.org` (Han et al., 2014). The crosshairs mark the approximate position of Earth on this plot. Although *Kepler*'s primary goal was to determine the frequency of Earth-like planets around Sun-like stars, very few such planets were detected. This is due to observational biases: *Kepler* is far more sensitive to larger, shorter-period planets, as can be seen from this plot. While the few detections in the HZ can be used to infer an estimate for $\eta_\oplus$, a better observing strategy is to target M dwarfs in search of habitable planets.

aligning the spacecraft along the plane of the ecliptic, solar radiation pressure could be used to stabilize the roll motion for short periods of time. The remaining drift could be counteracted by the periodic firing of the spacecraft's thrusters to achieve pointing precision within one or two orders of magnitude that of the original mission (Howell et al., 2014). On the software side, (Vanderburg and Johnson, 2014) showed how the instrumental noise introduced by the drift could be largely removed, recovering photometric precision close to that of the original mission for the brightest stars. This new mode of operation, known as *K2*, required a completely different observing strategy. Since the spacecraft must now point along the ecliptic, its field of view changes every three months in order to keep the boresight of the telescope perpendicular to the solar radiation pressure vector. The *Kepler* team thus organized a Guest Observer (GO) program, a community-led effort to select new targets for observation every quarter. Given the recent estimates of planet occurrence rates around M dwarfs, the community quickly realized that *K2* provided the perfect opportunity to study these planets *en masse*.

However, even with the hardware and software solutions described above, *K2* light curves are still significantly noisier than those of the original *Kepler* mission, complicating the detection of small planets. In this part of my thesis, I discuss techniques I have developed to improve on the de-trending of *K2* light curves and permit the detection of small planets, with a focus on those residing in the habitable zone of M dwarfs. In the following chapters, I will introduce `EVEREST`, my pipeline for removing systematics from *K2* light curves. `EVEREST`

`EVEREST` has already led to the detection of several new *K2* planets (e.g., Mann et al., 2017; Barragán et al., 2017; Fridlund et al., 2017; Jones et al., 2017). In collaboration with Ethan Kruse (Kruse et al., 2017, in prep.), we have discovered several hundred previously unknown, small *K2* planets, in great part due to the enhanced photometric precision of `EVEREST` light curves. Figure 29 shows one such planet. In Figure 30, I plot all 800+ planets detected with `EVEREST` in *K2* campaigns 0–8, approximately 300 of which are new. The green square shows the "optimistic" habitable zone: of the seven planetary systems shown in that region, four are new detections of small planets in the HZ of M dwarfs.

The first two chapters in this part discuss the development of the pipeline and its implementation. The final chapter applies the pipeline to TRAPPIST-1, an ultracool dwarf hosting seven terrestrial-size planets, three of which are in the HZ. I will show how `EVEREST` enabled the confir-

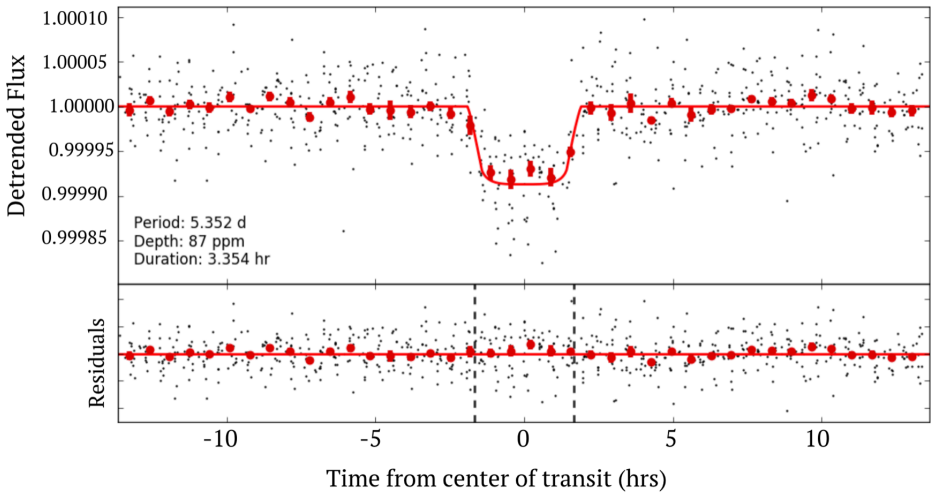

Figure 29: A small 5d-period planet candidate discovered in *K2* using `EVEREST` (Kruse et al., 2017, in prep.). `EVEREST` light curves have higher photometric precision than those produced by other pipelines, enabling the detection of smaller transiting exoplanets.

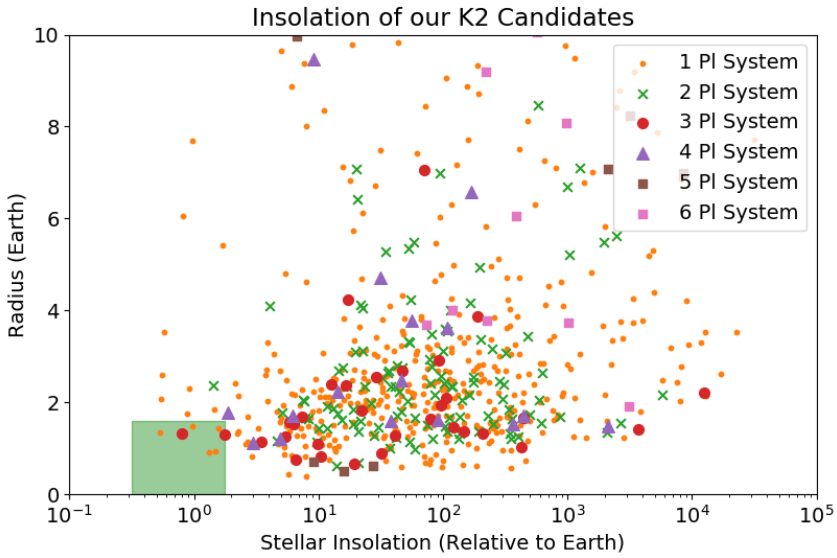

Figure 30: Radius versus irradiation of the *K2* planet candidates discovered using `EVEREST` (Kruse et al., 2017, in prep.). Most of the planets plotted here have been previously reported in the literature, but ∼300 are new discoveries. The green box at the lower left is the "optimistic" habitable zone (Chapter 1), hosting 7 planets, two of which are in 3-planet systems and four of which were previously undetected.

mation of the outermost planet and provided a glimpse into the complex resonant structure of the system.

5 EVEREST 1.0

Portions of this chapter were originally published in collaboration with E. Agol, E. Kruse, R. Barnes, A. Becker, D. Foreman-Mackey, and D. Deming in the October 2016 edition of the Astronomical Journal (Luger et al., 2016, AJ, Vol. 152, 100; © 2016 American Astronomical Society), and are reproduced below with permission of the American Astronomical Society.

5.1 Introduction

As I discussed above, because of its reduced pointing accuracy, *K2* raw aperture photometry is between 3 and 4 times less precise than that of the original *Kepler* mission. The data displays strong instrumental artefacts with different timescales, including a ~ 6 hour trend, which severely compromise its ability to detect small transits. Recently, several authors have developed powerful methods to correct for the systematics in *K2* data, often coming within a factor of ~ 2 of the original *Kepler* precision. In particular, the `K2SFF` pipeline (Vanderburg and Johnson, 2014) decorrelates *K2* aperture photometry with the centroid position of the stellar images. Centroids are determined based either on the center-of-light or via a Gaussian fit to the stellar PSF. The motion of the centroids is then fit with a polynomial and transformed into a single parameter that relates spacecraft motion to flux variations, which is then used to de-trend the data. Similar methods are employed in the `K2VARCAT` pipeline (Armstrong et al., 2015), developed specifically for variable *K2* stars, the `K2P`^{2} pipeline (Lund et al., 2015), which uses an intelligent clustering algorithm to define custom apertures, and in the pipeline of Huang et al. (2015), which employs astrometric solutions to the motion of *K2* targets, determining the X and Y motion of a target from the behavior of multiple stars on the same spacecraft module. Finally, the `K2SC` pipeline (Aigrain et al., 2015, 2016) and the pipeline of Crossfield et al. (2015) both employ a Gaussian process (GP) to remove instrumental noise, using the X and Y coordinates of the target star as the regressors to derive a model for the instrumental systematics. The nonparametric nature of the GP results in a flexible model with increased de-trending power especially for dim *K2* targets.

In one way or another, all of these techniques rely on numerical methods to identify and remove correlations between the stellar position and the intensity fluctuations. Even when a nonparametric technique such as a GP is used, assumptions are still made about the nature of the correlations

between spacecraft motion and instrumental variability. Moreover, the process of determining the stellar centroids is prone to uncertainties and relies on assumptions about the shape of the stellar PSF.

A powerful alternative to these methods is pixel level decorrelation (PLD), a method developed by Deming et al. (2015) to correct for systematics in *Spitzer* observations of transiting hot Jupiters. The tenet of PLD is that the best way to correct for noise introduced by the motion of the stellar image does not involve actual measurements of the position of the star. The centroid of the stellar image is, after all, a secondary data product of photometry, and is subject to additional uncertainty. PLD skips these two numerical steps (i.e., fitting for the stellar position and solving for the correlations) by operating on the *primary* data products of photometry, the intensities in each of the detector pixels. These intensities are normalized by the total flux in the chosen aperture then used as basis vectors for a linear least-squares (LLS) fit to the aperture-summed flux. Since astrophysical signals (stellar variability, planet transits, stellar eclipses, etc.) are present in all of the pixels in the aperture, the normalization step removes astrophysical information from the basis set, ensuring that PLD is sensitive only to the signals that are *different* across the aperture. PLD is therefore an "agnostic" method of performing robust flat-fielding corrections with minimal assumptions about either the nature of the intra-pixel variability or the correlation between spacecraft jitter and intensity fluctuations. I note that my method is similar to that of Foreman-Mackey et al. (2015) and Montet et al. (2015), who use the principal components of the variability among the full set of *K2* campaign 1 light curves as "eigen light curve" regressors. However, rather than deriving my basis vectors from other stars in the field, whose light curves contain undesired astrophysical signals, my basis vectors are derived solely from the pixels of the star under consideration.

In this chapter I build on the PLD method of Deming et al. (2015), extending it to higher order in the pixel fluxes and performing principal component analysis (PCA) on the basis vectors to limit the flexibility of the model and thus prevent overfitting. I further couple PLD with a GP to disentangle astrophysical and instrumental variability.

I apply my pipeline, `EVEREST` (*EPIC Variability Extraction and Removal for Exoplanet Science Targets*), to the entire set of *K2* light curves from campaigns 0-7 and generate a publicly-available database of processed light curves. My code is open-source and is available online at `github.com/`

`rodluger/everest`.

This chapter is organized as follows: in §5.2 I review the basics of PLD and derive my third-order PLD/PCA/GP model, and in §5.3 I describe my pipeline in detail. I present my results and a general discussion in §5.4.

5.2 Pixel Level Decorrelation

5.2.1 First Order PLD

The linear PLD model developed in Deming et al. (2015) is given by the expression

$$m_i = \sum_l a_l \frac{p_{il}}{\sum_k p_{ik}} + \alpha + \beta t_i + \gamma t_i^2 \tag{115}$$

where m_i denotes the noise model at time t_i, p_{il} denotes the flux in the l^{th} pixel at time t_i, and a_l is the linear PLD coefficient for the l^{th} pixel. Both sums are taken over all pixels in the aperture; the last three terms are a polynomial in time used to capture temporal variations in the flux baseline due to intrinsic variability of the star.

The coefficients are obtained by minimizing the sum of the squares of the difference between the model m_i and the simple aperture photometry (SAP) flux, y_i:

$$\frac{\partial \chi^2}{\partial a_l} = 0 \tag{116}$$

where

$$\chi^2 = \sum_i \frac{(y_i - m_i)^2}{\sigma_i^2}. \tag{117}$$

In the expression above, σ_i are the standard errors of the flux and

$$y_i = \sum_k p_{ik}. \tag{118}$$

Framed in this manner, computation of the PLD model is a linear regression problem; the coefficients are readily found by simultaneously solving Equation (116) for all coefficients a_l, as

well as α, β, and γ.

5.2.2 Higher Order PLD

Deming et al. (2015) found that the pointing jitter of the *Spitzer* telescope was sufficiently small that extending PLD to higher order in the pixel fluxes was unnecessary. However, because of the large pointing variations of the *K2* spacecraft on short timescales, I find it necessary to extend PLD to higher order. Keeping terms up to third order in the pixel fluxes, we may express this model as

$$\begin{aligned} m_i = &\sum_l a_l \frac{p_{il}}{\sum\limits_k p_{ik}} + \\ &\sum_l \sum_m b_{lm} \frac{p_{il} p_{im}}{(\sum\limits_k p_{ik})^2} + \\ &\sum_l \sum_m \sum_n c_{lmn} \frac{p_{il} p_{im} p_{in}}{(\sum\limits_k p_{ik})^3} + \alpha + \beta t_i + \gamma t_i^2 \end{aligned} \tag{119}$$

where once again index i denotes time and all other indices denote the pixel number. The PLD coefficients are now a_l (one per pixel), b_{lm} (one per pixel pair), and c_{lmn} (one per group of three pixels). Despite the added complexity, the model remains linear and may be solved in a similar fashion as before.

In Figure 31, I illustrate the de-trending technique for EPIC 201367065 (K2-3), a star host to three known transiting planets (Crossfield et al., 2015). The top panel shows the normalized raw aperture-summed flux after subtracting off the background; note the large systematics at very short ($\sim$ 6 hr) timescales. The next panel shows the flux de-trended with first order PLD (Equation 115). The scatter is reduced by a factor of 2.9 and the seven transits become visible by eye. The following two panels show the results of second and third order PLD, which improve the scatter over the raw data by factors of 4.7 and 5.2, respectively. Note, importantly, that even data collected during thruster fire events (the outlier points seen above the continuum every $\sim$ 6 hr in the top two plots), is properly corrected by higher order PLD.

One might wonder why not go to even higher order in the pixel coefficients. While possible in principle, the number of regressors increases steeply with the PLD order. For a typical *K2* star, this number is around 5×10^3 for third order PLD, and would increase to $\sim 5 \times 10^4$ for fourth

order and $\sim 3 \times 10^5$ for fifth order PLD. While such a large number of regressors can become computationally expensive, the most serious drawback of higher order PLD is that it can become so flexible as to lead to overfitting. As I discuss in §5.3.3 below, PLD can sometimes overstep and remove part of the white noise component of the light curve. This is particularly problematic when the number of regressors becomes very large, in which case PLD can lead to artificially low scatter in the de-trended light curve, washing out astrophysical signals (including transits) in the process.

To avoid these issues, in Figure 31 I computed the PLD basis vectors from the 10 brightest pixels in the aperture. In practice, however, I obtain better results by instead performing principal component analysis (PCA) on the first, second, and third order fractional pixel fluxes (the terms in Equation 119). This yields a set of N basis vectors $\mathbf{x}$ that describe most of the instrumental variability in the light curve. I use these to construct my *design matrix*

$$\mathbf{X} = \begin{pmatrix} x_{0,0} & x_{0,1} & ... & x_{0,N} & 1 \\ x_{1,0} & x_{1,1} & ... & x_{1,N} & 1 \\ ... & ... & ... & ... & ... \\ x_{M,0} & x_{M,1} & ... & x_{M,N} & 1 \end{pmatrix} \tag{120}$$

where M denotes the total number of data points along the time dimension. I choose the number of principal components N that maximize the de-trending power while preventing overfitting. I explain my method in §5.3.3 below; typically, $N < 200$.

Since the problem is still linear, my model is simply

$$\mathbf{m} = \mathbf{X} \cdot \mathbf{c} \tag{121}$$

where $\mathbf{c}$ is the vector of coefficients, one for each basis vector in $\mathbf{X}$. Their values are given by solving the generalized least-squares (GLS) problem

$$\mathbf{c} = \left(\mathbf{X}^\top \cdot \mathbf{K}^{-1} \cdot \mathbf{X}\right)^{-1} \cdot \left(\mathbf{X}^\top \cdot \mathbf{K}^{-1} \cdot \mathbf{y}\right) \tag{122}$$

where $\mathbf{K}$ is the covariance matrix of the data and $\mathbf{y}$ is the SAP flux given by Equation (118). Note that for a diagonal covariance $\mathbf{K}_{ij} = \delta_{ij}\sigma_i^2$, Equation (122) is mathematically equivalent to Equations (116) and (117) from before.

5.2.3 Gaussian Process Regression

In Equation (120) I purposefully neglected the temporal polynomial terms. In principle, modeling stellar variability should not be necessary if all we wish to remove is the instrumental noise. By virtue of using the *fractional* pixel fluxes as basis vectors (from which the astrophysical signal has been removed by the normalization), PLD should fit out instrumental variability only, obviating the need for an extra temporal term. However, in practice, this is not always the case. To illustrate this, in Figure 32 I plot a quarter of data of KOI 1275, a star with one planet candidate discovered by the original *Kepler* mission (Borucki et al., 2011). I chose this star over one from the *K2* mission simply because it is easier to discriminate between instrumental and astrophysical signals by eye for this light curve, though my arguments apply equally well to *K2*.

Figure 32**a** shows the raw simple aperture photometry (SAP) data in black (top panel); in red I plot a simple first-order PLD fit with no polynomial term (obtained from the solution to Equation 116). Below it are the residuals of the PLD fit (second panel, black) and the fully de-trended data after smoothing with a GP to remove stellar noise (third panel). The bottom panel shows the de-trended data folded on the period of the planet candidate, with the 1-hr median shown in red. Note that the PLD fit is quite good, as it removes the low-frequency arc as well as the flux jump at $t \sim 375$ d and the thermal features at $t \sim 385$ d and $t \sim 400$ d. Since the star isn't significantly variable, the temporal term is not necessary.

However, this is not the case when the star is variable. In the next three columns I multiply the pixel level light curve by a sinusoid with a period of 25 days to simulate stellar variability. I choose an amplitude comparable to the amplitude of the instrumental signal, and de-trend the light curve three different ways: PLD only (**b**), PLD plus a tenth order polynomial (**c**) and PLD plus a GP (**d**).

The PLD-only fit (Figure 32**b**) is quite poor. The general shape of the PLD curve is mostly preserved and, as expected, the stellar signal is still present in the residuals, but the fit increases the white noise by a factor of ~ 4, all but washing out the transits (bottom panel). This happens

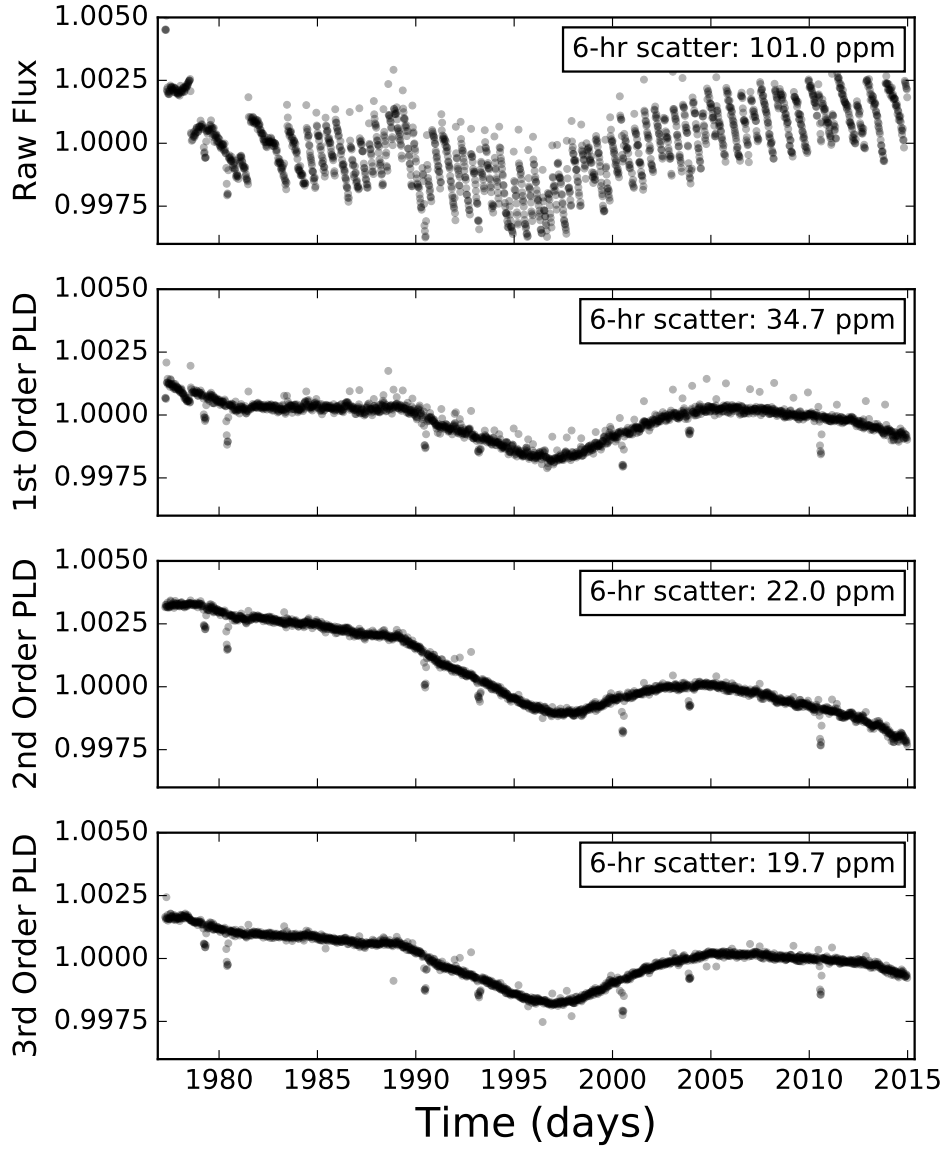

Figure 31: PLD applied to a portion of the data for EPIC 201367065 (K2-3). The top panel is the background-subtracted, normalized SAP flux in a large 35-pixel aperture centered on the target. The bottom three panels show the normalized PLD-de-trended flux for 1st, 2nd, and 3rd order PLD, respectively, using only the 10 brightest pixels. PLD increases the 6-hr photometric precision by factors of 2.9 (1st order), 4.7 (2nd order), and 5.2 (3rd order).

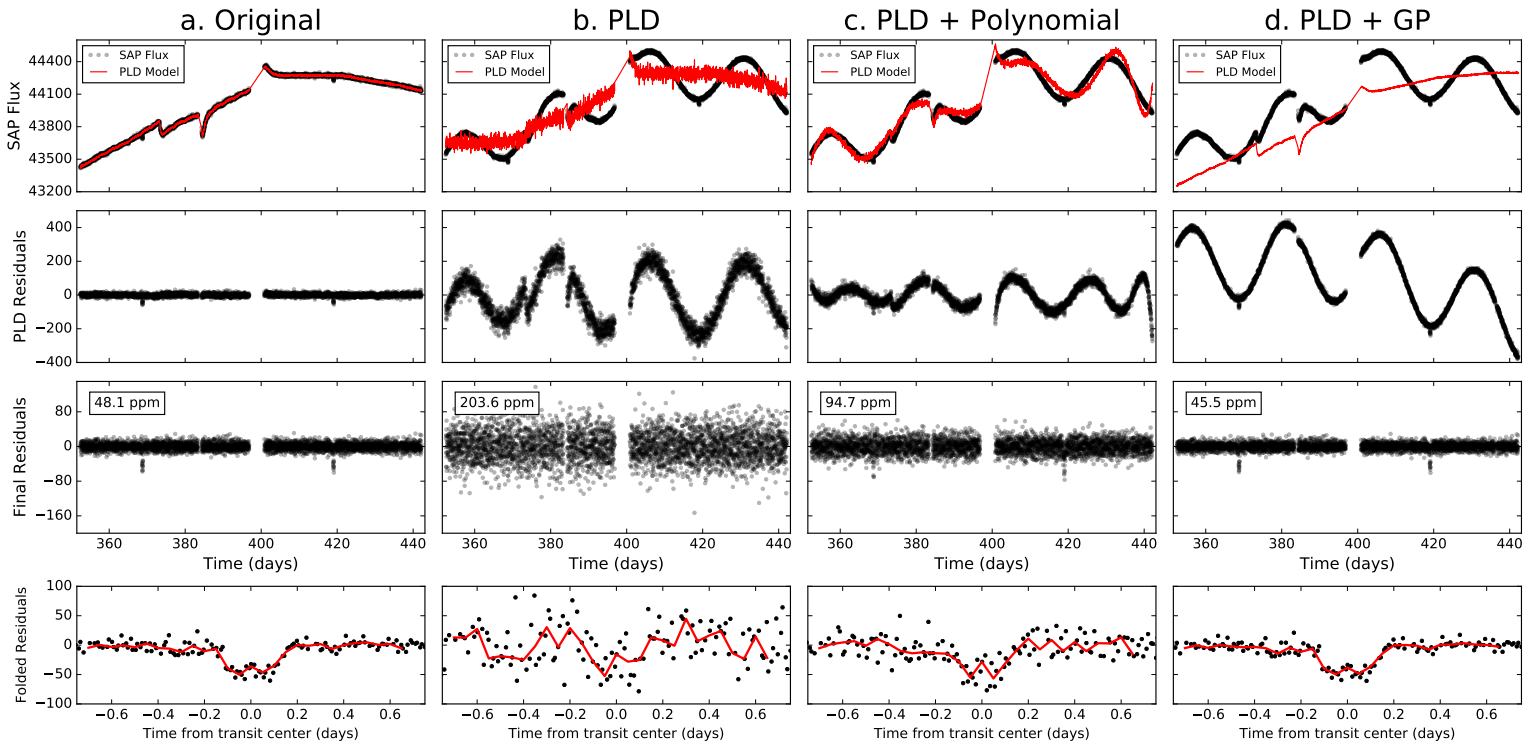

Figure 32: Different de-trending techniques for quarter 4 of KIC 8583696 (KOI 1275), a planet candidate host from the original *Kepler* mission. The original data are shown in the left column; in the other columns I artificially injected a sinusoidal signal with a period of 25 days and an amplitude comparable to that of the instrumental variability. The top row shows the raw SAP data (black) and the first order PLD model (red); the residuals of the fit are indicated directly below. The third row shows the final residuals after smoothing with a GP to eliminate low-frequency stellar variability. Finally, the bottom row shows these residuals folded on the orbital period of the planet candidate (black), with the 1-hr median indicated in red. Combining PLD with a GP ensures PLD fits out only the instrumental variability without inflating the white noise.

because of a serious limitation of the LLS problem framed in Equation 116, which yields the PLD coefficients that minimize the scatter of the PLD term m_i about the flux y_i. If m_i is not a suitable approximation to the flux, as in the case of a highly variable star, the χ^2 of the fit will necessarily be large. It is not hard to show that χ^2 can be substantially reduced by inflating the white noise component of m_i, leading to the large scatter seen in the de-trended data. Absent a model for the astrophysical variability, PLD naturally exchanges correlated noise for white noise, which severely compromises its de-trending power.

A straightforward way to improve the quality of the fit is to include the polynomial term to capture the stellar variability, as in Deming et al. (2015). I do this in Figure 32**c**, where I use a tenth order polynomial. While the fit is significantly improved, PLD still inflates the white noise component, since the polynomial is not flexible enough to capture all of the stellar signal. Though the transit is visible in the bottom plot, the quality of the de-trended data is significantly degraded when compared to that of the non-variable star.

One might imagine that a higher order polynomial would do the trick. However, there is an alternate approach, one that naturally follows from Equation 122. Rather than model the stellar signal explicitly, I instead treat it non-parametrically by performing Gaussian process regression, in which I use a GP to estimate the covariance matrix $\mathbf{K}$. Provided $\mathbf{K}$ is a reasonable approximation to the true covariance of the stellar signal, Equation 122 will yield a set of PLD coefficients that fit out only the instrumental component of the noise. I illustrate this in Figure 32**d**, where I use a Matérn-3/2 kernel (see Equation 125 below) with amplitude $\alpha_m = 100$ and timescale $\tau_m = 20$ days to model the stellar signal. Unlike in the previous two cases, the PLD model (top panel, red) no longer attempts to fit out the stellar signal; in fact, the curve is almost identical to the fit to the original data (Figure 32**a**). The astrophysical signal is recovered to high fidelity by subtracting the PLD term (second panel). After removing it, the final residuals (third panel) show a comparable (in fact, marginally smaller) RMS to the residuals of the original data. It is also clear from the bottom panel that the transit shape and depth are well preserved.

I note that another advantage of GP regression is that the PLD model is relatively insensitive to the particular choice of kernel function and values of its hyperparameters. This can be seen in the example above; even though the stellar variability signal was generated from a 25-day period sinusoidal function, a radial Matérn-3/2 kernel with a timescale of 20 days was sufficient to ensure

the instrumental signal was removed without inflating the white noise. This is important because the stellar signal is generally unknown *a priori*, and any attempt to estimate it from the data will be subject to how well one can discriminate between instrumental and stellar effects.

Finally, I caution that GP regression assumes Gaussian noise. Deviations from Gaussianity, such as the presence of outliers, can invalidate the assumptions of the method and lead to poor fits to the data. Although I remove outliers prior to optimizing the GP, users accessing my de-trended light curves are encouraged to check for Gaussianity with methods such as an Anderson-Darling test. In §5.3 I describe my iterative procedure for removing outliers and optimizing the GP kernel.

5.3 Methods

In §5.2 I outlined the basics of principal component regression on the fractional pixel fluxes using a GP. In this section, I describe how I apply this method to *K2* data.

5.3.1 Pre-Processing

For each star in the *K2* EPIC catalog, I download the target pixel files and select aperture number 15 from the `K2SFF` catalog (Vanderburg and Johnson, 2014; Vanderburg, 2014), which is derived by fitting the Kepler pixel response function (PRF) to the stellar image. The `K2SFF` data contain twenty such apertures of varying sizes for each star. I find that aperture 15 is typically the best compromise between having enough pixels to generate a good basis set for PLD while preventing excess contamination from neighboring stars. I note that when contamination is not an issue, the choice of aperture has little effect on my results. As in Vanderburg and Johnson (2014), I then compute the median per-timestamp flux in the pixels of the image that lie outside the aperture and subtract it from all pixels to remove the background signal. In this chapter, I consider only the long cadence (LC) *K2* dataset.

Next, I perform iterative sigma-clipping to mask large outliers. Just as PLD can artificially inflate the white noise in an attempt to remove stellar variability (§5.2.3), it can do so in the presence of short timescale features such as eclipses, transits, flares, or cosmic ray hits. Deming et al. (2015) deal with this by explicitly adding a transit term to the PLD model (Equation 115), but this requires knowledge of the transit parameters. Since I do not know *a priori* whether there are transits in a given light curve, I instead mask all features in the light curve that are not well

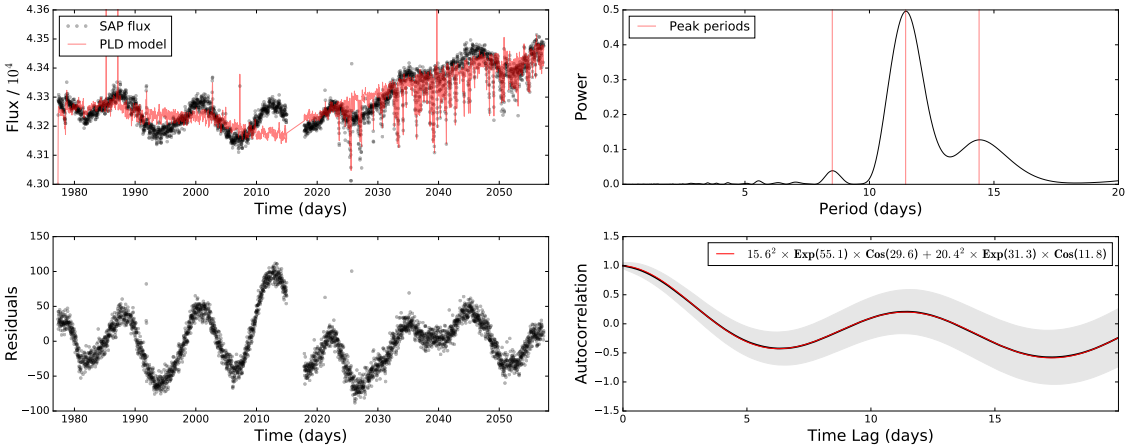

Figure 33: GP optimization procedure for EPIC 201497682. In the top left panel I plot the raw SAP flux (black) and a ten chunk, first order PLD fit (red); the residuals are shown in the panel below. These are used to compute the power spectrum of the stellar signal (top right), and its autocorrelation function (bottom right, black curve). Different kernels are then fit to the autocorrelation function, and the one with the lowest χ^2 value is chosen for the de-trending step (red curve). The grey envelope about the autocorrelation curve is the ad hoc standard error assumed to compute χ^2.

modeled by either the PLD terms or the GP. I then train my model on the masked light curve to compute the PLD coefficients and use those to de-trend the full, unmasked light curve.

I detect outliers by dividing the light curve into five chunks and de-trending each with a first order PLD model. I use a 2-day Matérn-3/2 covariance (see §5.3.2) with amplitude equal to the median standard deviation of the flux in all 2-day segments; again, I note that my results are relatively insensitive to the particular value of these parameters. Next, I perform a median absolute deviation (MAD) cut at 5σ to identify outliers. I then re-compute the PLD model by masking those outliers and de-trend the full light curve, identifying a new set of outliers. This process is repeated until no new outliers are identified. I find that this identifies deep transits, eclipses, and other outliers that could affect the quality of the PLD fit.

5.3.2 GP Optimization

The next step is to compute the covariance matrix $\mathbf{K}$, which I do using the `george`[11] Python package (Ambikasaran et al., 2016). I parametrize it as

$$\mathbf{K}_{ij} = k_w(t_i, t_j) + k_t(t_i, t_j) \tag{123}$$

where

$$k_w(t_i, t_j) = \sigma^2 \delta_{ij} \tag{124}$$

is a white noise kernel with standard deviation σ and k_t is either an additive or multiplicative combination of one or more of an exponential kernel k_e, a Matérn-3/2 kernel k_m, and a cosine kernel k_c:

$$\begin{aligned} k_e(t_i, t_j) &= \alpha_e e^{-|t_i - t_j|/\tau_e} \\ k_m(t_i, t_j) &= \alpha_m \left(1 + \sqrt{3(t_i - t_j)^2}\right) e^{-\sqrt{3(t_i - t_j)^2}/\tau_m} \\ k_c(t_i, t_j) &= \alpha_c \cos\left(\frac{2\pi}{P}(t_i - t_j)\right). \end{aligned} \tag{125}$$

The choice of kernel depends on the properties of the stellar signal, which I do not know *a priori*, since it is mixed with the instrumental signal in the light curve. I therefore adopt an iterative procedure, where I guess at the initial kernel form and hyperparameters and use it to de-trend the light curve with PLD, thus obtaining an approximate stellar component. I then train the GP on this stellar component and use it to run my PLD analysis again. While this can be repeated multiple times, I find that two iterations are typically enough for most targets. This procedure is illustrated in Figure 33.

My initial kernel is a Matérn-3/2 kernel with $\tau_m = 2$ days and amplitude equal to the median variance in 2-day chunks of the light curve. I split the light curve into 5-10 roughly equal chunks and de-trend each of them with first order PLD using Equations 121 and 122. Since I use first order PLD at this step, I do not perform PCA here. I then compute the autocorrelation function of the

[11] `http://dan.iel.fm/george/current/`

de-trended data and perform least-squares fits on different additive and multiplicative combinations of the kernels in Equation 125, using the peaks in a Lomb-Scargle periodogram as the initial guess for the periods in the cosine kernels. I choose the kernel and corresponding hyperparameters that result in the best fit and compute its overall amplitude as well as the amplitude of the white noise kernel k_w by maximizing the marginal likelihood $\mathcal{L}$ of the data given the model,

$$\log \mathcal{L} = -\frac{1}{2}\mathbf{y}^\top \cdot \mathbf{K}^{-\mathbf{1}} \cdot \mathbf{y} - \frac{1}{2}\log|\mathbf{K}| - \frac{n}{2}\log 2\pi \tag{126}$$

where $\mathbf{y}$ is the de-trended flux and n is the number of data points in $\mathbf{y}$ (Rasmussen and Williams, 2006). In principle, one may optimize the parameters of each of the kernel combinations in this way to obtain the best estimate of the covariance matrix. However, such a procedure is computationally expensive. After considerable experimentation, I find that the method outlined above—which takes on the order of one minute on a single core of a 2.66 GHz machine—is sufficient to ensure PLD removes only the instrumental component of the noise for most targets.

5.3.3 Design Matrix Optimization

The next step in the pipeline is to construct the design matrix (Equation 120). In order to do so, one must choose the number of PLD principal components to use in the regression. This number must be large enough to capture most of the instrumental variability but not so large as to lead to overfitting. In fact, one must take care to prevent PLD (or any other de-trending technique) from fitting out the white noise component of the light curve. As the number of basis vectors grows to become very large, there begin to exist linear combinations of these vectors that can artificially remove white noise from the data, improving the apparent quality of the fit but leading to spurious results. This is best illustrated by considering the example in Figure 31, where I kept only 285 basis vectors in the bottom panel. Had I de-trended with all 8435 components, I would obtain the light curve shown in Figure 34. While the median scatter improved by a factor of ~ 4, the in-transit scatter increased by a factor of several thousand. This is because I masked the transits of K2-3b, c, and d during the de-trending step (see §5.3.1). The poor performance of the extrapolated model betrays its terrible predictive power, a clear sign of overfitting.

In order to limit the flexibility of my PLD model, I implement a simple cross-validation scheme. I divide the light curve into training sets (large chunks of data that I use to compute the PLD coef-

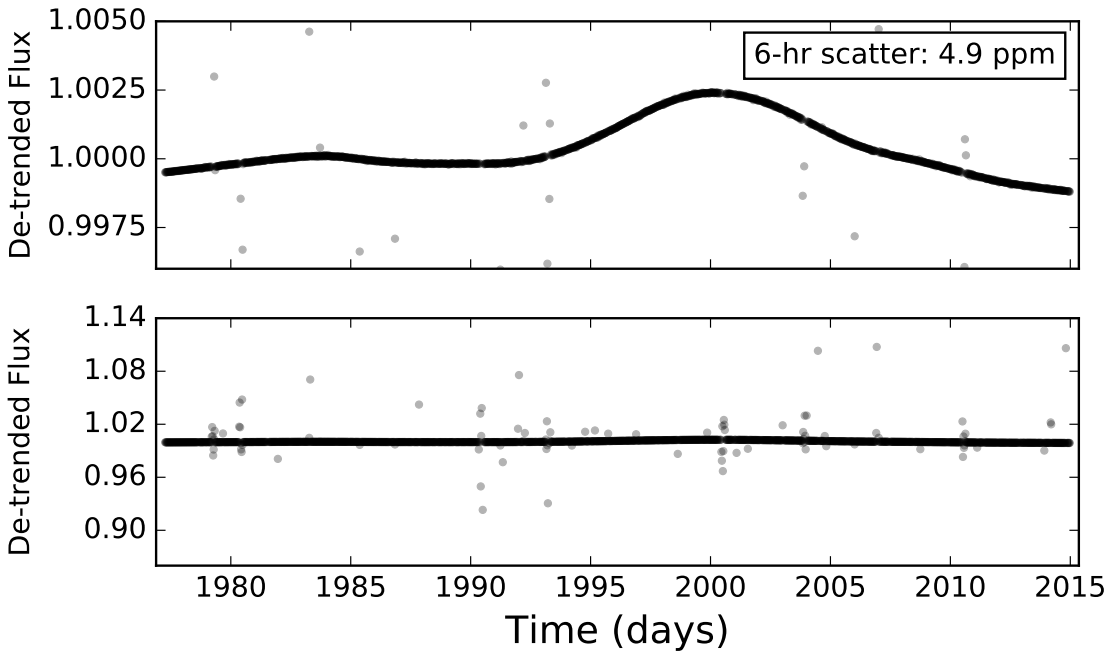

Figure 34: *Top*: Third order PLD applied to EPIC 201367065, but this time keeping *all* basis vectors. Compare to Figure 31. While the median scatter improved by a factor of about 4, the scatter in the transits (which were masked during the de-trending) increased by a factor of several thousand. *Bottom*: The same figure, but zoomed out to show the in-transit scatter.

ficients) and validation sets (small randomly selected chunks that I de-trend using the coefficients computed from the training sets). I then compute the 6-hr precision in the validation sets for a range of design matrix sizes (details below). Since the validation data is never used to compute the PLD coefficients, the model cannot overfit in the validation set. Instead, as the number of regressors becomes too large and PLD begins to fit out astrophysical signals and/or white noise, the scatter in the validation set will grow. This process is illustrated in Figure 35, where I plot the scatter in the training set (blue) and the scatter in the validation set (red) as a function of the number of principal components n_{pc}.

Initially, as n_{pc} increases, the scatter in both the training and validation data decreases. As the number of components increases further, the training precision continues to improve, eventually surpassing the photon noise limit for $n_{pc} \gtrsim 1000$. This is obviously unphysical and a clear sign of overfitting. The scatter in the validation set, on the other hand, begins to increase steadily above $n_{pc} \sim 50$, indicating that the predictive power of the model worsens as the number of principal

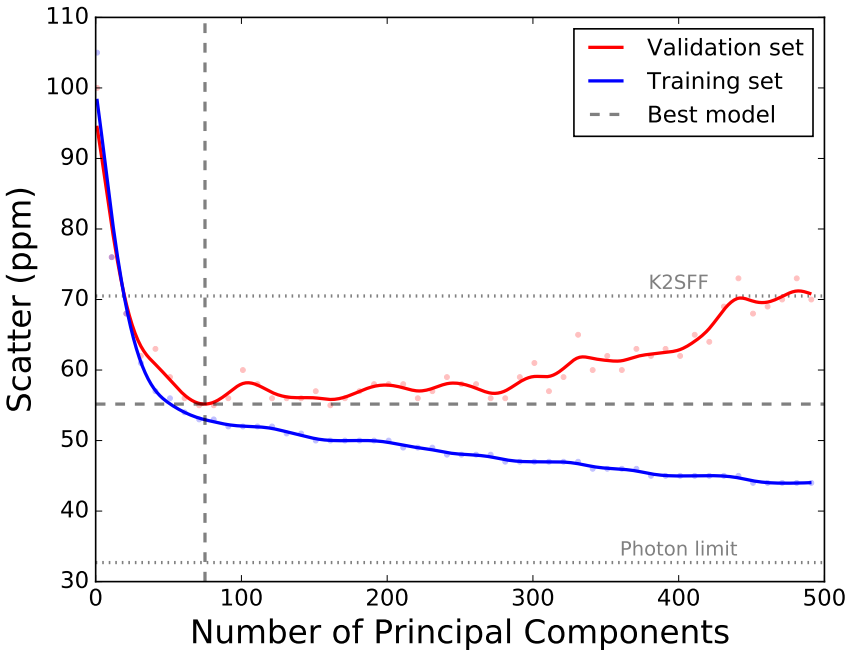

Figure 35: De-trended light curve precision as a function of the number of principal components for EPIC 201497682. The blue dots are the median 6-hr precision (in ppm) of the unmasked sections of the light curve (the training set); the red dots are the median precision in 6-hr chunks that were masked during the de-trending step (the validation set). Solid curves indicate my GP fit to the data points. Initially, the scatter decreases in both cases as the number of components is increased. However, above ~ 50 components, while the scatter in the training set continues to decrease, the scatter in the validation set (where the model is extrapolated) begins to grow. This is the signature of overfitting. I therefore choose 50 principal components for the de-trending, yielding a precision of 55 ppm (versus ~ 70 ppm for the `K2SFF` de-trended flux).

components is increased. The minimum in the red curve, which occurs for $n_{pc} \sim 80$, is the best one can do; above that point, PLD is likely to begin fitting out white noise.

I employ this cross-validation method for each EPIC target. In practice, I compute the 6-hr precision (see §5.4.3) in the validation sets fifty times for each value of n_{pc} and take the median. Each validation set is a group of 10 non-overlapping chunks that are masked when computing the PLD coefficients; each chunk is chosen randomly from the set of all contiguous 13-cadence segments of the light curve. My grid in n_{pc} typically ranges from 1 to 250 principal components, with a spacing of ten components. After computing the median scatter for each value of n_{pc}, I smooth the curves with a GP and choose the value of n_{pc} that minimizes the scatter in the validation set.

Finally, for campaign 1 I found it necessary to split all light curves into two separate timeseries,

with a breakpoint at BJD $-$ 2454833 = 2017, a mid-campaign data gap. Inspection of the top left panel of Figure 33 reveals a qualitative change in the instrumental systematics in the second half of the campaign, likely due to a reversal in the direction of the spacecraft roll (see, e.g., Aigrain et al., 2016). While similar reversals are also present in other campaigns, campaign 1 is the only one in which I obtain a de-trending improvement significant enough to justify the breakpoint.

5.3.4 De-trending

After optimizing the GP and choosing the order of the PLD, the number of principal components, and the number of light curve subdivisions, I de-trend each EPIC light curve by subtracting the model given by Equations 121 and 122. I then add the median value of the SAP flux to the result to obtain the final, de-trended light curves.

5.4 Results

5.4.1 Transit Injections

Since the PLD basis vectors are obtained from the *fractional* pixel fluxes, they do not in principle contain any astrophysical signals. De-trending with PLD should therefore preserve the transit shape and depth, a fact that is confirmed for the transits of several hot Jupiters observed with *Spitzer* (Deming et al., 2015). However, as I showed in §5.3.3, PLD can increase the scatter of regions of the data that are not used to compute the coefficients (such as transits that are flagged as outliers) when too many principal components are present (see Figure 34). If a transit is missed in the outlier detection step, the transit depth may also be affected, since PLD can exploit linear combinations of the white noise component to decrease the transit depth and therefore improve the likelihood of the fit.

In order to test whether my pipeline is robust against these effects, I ran a set of transit injection and recovery tests on a sample of EPIC targets that do not contain known transit events. I randomly selected 200 campaigns 0-6 stars from each magnitude bin in the range $8 \leq K_p \leq 18$ (K_p is the Kepler bandpass magnitude) for a total sample size of 2000 stars per run. I then injected synthetic transits of varying depths into the light curve by multiplying each pixel by a transit model generated by the `pysyzygy`[12] package, which calculates fast limb-darkened light curves based on

[12] `https://github.com/rodluger/pysyzygy`

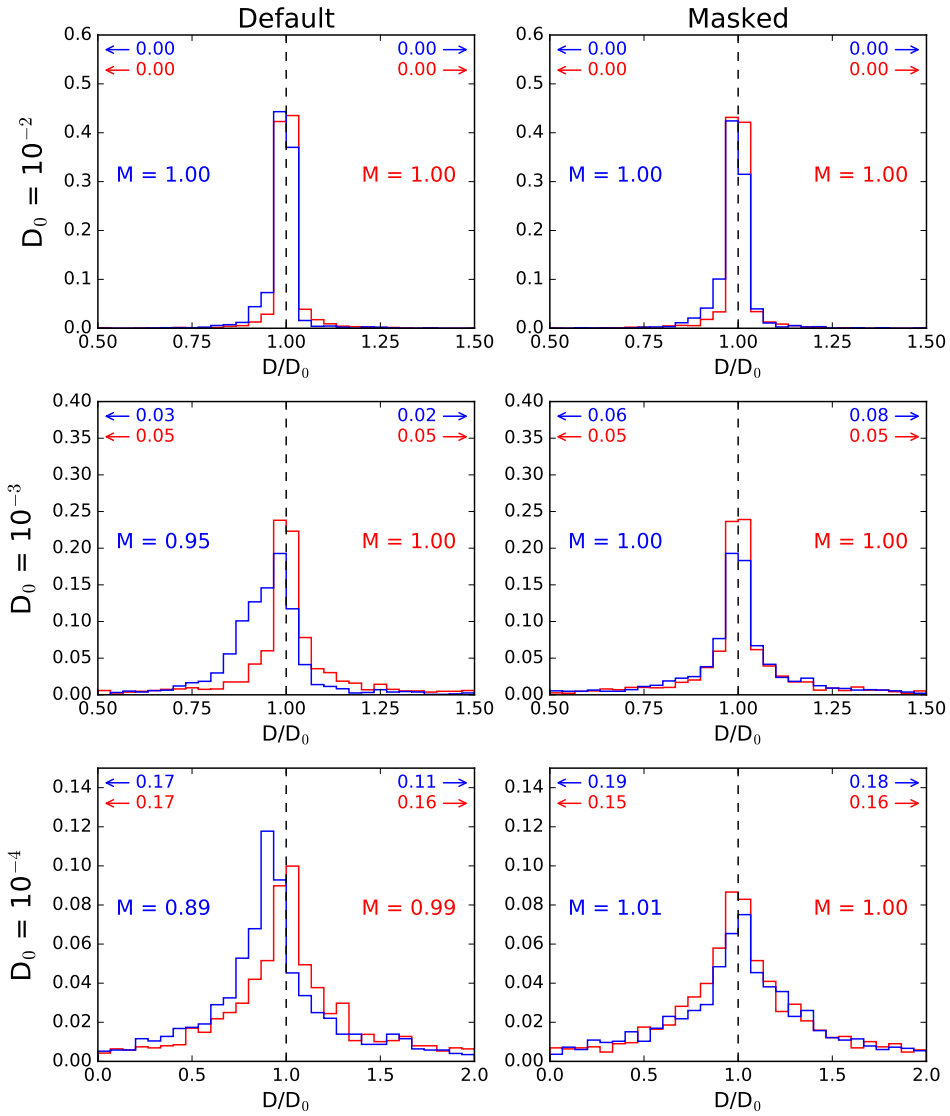

Figure 36: Transit injection results. Each panel shows the fraction of transits recovered with a certain depth ratio D/D_0 (recovered depth divided by true depth). Blue histograms correspond to the actual injection and recovery process performed with my pipeline; red histograms correspond to transits injected directly into the de-trended light curves and are shown for comparison. The values to the left and right of each histogram are the median D/D_0 for my pipeline and for the control run, respectively. The smaller values at the top indicate the fraction of transits recovered with depths lower and higher than the bounds of the plots. Finally, the two columns distinguish between the default runs (left) and runs where the transits were explicitly masked (right); the three rows correspond to different injected depths: 10^{-2}, 10^{-3}, and 10^{-4}. PLD preserves transit depths if the transits are properly masked; otherwise, a small bias toward smaller depths is introduced.

the Mandel and Agol (2002) transit model. I randomly chose orbital periods in the range 3–10 days and fixed the transit duration at 2.5 hours, assuming zero eccentricity and quadratic limb darkening parameters $a = 0.45$ and $b = 0.30$. I then ran my pipeline to de-trend the light curves.

Performing a full transit search is outside the scope of this chapter. Since my goal is to determine whether or not PLD can bias transit depths, I fix all the parameters except for the depth at their true values and recover the transit depth by linear regression, simultaneously fitting the baseline flux in the vicinity of the transit with a third-order polynomial. My results are shown in Figure 36.

Each panel shows two histograms of the recovered depth as a fraction of the true depth, D/D_0. In blue, I plot the recovery results after de-trending with my pipeline. As a control, I also injected transits directly into the *de-trended* data; I plot the corresponding distributions in red. Provided `EVEREST` does not affect transit depths or increase in-transit scatter, these two sets of distributions should be similar. The numbers at the top left and right of each panel respectively indicate the fraction of recovered depths below and above the limits of the plot. The median values M of each distribution are also shown.

Each row in the figure corresponds to a different injected depth: 10^{-2}, 10^{-3}, and 10^{-4}, ranging from a typical hot Jupiter to a roughly Earth-sized planet. The left column corresponds to the default runs of my pipeline; the right column corresponds to runs in which the transits were explicitly masked (more details on this below).

In the top left panel ($D_0 = 10^{-2}$), both distributions have medians equal to 1.00, corresponding to an unbiased recovery of the transit depth. Moreover, the spread is similar in both distributions, indicating that `EVEREST` does not introduce significant in-transit scatter. In the next two panels ($D_0 = 10^{-3}$ and 10^{-4}), however, a bias toward smaller transit depth is visible; the `EVEREST` de-trending causes transits of these small planets to be shallower by ~ 5 and $\sim 10\%$, respectively. This is because my sigma-clipping outlier detection technique (§5.3.1) is not effective at finding shallow transits, and thus these transits are not masked during the de-trending step. As I mentioned above, since I have no term in my design matrix (Equation 120) to explicitly model transits, the PLD model picks up the slack and partially fits out these features by inflating the white noise, slightly improving the quality of the overall fit. This is similar to the example in Figure 32, where PLD inflated the white noise to remove an astrophysical signal.

When the transits are properly masked (right column of the figure), both the bias and the increased scatter in the recovered depths disappear for transits of all depths. This shows that my cross-validation technique (§5.3.3) is correctly preventing overfitting and enforcing the high predictive power of the model in the masked regions of the dataset.

A particularly extreme example of transit depth alteration is illustrated in Figure 37, which shows the (folded and normalized) primary and secondary eclipses of the eclipsing binary EPIC 202072563. Three datasets are presented: the raw SAP flux (left), the default `EVEREST` light curve, and the masked `EVEREST` light curve. The `EVEREST` fluxes have been smoothed with a GP to remove astrophysical variability. It is clear from the second column that my outlier detection technique failed to mask all points during the eclipse, leading to a $\sim 20\%$ decrease in the eclipse depths. When properly masking the eclipses, however, the depth is correctly recovered (third column).

In order to remove the small bias introduced in the presence of unmasked transits and eclipses, I strongly encourage those making use of my light curves for transiting exoplanet characterization to run `EVEREST` while explicitly masking these transits. This can be done by specifying the transit times and durations when calling the `EVEREST` Python module; the code takes only a few seconds to run. The same applies to those using my light curves for transiting planet searches. Once features of interest are detected, one should run `EVEREST` again with those features masked to obtain unbiased estimates of the transit parameters.

It is important to note, however, that some transits with very low signal-to-noise could be completely fit out during the de-trending step, preventing their detection in the first place. As Foreman-Mackey et al. (2015) point out, this is an inevitable consequence of the de-trend-then-search method. It is *always* best to use a model that captures all the features of the data, allowing one to solve for instrumental noise, stellar variability, and transits *simultaneously*. To this end Foreman-Mackey et al. (2015) explicitly include a transit model in their design matrix, solving Equation 122 over a fine grid of periods and transit epochs. This eliminates the de-trending step in favor of a combined de-trending/planet searching step, which effectively circumvents the overfitting problems inherent to least-squares de-trending techniques. However, such an approach is very computationally expensive, given that each light curve must be processed once for every combination of transit parameters. I therefore defer this procedure to future work.

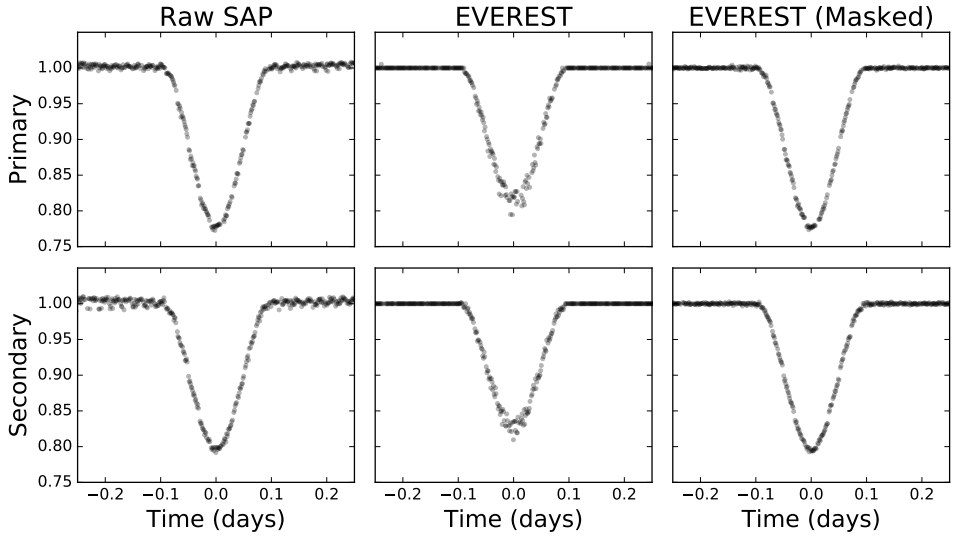

Figure 37: Primary and secondary eclipses of EPIC 202072563, folded on a period of 2.1237 days. From left to right, the columns show the raw SAP flux, the `EVEREST` flux, and the `EVEREST` flux obtained when explicitly masking the eclipses. Masking is critical to preserving the transit depth during de-trending for any pipeline that does not explicitly include a transit or eclipse model.

5.4.2 Limitations

I find that there are three particular situations in which PLD is likely to fail: saturated stars, stars in crowded apertures, and extremely variable stars. These limitations are inherent to the method itself and not specific to *K2* data.

5.4.2.1 Saturated Stars

The *Kepler* detectors begin to saturate for stars with $K_p \lesssim 11.5$, leading to flux bleeding along the pixel columns (Gilliland et al., 2010). Since the total charge is well conserved for stars dimmer than $K_p \approx 7$, this is not an issue for aperture photometry; transit depths and shapes are preserved in the SAP flux. However, the basic assumption of PLD—that the fractional pixel fluxes $p_{il} / \sum_k p_{ik}$ (see Equation 119) contain no astrophysical information—breaks down for these stars. This occurs because saturated pixels contain virtually no transit signal, as both the in-transit flux and the out-of-transit flux are above the saturation level, resulting in a relatively featureless timeseries. Since the total flux is conserved, the transit signal overflows into the adjacent pixels, traveling along the column until it reaches an unsaturated pixel. In these "overflow" pixels, the fractional

transit depth is much larger than the true depth, since it contains the total transit signal from all of the saturated pixels in that column. Consequently, the normalization $p_{il}/\sum_k p_{ik}$ will only *partially* remove the transit in the "overflow" pixels. Conversely, it will over-correct the saturated pixel fluxes, leading to an inverted transit shape in the PLD basis vectors corresponding to those pixels. This is illustrated in Figure 38, which shows the *fractional* pixel fluxes as a function of their position on the detector for the saturated hot-Jupiter host Kepler-3b. I again choose a *Kepler* target for illustrative purposes, though the idea applies equally to *K2*. Saturated pixels are indicated in red; "overflow" pixels are indicated in blue; both groups of pixels contain the transit signal. Unsurprisingly, PLD fails to properly de-trend this target, removing most of the transit signal along with the instrumental noise.

In Figure 39, I show an example of a saturated light curve of an eclipsing binary processed with `K2SFF` and `EVEREST`. The `K2SFF` pipeline removes a significant amount of the instrumental noise while preserving most of the astrophysical signal. `EVEREST`, on the other hand, fits out both the stellar variability and the eclipses, leading to spurious low scatter.

In principle, the saturation effects may be mitigated by collapsing each saturated column into a single timeseries; since charge is conserved within individual columns, this would enforce the correct transit depth in those "superpixels," which could then be used as basis vectors for PLD. However, I leave this procedure to future work. At present, I flag saturated stars in my database that are at risk of PLD overfitting. As a general rule, while saturation may occur for stars as faint as $K_p = 11.5$, I find that PLD preserves transit depths for most stars dimmer than $K_p = 11$.

5.4.2.2 Crowded Apertures

Another situation in which PLD performs poorly is in apertures containing significant contamination from other stars. In simple aperture photometry, the transit depth can be decreased in the presence of another star within the aperture. To correct for this, one can simply scale the depth by the (inverse of the) crowding metric, the ratio of the flux due to the planet host to the total flux in the aperture. However, this is not the case for data de-trended with PLD, for reasons I describe below.

In the case of a single star, the pixel flux p_{il} in the l^{th} pixel at the i^{th} time is the product of the (position-dependent) stellar signal a_{il} and the (position-independent) transit signal τ_i, such that

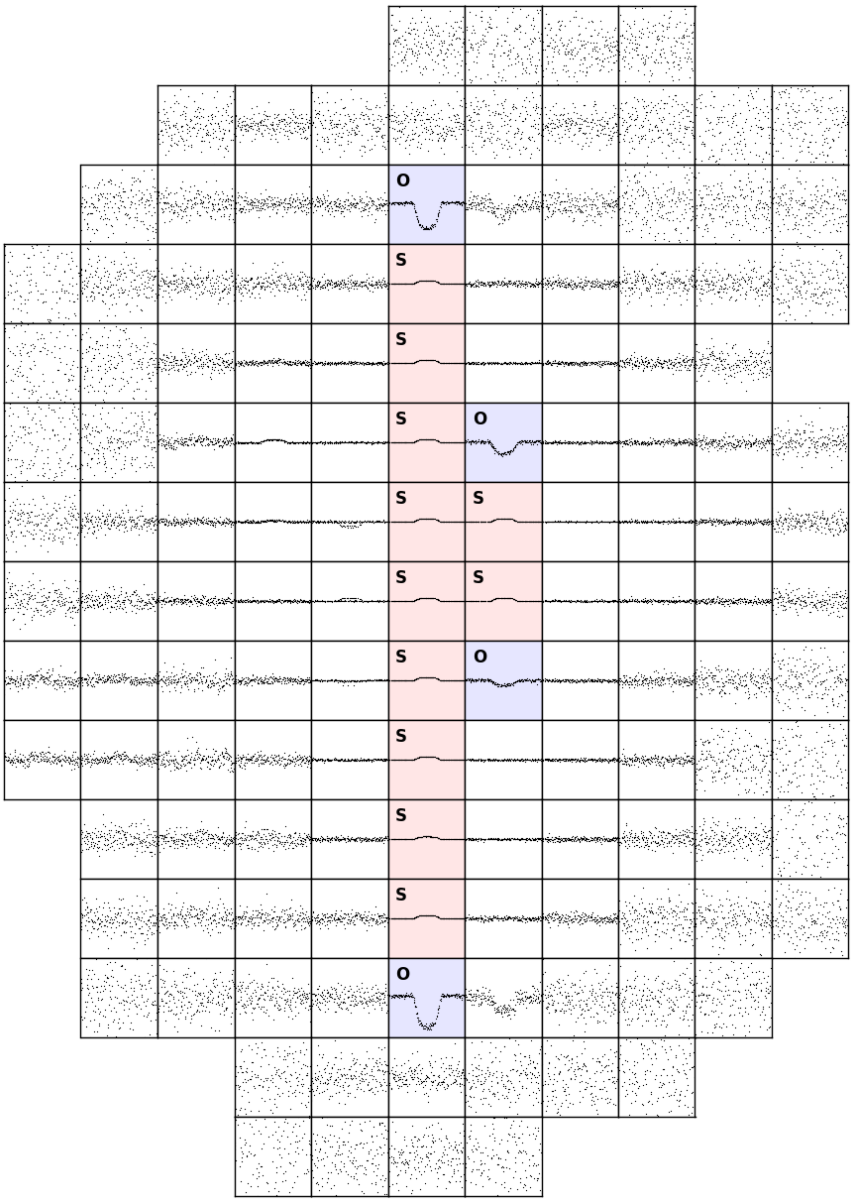

Figure 38: Fractional pixel fluxes $p_{il}/\sum_k p_{ik}$ for quarter 3 of Kepler-3, a $K_p = 9.2$ hot-Jupiter host observed by the original *Kepler* mission. The panels are arranged according to the positions of the pixels on the detector, and the data is smoothed and folded on the orbital period of Kepler-3b. Saturated pixels are highlighted in red and are labeled with an **S**; overflow pixels are highlighted in blue and labeled with an **O**. PLD fails for this system because the transit signal is present in several of the basis vectors.

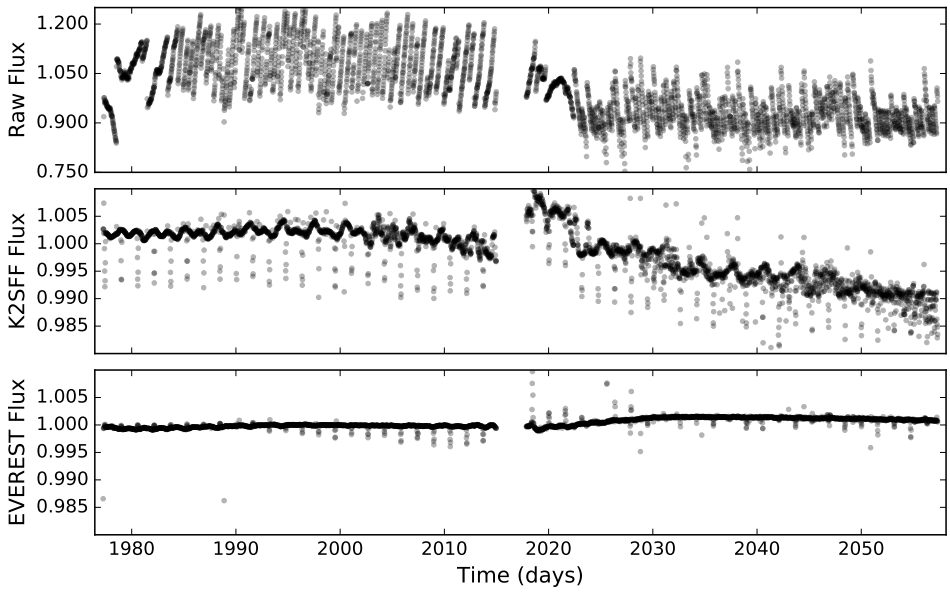

Figure 39: EPIC 201270464, a $K_p = 9.4$ saturated eclipsing binary. Plotted here is the raw flux (top), the `K2SFF` flux (center), and the `EVEREST` flux (bottom). PLD washes out the stellar variability along with most of the eclipses for some saturated stars.

the fractional pixel flux is

$$\frac{p_{il}}{\sum_k p_{ik}} = \frac{a_{il}\tau_i}{\sum_k a_{ik}\tau_i} = \frac{a_{il}}{\sum_k a_{ik}}, \tag{127}$$

which, as I have argued before, contains no transit signal. In the case of two stars, a_{il} and b_{il}, the first of which contains a transit, we may write the fractional flux as

$$\frac{p_{il}}{\sum_k p_{ik}} = \frac{a_{il}\tau_i + b_{il}}{\sum_k a_{ik}\tau_i + b_{ik}}. \tag{128}$$

Note that in this case the transit signal does not cancel, and the PLD basis vectors will contain some amount of transit information, leading to possible overfitting as in the case of a saturated aperture. For a dim contaminant star, $b_{il} \ll a_{il}$, we may write

$$\frac{p_{il}}{\sum_k p_{ik}} \approx \frac{a_{il}}{\sum_k a_{ik}} \left(1 + \frac{\Delta}{\tau_i}\right) \tag{129}$$

where

$$\Delta \equiv \left(\frac{b_{il}}{a_{il}} - \frac{\sum_k b_{ik}}{\sum_k a_{ik}} \right) \tag{130}$$

is a measure of the difference between the crowding in a given pixel and the crowding in the entire aperture. In general, the larger the value of Δ, the more power PLD will have to fit out the transit signal. This is the case for *bright contaminant sources near the edge of the aperture*, for which the value of b_{il}/a_{il} varies greatly across the aperture. Interestingly, for contaminant sources co-located with the transiting planet host (as in the case of binary stars or stars that are aligned but unresolved), the quantity b_{il}/a_{il} is constant across the detector and $\Delta = 0$, leading to a PLD basis set that does not contain transit information. This can also be shown from the exact expression (Equation 128). Crowding is therefore only a concern when there is a bright contaminant sufficiently separated from the transiting planet host. In practice, I find that PLD begins to overfit for contaminants that are separated by more than one pixel from the target and are either brighter than or within ~ 2 magnitudes of the target.

It is possible to correct for the effects of crowding if the quantity b_{il}/a_{il} is known, even if just approximately. However, this requires careful modeling of the stellar PSF and is beyond the scope of this chapter. In my database, I flag sources that may suffer from overfitting due to crowded fields.

5.4.2.3 Extremely variable stars

Finally, I note that `EVEREST` performs poorly for stars with high amplitude and high frequency variability. This is particularly the case for stars that are variable on timescales $\lesssim 1$ day, such as RR Lyrae variables or very short period eclipsing binaries. In these situations, the power of the GP model exceeds that of the PLD model at short timescales, and nearly all of the light curve variability (astrophysical *and* instrumental) is captured by the GP. As I rely only on the PLD model to produce my de-trended light curves, these often still contain the instrumental signal, in addition to artifacts introduced by tension between the GP and PLD models. The reader is referred to the `K2VARCAT` pipeline (Armstrong et al., 2015, 2016) for de-trended light curves of extremely variable stars.

5.4.3 Photometric Precision

The formal photometric noise metric developed by the *Kepler* team is the Combined Differential Photometric Precision (CDPP) (Christiansen et al., 2012), which is computed by the *Kepler* pipeline on transit-like timescales of 3, 6 and 12 hours. The CDPP evaluated for a given duration is defined so that it is equivalent to the depth of a transit of that duration that would yield a SNR of 1. In this section I adopt a proxy of the 6-hr CDPP that is easier to calculate than the formal metric defined in Christiansen et al. (2012). My approach is very similar to the approaches of Gilliland et al. (2011), Vanderburg and Johnson (2014) and Aigrain et al. (2016). In order to prevent correlated stellar noise from inflating the white noise calculation, I apply a 2-day, quadratic Savitsky-Golay (Savitzky and Golay, 1964) high-pass filter to the de-trended flux, clipping outliers at 5σ. I then compute the standard deviation of all contiguous 13-cadence chunks of data, take the median, and divide by $\sqrt{13}$ as in Vanderburg and Johnson (2014) to obtain the approximate 6-hr photometric precision of the data, which I henceforth refer to as the CDPP.

In Figure 40 I plot the CDPP of my de-trended fluxes for campaign 0-7 stars, excluding saturated targets (stars whose brightest pixel's median flux $> 1.6 \times 10^5$e/s) and stars in highly crowded apertures (stars with $\Delta K_p < 5$ neighbors inside the aperture or brighter neighbors within 2 pixels of the edge of the aperture), since PLD is likely to lead to artificially low CDPP in those cases. `EVEREST` CDPP values are plotted as blue dots. For comparison, in yellow I plot the CDPP calculated for the raw SAP fluxes of all stars observed by the original *Kepler* mission. `EVEREST` recovers the raw *Kepler* CDPP for most $K_p \lesssim 13$ stars and yields light curves with CDPP within about 1.5 times that of *Kepler* for $13 \lesssim K_p \lesssim 15$. The sparser clump of stars above the bulk group between $K_p = 11$ and 14 are likely giants, whose short-timescale pulsations are not efficiently captured by the high-pass filter and thus appear to be more noisy; this clump is also present in the *Kepler* distribution.

In Figure 41 I again plot the CDPP as a function of K_p, but this time on a logarithmic scale, comparing the `EVEREST` values (blue dots) to the raw *K2* CDPP (red dots) and the approximate 6-hr photon limit (dashed yellow line), which I calculate as

$$\mathrm{CDPP}_{phot} = \frac{10^6}{\sqrt{21600 \times \bar{F}}} \tag{131}$$

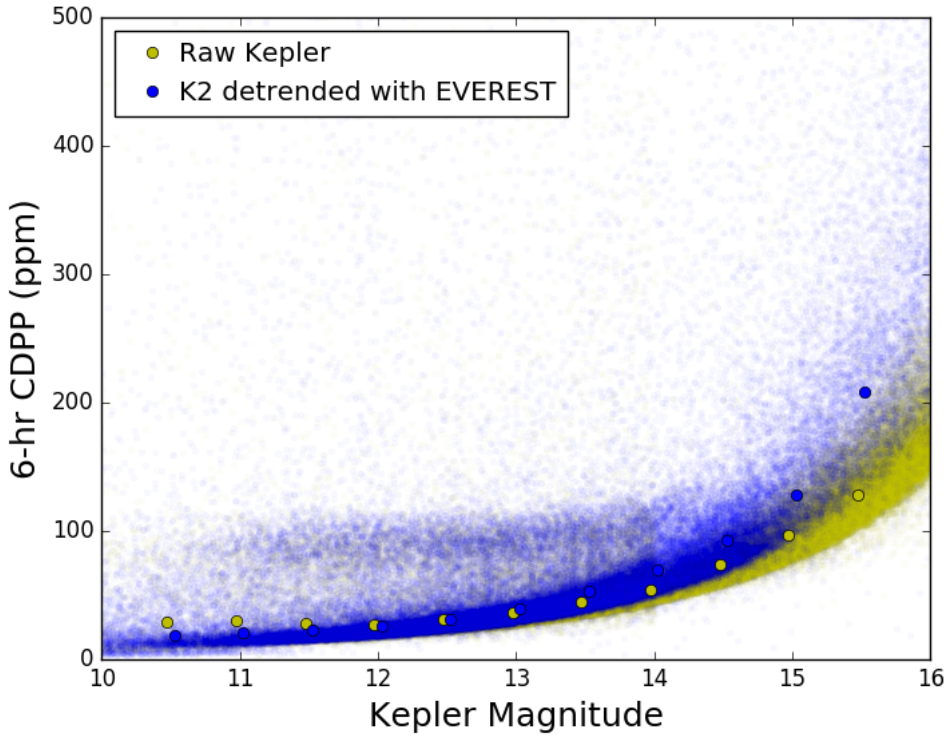

Figure 40: 6-hr photometric precision as a function of Kepler magnitude K_p for all stars observed by *Kepler* (yellow dots) and for all unsaturated, non-crowded *K2* targets in campaigns 0-7 de-trended with `EVEREST` (blue dots). The median values are indicated for 0.5 magnitude-wide bins with filled circles. My pipeline recovers the original *Kepler* precision for stars brighter than $K_p \approx 13$.

where $\bar{F}$ is the average SAP flux in e/s at a particular value of K_p. The dashed red and solid blue lines indicate the median CDPP in 0.5 magnitude-wide bins in the raw and de-trended light curves, respectively. `EVEREST` leads to nearly an order-of-magnitude improvement in the CDPP for the brightest targets, approaching the photon limit for $K_p \lesssim 13$ dwarfs.

5.4.4 Comparison to Other Pipelines

In this section I compare the precision of my de-trended light curves to that of the `K2VARCAT` (Armstrong et al., 2015), `K2SFF` (Vanderburg and Johnson, 2014), and `K2SC` (Aigrain et al., 2016) light curves. Though other pipelines exist (e.g., Lund et al., 2015; Huang et al., 2015; Crossfield et al., 2015), here I focus on those that are also available for download at the MAST HLSP *K2*

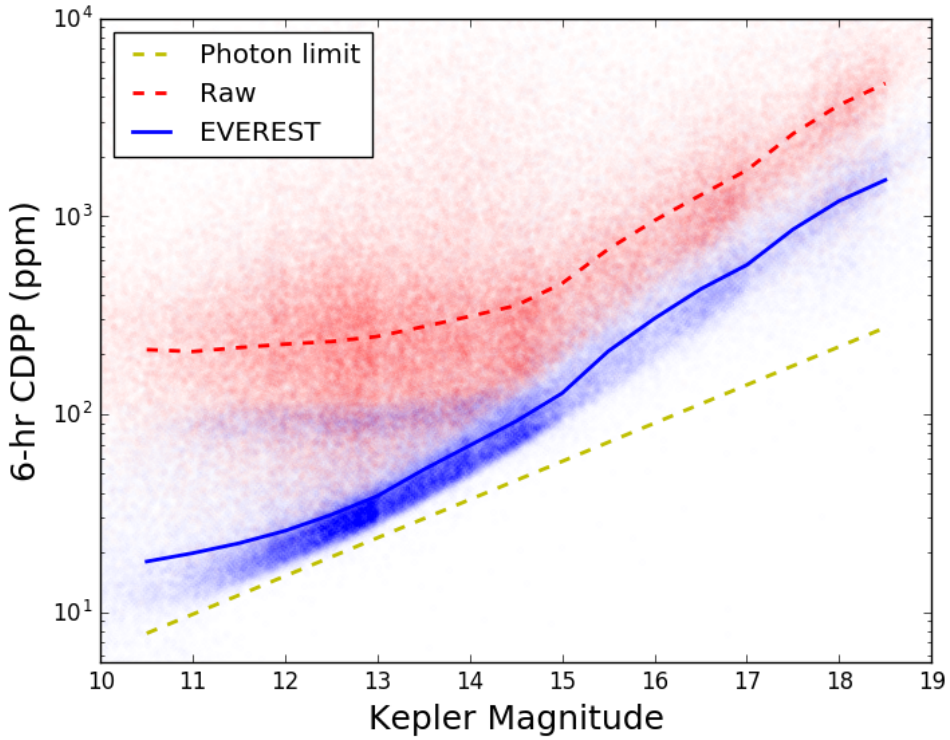

Figure 41: A comparison of the raw *K2* 6-hr precision (red dots) and the `EVEREST` precision (blue dots) as a function of K_p. The lines indicate the median in 0.5 magnitude-wide bins. I also plot the approximate photon limit (dashed yellow line) for reference. `EVEREST` leads to an order-of-magnitude improvement in the CDPP for the brightest stars.

archive.[13] For each of the pipelines, I download all available de-trended light curves and compute the CDPP as described in §5.4.3. In Figures 42-45, I plot a pairwise comparison of the precision achieved for all campaign 0-7 light curves that are available in both my catalog and that of the other pipeline. These plots are based on Figure 10 in Aigrain et al. (2016); the plots show the relative 6-hr CDPP difference between my light curves and those produced by the `K2VARCAT`, `K2SFF`, and `K2SC` pipelines as a function of Kepler magnitude. Results for individual light curves are indicated by blue dots, and the median in half magnitude-wide bins is shown as a black line. As before, I trim 5σ outliers and apply a high-pass filter prior to computing the CDPP. I again exclude from the comparison saturated stars and stars in highly crowded apertures.

Figure 42 shows the comparison to the `K2VARCAT` pipeline. My de-trended precision is systematically better at all magnitudes shown; on average, I achieve about double the precision (half the

[13] https://archive.stsci.edu/k2/hlsp/

CDPP). There are very few cases where K2VARCAT yields lower CDPP values than EVEREST (points above zero).

Figure 43 shows the same plot, but comparing EVEREST to K2SFF. Again, EVEREST yields higher median precision at all magnitudes; EVEREST light curves have $\sim$ 20% less scatter on average. The gain in precision changes with K_p, but also with campaign; this can be seen in Figure 44. Note, in particular, that the EVEREST precision is significantly higher relative to that of K2SFF in campaigns 0-2; in campaign 2, particularly, the average star has half the scatter in the EVEREST light curves. In the more recent campaigns, the precision gain relative to K2SFF is smaller, but EVEREST still yields higher median precision at all magnitudes. In general, EVEREST light curves also have significantly fewer outliers than K2SFF light curves. This is clear from Figures 46 and 47 in the following section: K2SFF light curves often display a band of outliers above or below the light curve continuum, most likely associated with thruster firings. PLD naturally corrects for these, obviating the need to discard the several hundred thruster firing data points in each campaign.

Finally, in Figure 45 I plot the comparison to the PDC version of the K2SC pipeline, whose performance is the best out of the three I consider here. In order to ensure the two datasets are compared on equal footing, I use the systematics-corrected K2SC fluxes rather than the fully whitened fluxes; I obtain these by summing the FLUX and TREND_T columns in the dataset. I then apply a Savitsky-Golay filter, as I do to the EVEREST data, to remove the stellar components of the variability. Since campaigns 0-2 are not available in the K2SC catalog, this plot shows the comparison for campaigns 3-5 only. Once again, EVEREST light curves have higher precision than those of K2SC by $\sim$ 10% for bright stars and $\sim$ 5% for stars in the range $16 \lesssim K_p \lesssim 11$.

All plots shown here display a significant amount of scatter. While EVEREST yields the lowest median CDPP in all cases, I recommend comparing light curves from the different pipelines when the highest precision is desired for specific targets. Moreover, at this time EVEREST performs poorly for saturated targets and for those in very crowded fields.

5.4.4.1 Example Light Curves

In Figures 46 and 47 I plot my de-trended light curves for two *K2* planet candidate hosts, EPIC 201367065 and 205071984. These were chosen specifically to illustrate the advantages of the PLD technique and are, in a sense, best case scenarios. Both stars have $K_p \approx 12$ and are the only bright

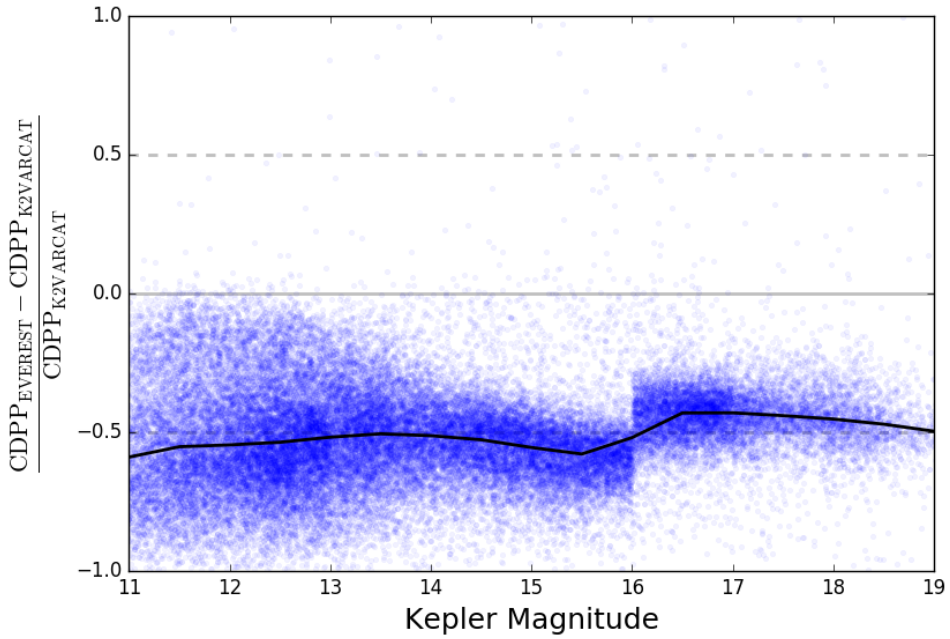

Figure 42: Relative 6-hr CDPP difference between `EVEREST` and `K2VARCAT` light curves for campaigns 0-7. Blue dots show differences for individual stars, while the black line indicates the median in 0.5 magnitude-wide bins. Negative values indicate higher precision in the `EVEREST` light curves; compare to Figure 10 in Aigrain et al. (2016). On average, `EVEREST` yields light curves with half the scatter for all Kepler magnitudes $K_p > 11$.

sources in their apertures. (There are two faint sources, $K_p \approx 17$ and $K_p \approx 18$, at the edge of the aperture of EPIC 205071984, but they are faint enough as to not affect the de-trending.) In each of the figures, I plot the `K2SFF` light curves at the top (red) and the `EVEREST` light curves at the bottom (black). To the right, I plot the folded transits of their planet candidates after removing the stellar variability signal with a GP.

In both cases, the `EVEREST` precision is a factor of about 2 higher than that of `K2SFF`: `EVEREST` yields 30.9 ppm for EPIC 201367065 and 24.0 ppm for EPIC 205071984. This is visible in both the full light curve and in the folded transits, which display significantly less scatter. Note, importantly, that the greater de-trending power of `EVEREST` does not affect the low-frequency stellar variability, which is present equally in both sets of light curves.

The corresponding light curves in the `K2VARCAT` database have CDPP values of 43.1 and 63.4, respectively. At the time of writing, these light curves are not present in the `K2SC` catalog.

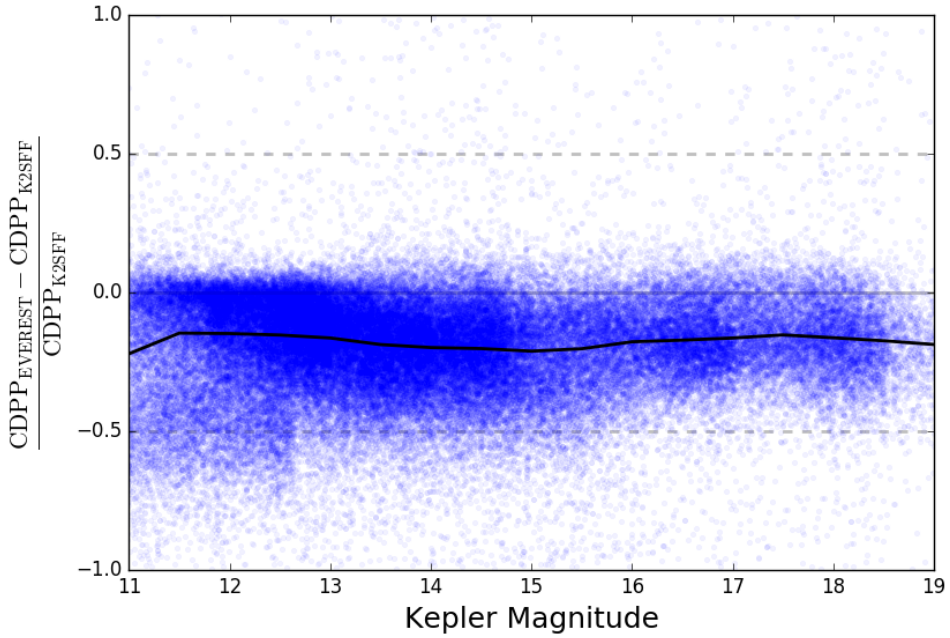

Figure 43: Same as Figure 42, but comparing `EVEREST` to `K2SFF`. Once again, negative values correspond to higher precision in the `EVEREST` light curves. My pipeline yields higher precision light curves for most *K2* stars and does better on average for all Kepler magnitudes $K_p > 11$.

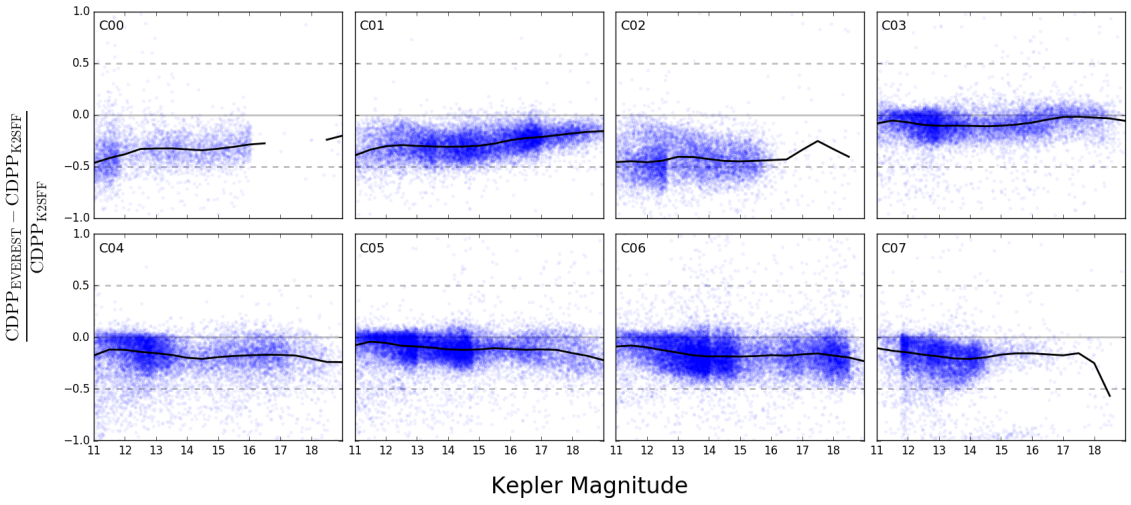

Figure 44: A comparison of the `EVEREST` and `K2SFF` CDPP for each individual campaign. Note the marked difference between campaigns 3-7 and campaigns 0-2. For campaign 2, in particular, the relative improvement is close to 0.5, corresponding to an average `EVEREST` precision a factor of 2 higher than `K2SFF`.

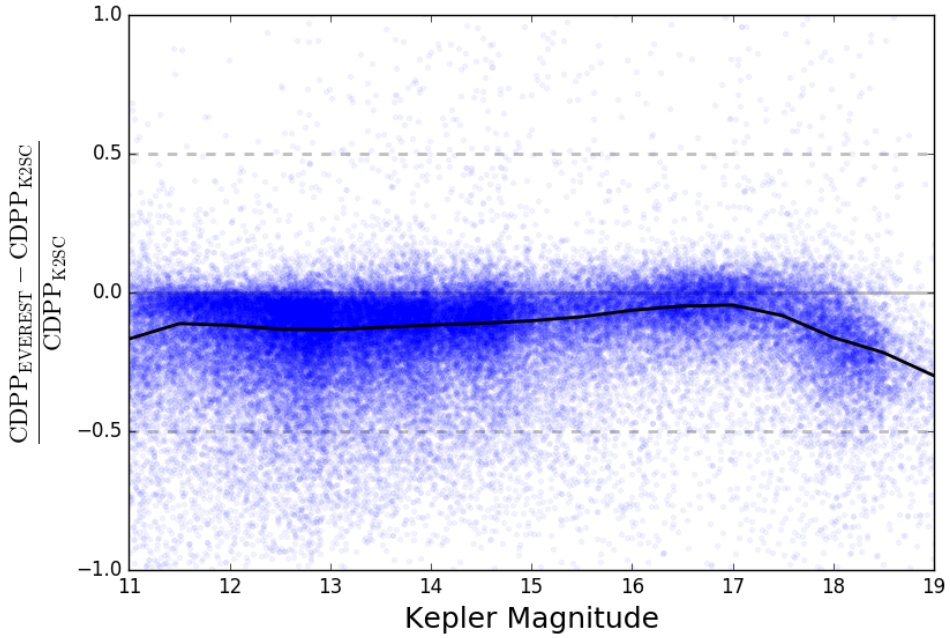

Figure 45: Same as Figure 42, but comparing `EVEREST` to `K2SC`. To ensure both sets of light curves are on the same footing, the `K2SC` CDPP is computed from the PDC flux corrected for the instrumental systematics *only*. As before, a Savitsky-Golay filter is then applied to both sets of light curves. The median relative difference is once again negative everywhere, indicating that `EVEREST` yields higher precision light curves at all magnitudes.

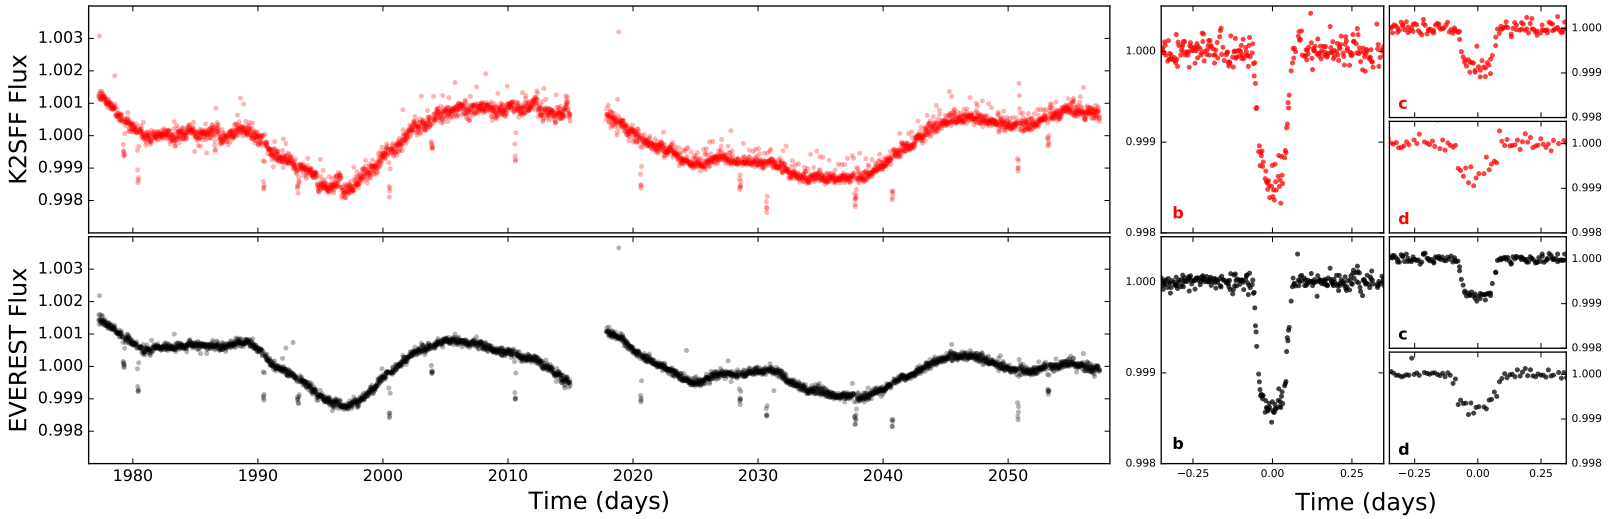

Figure 46: De-trended light curves for the campaign 1 star EPIC 201367065 (K2-3, Crossfield et al., 2015). *Top*: The de-trended `K2SFF` flux (left) and the GP-smoothed flux folded on the periods of the planets b, c, and d (right). *Bottom*: The de-trended `EVEREST` flux. The 6-hr CDPP is 30.9 ppm for `K2SFF` and 16.6 ppm for `EVEREST`, a factor of ~ 2 improvement.

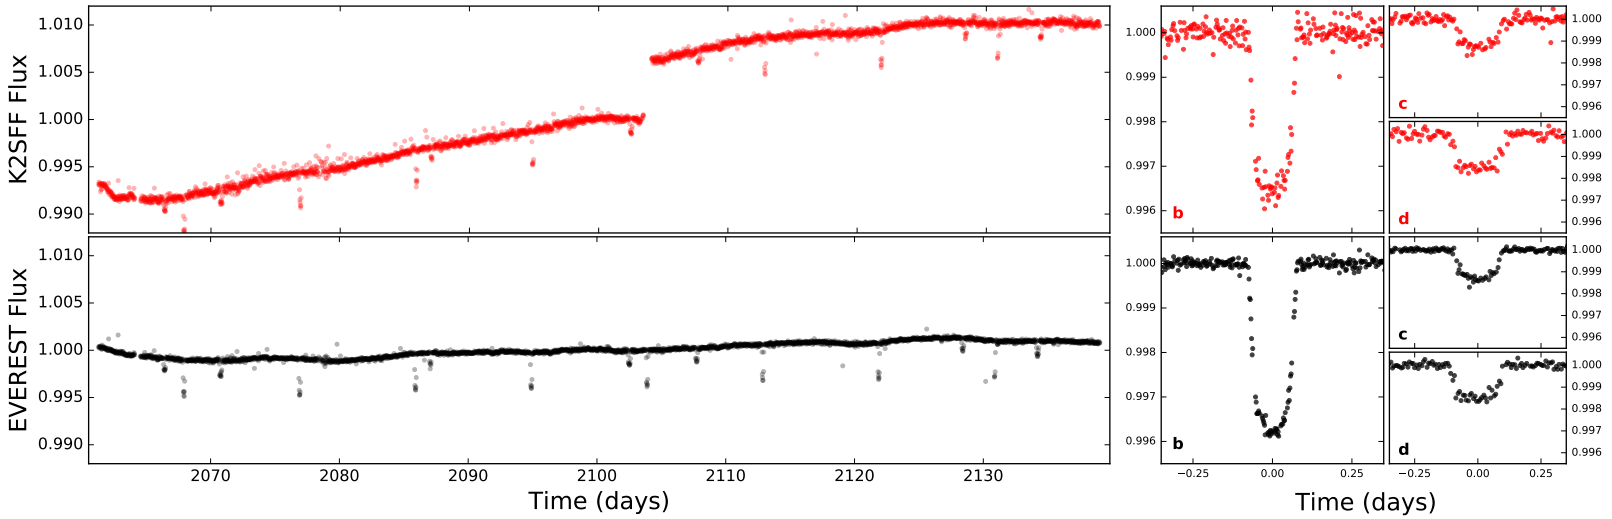

Figure 47: De-trended light curves for EPIC 205071984, a campaign 2 star with three known planet candidates (Sinukoff et al., 2015). As in Figure 46, the `K2SFF` light curve and the folded transits of EPIC 205071984.01 (b), 205071984.02 (c), and 205071984.03 (d) are shown at the top; the equivalent plots for `EVEREST` are shown at the bottom. The 6-hr CDPP is 56.1 ppm for `K2SFF` and 24.0 ppm for `EVEREST`, a factor of $\gtrsim 2$ improvement.

5.5 Conclusions

In this chapter I introduced `EVEREST`, a pipeline developed to yield the highest precision light curves for *K2* stars. `EVEREST` builds on the pixel level decorrelation (PLD) technique of Deming et al. (2015), extending it to third order in the pixel fluxes and combining it with principal component analysis to yield a set of basis vectors that together capture most of the instrumental variability in the data. Gaussian process (GP) regression is then performed to remove the instrumental systematics while preserving astrophysical signals. In order to prevent overfitting, we developed a method to determine the ideal number of principal components to use in the fit, yielding reliable, high precision de-trended light curves for all *K2* campaigns to date.

I validated my model by performing transit injection and recovery tests and showed that when transits were properly masked by my iterative sigma-clipping technique, I recovered the correct depths without any bias. When transits were not masked (the case of many low signal-to-noise transits, which are missed by my outlier detection step), PLD de-trending resulted in somewhat shallower transits by $\lesssim 10\%$. I therefore strongly encourage those making use of my light curves for transiting exoplanet characterization to run `EVEREST` while explicitly masking these transits. My code is implemented in Python, is user-friendly, and takes only a few seconds to run for a given target. The same applies to those using my light curves for transiting planet searches. Once features of interest are detected, one should run `EVEREST` again with those features masked to obtain unbiased estimates of the transit parameters. While the decreased transit depth can in principle preclude the detection of very low signal-to-noise planets in the `EVEREST` light curves, I find that the increased precision of these light curves relative to other pipelines is sufficient to enable the detection of previously unknown, small transiting planets around *K2* host stars (Kruse et al., 2017, in prep.). Since `EVEREST` preserves stellar signals, these light curves should also greatly aid in stellar variability and asteroseismology studies.

For stars brighter than $K_p \approx 13$, I found that `EVEREST` recovers the photometric precision of the original *Kepler* mission; for fainter stars, the median precision is within a factor of 2 of that of the original mission. I further compared my de-trended light curves to those produced by the other *K2* pipelines available at the MAST HLSP *K2* archive. `EVEREST` light curves have systematically higher precision than `K2SFF`, `K2VARCAT` and `K2SC` for all Kepler magnitudes $K_p > 11$. Currently, `EVEREST` performs poorly for saturated targets and for those in highly crowded fields.

My catalog of de-trended light curves is publicly available at `https://archive.stsci.edu/prepds/everest/` and will be constantly updated for new *K2* campaigns. The code used to de-trend the light curves is open source under the MIT license and is available at `https://github.com/rodluger/everest`, along with user-friendly routines for downloading and interacting with the de-trended light curves. A static release of version `1.0` of the code is also available at `http://dx.doi.org/10.5281/zenodo.56577`. Since the only inputs to `EVEREST` are the pixel level light curves, the techniques developed here can be generally applied to light curves produced by any photometry mission, including the upcoming *TESS* mission (Ricker et al., 2015), to remove instrumental noise and enable the detection of small transiting planets.

6 EVEREST 2.0

6.1 Introduction

In the previous chapter, I developed the first version of my $K2$ pipeline, `EVEREST` (EPIC Variability Extraction and Removal for Exoplanet Science Targets). I employed a variant of pixel level decorrelation (PLD) based on the method of Deming et al. (2015), a data-driven approach that uses a star's own pixel-level light curve to remove instrumental effects. I showed that `EVEREST` recovered the original *Kepler* precision for stars brighter than *Kepler*-band magnitude $Kp \approx 13$, yielding higher average precision than any publicly available $K2$ catalog for unsaturated stars.

In this chapter, I present an update to my pipeline, which I refer to as `EVEREST 2.0`. By combining PLD with spacecraft motion information obtained from nearby stars, this update improves the precision of $K2$ light curves at all magnitudes relative to version `1.0` and addresses certain issues with overfitting. `EVEREST 2.0` also reliably de-trends saturated stars and stars observed in short cadence mode, obtaining comparable or even higher de-trending power for these targets.

This chapter is organized as follows: in §6.2 I derive the mathematical framework of the `EVEREST 2.0` model, and in §6.3 I describe the implementation of my pipeline in detail. I present my results in §6.4 and some additional remarks in §6.5. In §6.6 I briefly discuss how to use the `EVEREST` catalog and code, and in §6.7 I summarize the work.

6.2 The PLD Model

Here I describe the mathematical formulation of the `EVEREST` pixel level decorrelation (PLD) model. In PLD, products of the fractional fluxes in each pixel of the target aperture are used as regressors in a linear model:

$$\begin{aligned}\mathbf{m} = &\sum_i a_i \frac{\mathbf{p}_i}{\sum\limits_n \mathbf{p}_n} + \\ &\sum_i \sum_j b_{ij} \frac{\mathbf{p}_i \mathbf{p}_j}{(\sum\limits_n \mathbf{p}_n)^2} + \\ &\sum_i \sum_j \sum_k c_{ijk} \frac{\mathbf{p}_i \mathbf{p}_j \mathbf{p}_k}{(\sum\limits_n \mathbf{p}_n)^3}.\end{aligned} \tag{132}$$

In the expression above, $\mathbf{m}$ is the instrumental model and $\mathbf{p}_i$ is the flux in the i^{th} pixel; both are vector quantities defined at an array of times $\mathbf{t}$. Each term corresponds to a different PLD order (first, second, and third) resulting from a Taylor expansion of the instrumental signal. The a_i, b_{ij}, and c_{ijk} are the linear weights of the model, which we seek to obtain below. For a detailed discussion of the theory behind PLD, see Deming et al. (2015) and Chapter 5. Below I simply discuss its mathematical implementation.

6.2.1 Regularized Regression (`rPLD`)

Given a timeseries $\mathbf{y}$ with N_{dat} data points, we wish to find the linear combination of N_{reg} regressors that best fits the instrumental component of $\mathbf{y}$. Expressed in vector form, my linear model is thus

$$\mathbf{m} = \mathbf{X} \cdot \mathbf{w}, \tag{133}$$

where $\mathbf{X}$ is the $(N_{dat} \times N_{reg})$ design matrix constructed from the set of regressors (the fractional pixel fluxes in Equation 132) and $\mathbf{w}$ is the $(N_{reg} \times 1)$ vector of weights (the set $\{a_i, b_{ij}, c_{ijk}\}$). If $\mathbf{w}$ is known, the de-trended light curve is simply

$$\mathbf{y}' = \mathbf{y} - \mathbf{m}. \tag{134}$$

In Chapter 5, I obtained $\mathbf{w}$ by maximizing the likelihood function

$$\begin{aligned} \log \mathcal{L}_0 = &- \frac{1}{2} \left(\mathbf{y} - \mathbf{X} \cdot \mathbf{w}\right)^\top \cdot \mathbf{K}^{-\mathbf{1}} \cdot \left(\mathbf{y} - \mathbf{X} \cdot \mathbf{w}\right) \\ &- \frac{1}{2} \log |\mathbf{K}| - \frac{N_{dat}}{2} \log 2\pi, \end{aligned} \tag{135}$$

where $\mathbf{K}$ is the $(N_{dat} \times N_{dat})$ covariance matrix of the data and $\mathbf{y}$ is the $(N_{dat} \times 1)$ simple aperture photometry (SAP) flux. Since the number of third order PLD regressors can be quite large (on the order of several thousand for a typical star, which is larger than the number of data points), the problem is ill-posed, meaning that a unique solution does not exist and maximizing $\log \mathcal{L}_0$ is likely to lead to overfitting. I thus constructed $\mathbf{X}$ from the (smaller) set of N_{pc} principal components of the PLD regressors. I chose N_{pc} by performing cross-validation, which aims to maximize the predictive power of the model while minimizing overfitting.

However, while principal component analysis (PCA) yields a set of components that captures the most variance among the PLD vectors, there is no guarantee that the principal components are the ideal regressors in the PLD problem. Dimensionality reduction techniques such as PCA inevitably lead to information loss, and so it is worthwhile to consider alternative regression methods to fully exploit the potential of PLD.

A common regression method for ill-posed problems is regularization, in which a prior is imposed on the values of the weights $\mathbf{w}$. Since overfitting occurs when $\mathbf{w}$ becomes very large, regularization recasts the problem by adding a penalty term to the likelihood that increases with increasing $|\mathbf{w}|$. While many forms of regularization exist, I focus on L2 regularization, since it has an analytic solution. Recently, Wang et al. (2016) used L2 regularization in a "causal pixel model" to de-trend light curves from the original *Kepler* mission. L2 regularization is mathematically equivalent to placing a Gaussian prior on each of the weights $\mathbf{w}$, so that the posterior likelihood function becomes

$$\log \mathcal{L} = \log \mathcal{L}_0 - \frac{1}{2}\mathbf{w}^\top \cdot \mathbf{\Lambda}^{-1} \cdot \mathbf{w} - \frac{1}{2}\log|\mathbf{\Lambda}|\,, \tag{136}$$

where $\mathbf{\Lambda}$ is the $(N_{reg} \times N_{reg})$ regularization matrix, which I choose to be diagonal for simplicity and for computational efficiency:

$$\Lambda_{m,n} = \lambda_n^2 \delta_{mn}. \tag{137}$$

Each element λ_n^2 in $\mathbf{\Lambda}$ is the variance of the zero-mean Gaussian prior on the weight of the corresponding column of the design matrix, $\mathbf{X}_{*,n}$. In practice I find that if I choose the λ_n correctly, this model has a higher predictive power than the PCA model adopted in Chapter 5.

Given this formulation, our task is to find the weights $\hat{\mathbf{w}}$ that maximize the posterior probability $\mathcal{L}$. Differentiating Equation (136) with respect to $\mathbf{w}$, we get

$$\begin{aligned}\frac{\mathrm{d}\log\mathcal{L}}{\mathrm{d}\mathbf{w}} &= \mathbf{X}^\top \cdot \mathbf{K}^{-1} \cdot \mathbf{y} \\ &\quad - \left(\mathbf{X}^\top \cdot \mathbf{K}^{-1} \cdot \mathbf{X} + \mathbf{\Lambda}^{-1}\right) \cdot \mathbf{w}.\end{aligned} \tag{138}$$

By setting this expression equal to zero, we obtain the maximum *a posteriori* prediction for the

weights,

$$\hat{\mathbf{w}} = \left(\mathbf{X}^\top \cdot \mathbf{K}^{-1} \cdot \mathbf{X} + \mathbf{\Lambda}^{-1}\right)^{-1} \cdot \mathbf{X}^\top \cdot \mathbf{K}^{-1} \cdot \mathbf{y} \tag{139}$$

with corresponding model

$$\mathbf{m} = \mathbf{X} \cdot \left(\mathbf{X}^\top \cdot \mathbf{K}^{-1} \cdot \mathbf{X} + \mathbf{\Lambda}^{-1}\right)^{-1} \cdot \mathbf{X}^\top \cdot \mathbf{K}^{-1} \cdot \mathbf{y}. \tag{140}$$

In what follows, I refer to the implementation of PLD with regularized regression as `rPLD`.

6.2.2 Cross-validation

Similarly to Chapter 5, I solve for $\mathbf{\Lambda}$ by cross-validation. For each value of $\mathbf{\Lambda}$, the model is trained on one part of the light curve (the training set) and used to de-trend the other part of the light curve (the validation set); see §6.3.7 for details. The value of $\mathbf{\Lambda}$ that results in the minimum scatter in the validation set is then chosen for the final de-trending step.

In principle, each of the λ_n in $\mathbf{\Lambda}$ could take on a different value, but solving for each one requires minimizing an N_{reg}-dimensional function and is not computationally tractable. Instead, I simplify the problem by requiring that all regressors of the same order have the same regularization parameter λ. Provided we write the third order design matrix in the form

$$\mathbf{X} = \begin{pmatrix} \mathbf{X_1} & \mathbf{X_2} & \mathbf{X_3} \end{pmatrix}, \tag{141}$$

where $\mathbf{X_n}$ is the matrix of n^{th} order regressors, we may express the regularization matrix as

$$\mathbf{\Lambda} = \begin{pmatrix} \mathbf{\Lambda_1} & & \\ & \mathbf{\Lambda_2} & \\ & & \mathbf{\Lambda_3} \end{pmatrix} \tag{142}$$

where $\mathbf{\Lambda_n} = \lambda_n^2 \mathbf{I}$ is the n^{th} order regularization matrix and λ_n^2 is the variance of the prior on the n^{th} order regressors.

A typical $K2$ star with 30 aperture pixels has $N_{reg} \sim 5{,}000$ regressors and $N_{dat} \sim 500$ data points in each cross-validation light curve segment (see §6.3.7). Evaluating the matrix inverse in

Equation (140) is thus computationally expensive, and becomes prohibitive during cross-validation, since this must be done for every set of λ_n's. Fortunately, we can reduce the number of calculations with some linear algebra. First, we apply the Woodbury matrix identity (e.g., Golub and Van Loan, 1996) to Equation (140), obtaining

$$\mathbf{m} = \mathbf{X} \cdot \mathbf{\Lambda} \cdot \mathbf{X}^\top \cdot \left(\mathbf{X} \cdot \mathbf{\Lambda} \cdot \mathbf{X}^\top + \mathbf{K}\right)^{-1} \cdot \mathbf{y}. \tag{143}$$

Next, note that

$$\begin{aligned}
\mathbf{X}\cdot\mathbf{\Lambda}\cdot\mathbf{X}^\top &= \begin{pmatrix} \mathbf{X_1} & \mathbf{X_2} & \mathbf{X_3} \end{pmatrix} \begin{pmatrix} \mathbf{\Lambda_1} & & \\ & \mathbf{\Lambda_2} & \\ & & \mathbf{\Lambda_3} \end{pmatrix} \begin{pmatrix} \mathbf{X_1^\top} \\ \mathbf{X_2^\top} \\ \mathbf{X_3^\top} \end{pmatrix} \\
&= \lambda_1^2 \mathbf{X_1} \cdot \mathbf{X_1^\top} + \lambda_2^2 \mathbf{X_2} \cdot \mathbf{X_2^\top} + \lambda_3^2 \mathbf{X_3} \cdot \mathbf{X_3^\top} \\
&= \sum_n \lambda_n^2 \mathbf{X_n^2},
\end{aligned} \tag{144}$$

where I have defined

$$\mathbf{X_n^2} \equiv \mathbf{X_n} \cdot \mathbf{X_n^\top}. \tag{145}$$

We may thus re-write the maximum *a posteriori* model as

$$\mathbf{m} = \sum_n \lambda_n^2 \mathbf{X_n^2} \cdot \left(\sum_n \lambda_n^2 \mathbf{X_n^2} + \mathbf{K} \right)^{-1} \cdot \mathbf{y}. \tag{146}$$

The matrix that we must invert in Equation (146) has dimensions $(N_{dat} \times N_{dat})$, while that in Equation (140) is $(N_{reg} \times N_{reg})$. Since $N_{reg} \sim 10 N_{dat}$, casting the model in this form can greatly speed up the computation. In practice, I pre-compute the three matrices $\mathbf{X_n^2}$ at the beginning of the cross-validation step, so the only time-consuming operation in Equation (146) is the inversion.

6.2.3 Neighboring Stars (`nPLD`)

One of the downsides of PLD is that the regressors used in the linear model tend to be noisy. This is particularly a problem for faint targets, whose PLD vectors are dominated by photon noise.

Their light is also distributed over fewer pixels, resulting in a smaller set of vectors on which to regress. The effect of this is evident in Figure 10 of Chapter 5, which shows how `EVEREST 1.0` light curves for faint ($Kp \gtrsim 15$) stars are significantly noisier than those of the original *Kepler* mission. This decrease in de-trending power at the faint end affects most other $K2$ pipelines as well, as those usually regress on information derived (either directly or indirectly) from the motion of the stellar image across the detector, which is similarly noisy.

While the spacecraft motion (the dominant source of instrumental noise in $K2$) is imprinted at relatively low signal-to-noise ratio (SNR) on the light curves of any one star, the collective light curves of all the stars on the detector encode this information at very high SNR. Therefore, a straightforward way to improve the performance of PLD for faint targets is to include this information in the design matrix. To this end, in `EVEREST 2.0` I incorporate the PLD vectors of a set of other targets located on the same CCD module as the target of interest when performing the regression. I dub this method `nPLD`, for *n*eighboring *PLD*, and discuss its implementation in §6.3.4. The third-order design matrix (Equation 141) is now

$$\mathbf{X} = \begin{pmatrix} \mathbf{X_1} & \mathbf{X_1'} & \mathbf{X_2} & \mathbf{X_2'} & \mathbf{X_3} & \mathbf{X_3'} \end{pmatrix}, \tag{147}$$

where $\mathbf{X_n'}$ is the design matrix constructed from the n^{th} order PLD vectors of all the neighboring targets. For computational speed, I still solve for a single prior amplitude λ_n for each PLD order, but in principle one could assign different priors to the neighboring vectors. I discuss the implementation of `nPLD` in §6.3.

6.3 Implementation

6.3.1 Light Curves

As in Chapter 5, I downloaded all stars in the $K2$ EPIC catalog with long and/or short cadence target pixel files and adopted aperture #15 from the `K2SFF` catalog (Vanderburg, 2014; Vanderburg and Johnson, 2014). I masked all cadences with `QUALITY` flags 1-9, 11-14, and 16-17, though I still compute the model prediction on them. For campaigns 0-2, I remove the background signal as described in Chapter 5; for more recent campaigns, the background is removed by the *Kepler* team.

Next, I perform iterative sigma clipping to identify and mask outliers at 5σ. During each

iteration, I compute the linear (unregularized) PLD model and smooth it with a Savitsky-Golay filter (Savitzky and Golay, 1964), then identify outliers based on a median absolute deviation (MAD) cut. I implement this outlier-clipping step at the beginning of each cross-validation step (§6.3.7), each time computing the model with a higher (regularized) PLD order, to progressively refine the outlier mask.

6.3.2 GP Optimization

In order to compute the covariance matrix $\mathbf{K}$ for each target, I use a Gaussian process (GP), as I did in Chapter 5. GP optimization can be costly, especially when performing model selection over a range of possible kernels and optimizing many hyperparameters simultaneously. For this reason, in Chapter 5 I cut corners and performed kernel selection based on fits to the autocorrelation function of the light curve, which I also used to fix the timescale and/or period of those kernel(s). I then ran a nonlinear minimizer to optimize the overall amplitude of the GP. In practice, this worked reasonably well, but often failed for light curves dominated by high frequency stellar variability. After much experimentation, I decided to forego the kernel selection step in favor of using a single carefully optimized Matérn-3/2 kernel with an added white noise term:

$$\mathbf{K}_{ij} = \alpha \left(1 + \frac{\sqrt{3(t_i - t_j)^2}}{\tau}\right) e^{-\frac{\sqrt{3(t_i - t_j)^2}}{\tau}} + \sigma^2 \delta_{ij}, \tag{148}$$

where the hyperparameters σ, α, and τ are the white noise amplitude, red noise amplitude, and red noise timescale, respectively, and t_i and t_j correspond to the timestamps of cadences i and j. I initialize the hyperparameters at random values and run a nonlinear optimizer to solve for the maximum likelihood (Equation 136), keeping the PLD model parameters fixed; I repeat this process several times and retain the highest likelihood solution. As with outlier clipping, I progressively optimize the GP at each of the three cross-validation steps, so that each time I train the GP on a light curve that is increasingly dominated by stellar variability (as opposed to instrumental systematics).

In principle, the quasi-periodic kernels used in Chapter 5 should be better suited to handling variable stars, but in practice I find that a properly optimized Matérn-3/2 kernel is flexible enough to fully capture the variability and prevent PLD overfitting. I discuss this further in §6.5.

6.3.3 Breakpoints

Because the instrumental noise properties are quite variable over the course of $K2$ campaigns, I find a significant improvement in the de-trending power of my regularized regression model when I subdivide light curves into two or three segments. This is in contrast to the PCA approach in Chapter 5, where I did not find it necessary to split the timeseries. For all campaigns except 4 and 7, I add a single breakpoint in the light curve near the mid-campaign point, where the spacecraft roll is at a minimum. For campaigns 4 and 7, I find it necessary to insert two breakpoints. I cross-validate and de-trend each light curve segment separately and mend them at the end. In order to mitigate flux discontinuities at the breakpoints, I train the model in each segment on an additional 100 cadences past the breakpoint to remove potential edge effects and offset the models in each segment so that they align at the breakpoint. While this method removes flux discontinuities, it can introduce discontinuities in the derivative of the flux, showing up as spurious "kinks" in the light curve. I remove these in a post-processing step (§6.4.7).

6.3.4 Neighboring Stars

In principle, the larger the number of neighboring PLD vectors I include in the `nPLD` design matrix, the higher the de-trending power of the model. However, adding regressors significantly increases computing time, so I would like to instead select a small set of high SNR regressors that capture most of the spacecraft motion information. Moreover, since I employ a single prior for all n^{th} order regressors, adding many foreign PLD vectors effectively dilutes the contribution of the target's own PLD vectors, which are the only ones that can correct instrumental signals arising from local pixel sensitivity variations, and in practice results in poorer quality light curves. After much experimenting, I obtain the highest average de-trending power when the number of neighboring stars is about ten. I therefore de-trend each $K2$ target with the aid of the PLD vectors of ten randomly selected bright ($11 \leq Kp \leq 13$) stars on the same detector module as the target. To minimize contamination of the target by outliers in its neighbors' fluxes, I linearly interpolate over all neighbor data with flagged `QUALITY` bits. Finally, for computational reasons, I neglect all cross terms of the form $\prod_{i \neq j} p_i p_j$, where p_i is the flux in the i^{th} pixel, when computing the neighbors' PLD vectors. Cross terms typically encode information specific to the sets of pixels from which they are computed and aid in correcting features such as thruster firing discontinuities

(Luger et al., 2016). Cross terms from stars other than the target in question are therefore of little help in the de-trending and can be safely neglected.

One potential pitfall of `nPLD` is that if the PLD assumptions break down for any of the neighboring targets, the PLD regressors may become contaminated with astrophysical information from that neighbor. This is not in general an issue, since overfitting would only occur if an astrophysical signal in the target star and in its neighbor had the same period and the same phase. However, in the (unlikely) case that PLD fails for the neighboring star and this star happens to be an eclipsing binary or a transiting exoplanet host, it is possible that its transit signals could get imprinted onto the target star's de-trended light curve, resulting in potential false positive planet detections down the line. The two cases relevant to $K2$ in which PLD could fail in such a way are for saturated stars and stars with bright contaminant sources in their apertures (Luger et al., 2016). As I show in §6.3.5 below, it is straightforward to adapt PLD to work reliably for saturated stars, thereby circumventing this issue. But while `EVEREST 2.0` is more robust against overfitting of crowded stars (§6.5), highly crowded apertures remain an issue for PLD. When de-trending with `nPLD`, I therefore select neighboring stars with no other known sources in their apertures that are bright enough ($\Delta Kp < 5$) to contaminate the PLD vectors.

6.3.5 Saturated Stars

As discussed in Chapter 5, PLD typically fails for stars with saturated pixels, resulting in overfitted light curves with artificially low scatter and suppressed astrophysical information (such as transits with significantly shallower depths). This happens because saturated pixels contain nearly no astrophysical information, as the signal overflows into adjacent pixels in the same column and is ultimately dumped into the pixels at the top and bottom of the bleed trails; these "tail" pixels ultimately contain more astrophysical information than the other pixels in the aperture. Since PLD implicitly assumes that astrophysical information is constant across the aperture, the method breaks down for these stars, and PLD vectors from pixels in the saturated columns become capable of fitting out the astrophysical information in the rest of the aperture.

An obvious workaround is to simply discard pixels in saturated columns from the set of PLD regressors. However, this does not work well in practice, since the remaining regressors often have much lower SNR than the SAP flux and thus have low de-trending power. I instead suggested

in Chapter 5 that collapsing saturated columns into single pixels — by co-adding the fluxes in each of the pixels and treating the resulting timeseries as a single PLD pixel — could reduce the effect of saturation, since charge is conserved along the bleed trail. While this ensures PLD does not overfit, it leads to the loss of some of the information about the vertical motion of the stellar point spread function (PSF) across the detector. This leads to significantly poorer de-trending, and therefore I did not employ this method in the first version of the pipeline. However, I find that including the PLD vectors of neighboring stars in the design matrix (i.e., `nPLD`) effectively restores the information lost when saturated columns are collapsed, leading to high quality de-trended light curves of saturated stars.

In practice, I collapse all columns containing one or more pixels whose flux comes within 10% of (or exceeds) the pixel well depth for the corresponding detector channel during more than 2.5% of the timeseries. I obtained the well depths from Table 13 of the Kepler Instrument Handbook. [14] To avoid aperture losses, I found it necessary to use aperture `K2SFF` #19 (the largest of the PSF-based apertures) for these stars. I further padded these apertures by two pixels at the top and bottom of saturated columns to ensure that none of the target flux bled out of the aperture; see §6.5.3.

As an example, in Figure 48 I plot the light curve of EPIC 202063160, a saturated campaign 0 eclipsing binary. The raw light curve is shown at the top and the light curve de-trended with `EVEREST 1.0` is shown at the center. Because three of the columns in the aperture contain saturated pixels (right panel), `EVEREST 1.0` almost completely fits out the eclipses. With column-collapsed `nPLD` (bottom), the eclipse is preserved and the instrumental signal is effectively removed.

6.3.6 Short Cadence

I treat short cadence targets in much the same way as long cadence targets, with the exception that I find it necessary to introduce more breakpoints in the light curves. This is due primarily to computational reasons (short cadence $K2$ light curves are over 10^5 cadences in length; computing Equation (146) for the entire light curve is not feasible). Moreover, I find that noise on sub-30 minute timescales is only properly removed when the size of the light curve segments is kept small. In practice, I find that on the order of 30 breakpoints result in the best de-trending. This might raise concerns of overfitting, but since short cadence light curves contains 30 times more data

[14] `archive.stsci.edu/kepler/manuals/KSCI-19033-001.pdf`

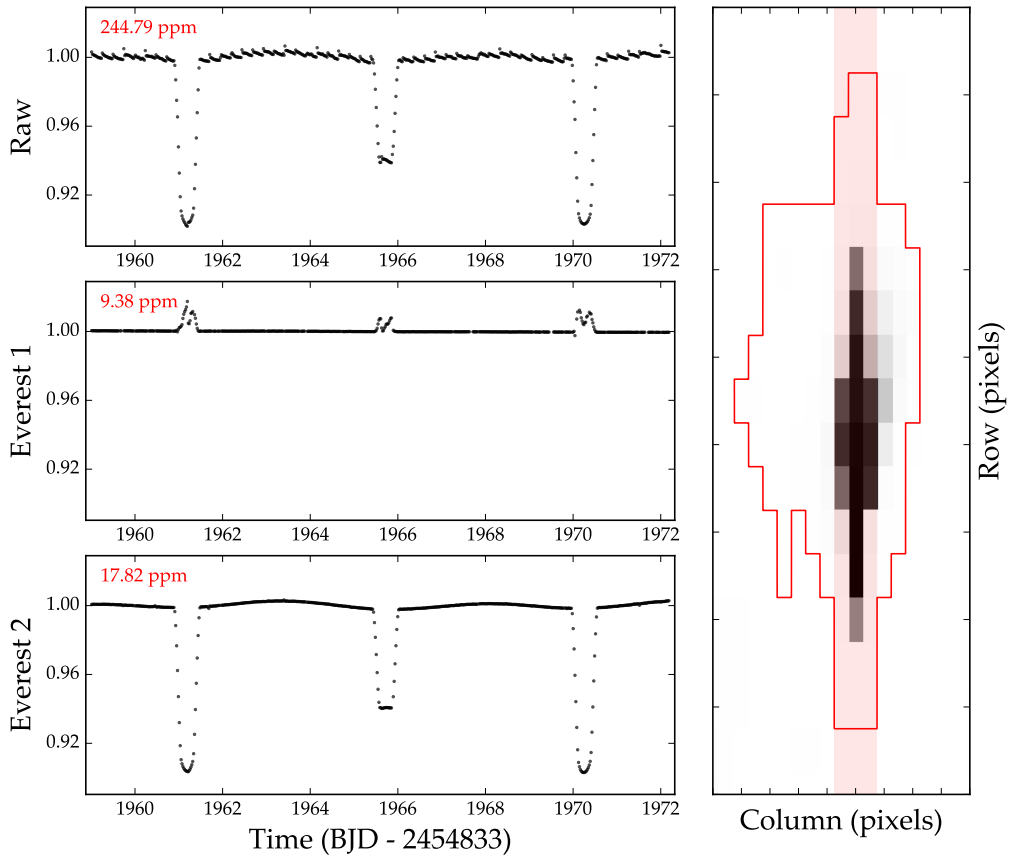

Figure 48: EPIC 202063160, a saturated Kp= 9.2 campaign 0 eclipsing binary. Shown is a portion of the raw light curve (top), the light curve "de-trended" with `EVEREST 1.0` (center), and the light curve de-trended with `EVEREST 2.0` (bottom); the y axis in each of these plots is the normalized flux. The pixel image is shown at the right on a linear scale, with the adopted aperture contour indicated in red. The three columns highlighted in red contain saturated pixels. Despite a great improvement in the precision, `EVEREST 1.0` leads to severe overfitting, causing the eclipses to all but disappear. By collapsing saturated columns, `EVEREST 2.0` correctly de-trends saturated stars without overfitting.

than long cadence light curves, and I split the latter into two segments, each of the short cadence segments has about twice as many cadences as the long cadence ones.

The major downside of such a large number of segments are the discontinuities that could be introduced at each breakpoint. As before, I overcompute the model into adjacent segments and match the models at the breakpoints, but some de-trended light curves display occasional jumps in either the flux or its derivative.

A second issue with short cadence light curves concerns deep transits and eclipses. As I discussed in Chapter 5, PLD may attempt to fit out these features if they are not properly masked, since doing so can result in a very large (but spurious) improvement in the photometric precision. With long cadence light curves, transit masking can be done by the user by simply re-computing the model with the appropriate cadences masked, since the transits are sparse and their presence does not significantly affect the cross-validation step. Moreover, outlier clipping usually masks most deep transits anyways, so this is hardly ever a problem. However, that is not the case with short cadence light curves, where transits and eclipses span upwards of fifty contiguous cadences. Since these features are so smooth, they are not flagged as outliers. And since the transit signal is no longer sparse — as it makes up a substantial fraction of the light curve segment — it is far more likely to bias the cross-validation step. In practice, I find that this leads to substantial *under*fitting of short cadence light curves with deep transits. As λ_n increases, PLD begins to fit out the transit and the scatter in the validation set grows, forcing the algorithm to select very low values of $\hat{\lambda}_n$ and resulting in de-trended light curves that still contain significant instrumental signals.

I therefore explicitly mask all deep transits and eclipses in the short cadence light curves *before* the cross-validation step. Since only deep transits are likely to bias the cross-validation, and since the number of short cadence light curves in each campaign is relatively small, these can easily be identified by inspection.

6.3.7 Cross-validation

The principal step in the de-trending process is determining the prior amplitudes λ_n in Equation (146), which I do by cross-validation. My method is analogous to that of Chapter 5, where I performed cross-validation to obtain the optimal number of principal components to regress on. However, here I seek to optimize a three-dimensional function $\sigma_{\rm v}(\lambda_1, \lambda_2, \lambda_3)$, where $\sigma_{\rm v}$ is the scat-

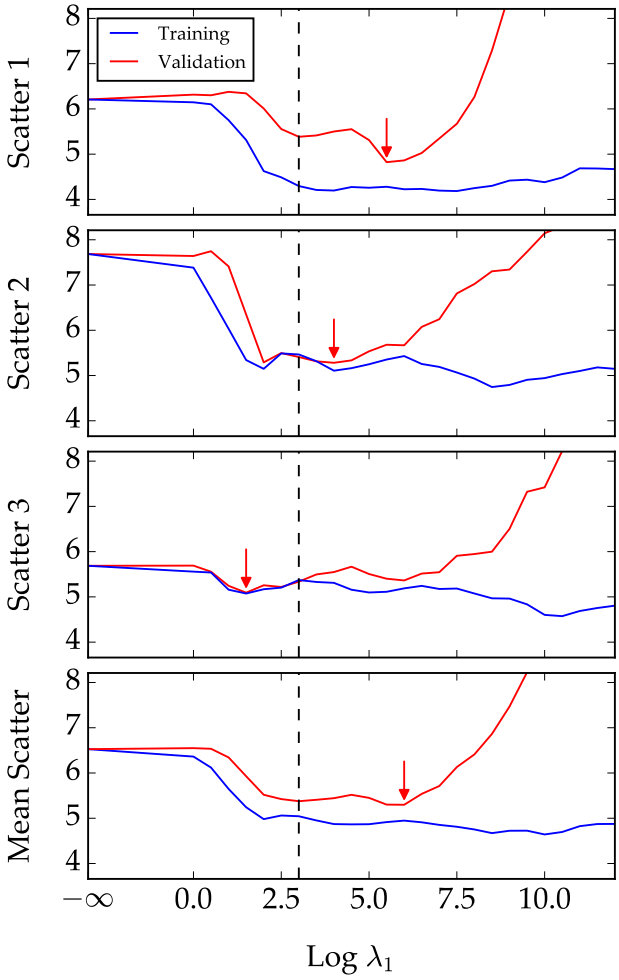

Figure 49: Cross-validation procedure for first order PLD on EPIC 206103150 (WASP-47 e), a campaign 3 planet host. Shown is the scatter σ_v in the validation set (red) and the scatter in the training set (blue) as a function of λ_1, the prior amplitude for the first order PLD weights, for each of three light curve sections; the mean scatter is shown at the bottom. Red arrows indicate the minima in the σ_v curves for each section; note that because of variable noise properties across the campaign, they all occur at different values of λ_1. The dashed vertical line indicates the value of $\hat{\lambda}_n$ obtained by the procedure outlined in the text, which establishes a compromise between slight underfitting in the first two segments and slight overfitting in the third.

ter in the validation set; this is a far more expensive calculation to do. While I could employ a nonlinear optimization algorithm (see below), in the interest of computational speed, I perform a simplification. Since I expect the first order PLD regressors to contain most of the de-trending information, with each successive PLD order providing a small correction term to the fit, I break down the minimization problem into three separate one-dimensional problems. First, I perform cross-validation on the first order PLD model by setting $\lambda_2 = \lambda_3 = 0$ to obtain the value of λ_1 that minimizes the validation scatter, $\hat{\lambda}_1$. I do this by computing the model for each value of λ_1 in a logarithmically-spaced grid with 36 points in the range $[10^0, 10^{18}]$, plus $\lambda_1 = 0$, and select the minimum (details below). I then repeat this process on the second order model by fixing λ_1 at this estimate and keeping $\lambda_3 = 0$. Finally, I solve for $\hat{\lambda}_3$ by fixing the first and second order parameters at their optimum values:

$$\begin{aligned}
\hat{\lambda}_1 &= \arg\min \sigma_{\mathrm{v}}(\lambda_1)\Big|_{\lambda_2=0,\ \lambda_3=0} \\
\hat{\lambda}_2 &= \arg\min \sigma_{\mathrm{v}}(\lambda_2)\Big|_{\lambda_1=\hat{\lambda}_1,\ \lambda_3=0} \\
\hat{\lambda}_3 &= \arg\min \sigma_{\mathrm{v}}(\lambda_3)\Big|_{\lambda_1=\hat{\lambda}_1,\ \lambda_2=\hat{\lambda}_2}
\end{aligned} \tag{149}$$

It is important to note that there is no *a priori* reason that this method should yield the global minimum of σ_{v}; in fact, it very likely does not. However, I explicitly allow for $\lambda_n = 0$ in my grid search, and thus this approximation cannot lead to overfitting, as it will always prefer a lower-order PLD model to one with higher scatter in the validation set.

As a proof of concept, I de-trended a sample of 2,700 randomly selected campaign 6 targets with `nPLD` using this approximation in the cross-validation step. I then repeated the de-trending by solving for the $\hat{\lambda}_n$ using Powell's method, initializing the solver at different points in the vicinity of the `nPLD` solution and keeping the solution with the lowest average 6 hr CDPP (combined differential photometric precision; Christiansen et al., 2012) for each target; I dub this method `pPLD`. In Figure 50 I plot the star-by-star CDPP difference between the two models, $(\mathrm{CDPP}_{\mathrm{pPLD}} - \mathrm{CDPP}_{\mathrm{nPLD}})/\mathrm{CDPP}_{\mathrm{nPLD}}$. While for some stars the CDPP improves substantially with `pPLD`, cross-validating with Powell's method leads to a less than one percent improvement in the CDPP on average. Given that this method is more computationally expensive, I adopt the grid search method outlined above when producing the `EVEREST 2.0` catalog.

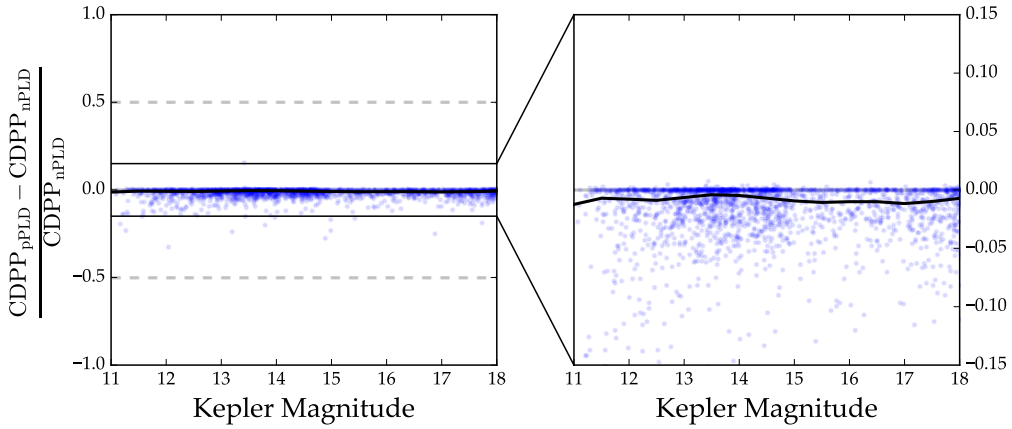

Figure 50: 6 hr CDPP comparison between de-trending with `pPLD` and de-trending with `nPLD` for a sample of 2,700 randomly selected campaign 6 stars. Plotted is the star-by-star difference in the CDPP values for each method, normalized to the `nPLD` CDPP (blue dots); stars with negative values have lower CDPP when de-trended with `pPLD`. The black line is the median CDPP difference in 0.5 magnitude-wide bins. `pPLD` leads to an average improvement in the CDPP of $\lesssim$1%.

In addition to being more computationally tractable, there are two major benefits to minimizing σ_v in this fashion. First, since I perform cross-validation three times (once for each PLD order), I am able to progressively refine the outlier masks (§6.3.1) and the GP hyperparameters (§6.3.2) for each target in between cross-validation steps. Second, it allows for some leeway in how I determine the minimum validation scatter. In Chapter 5, I sought to minimize the *median* scatter in groups of random 13-cadence segments of the light curve (the validation set). A potential issue with this method is that the noise properties of $K2$ light curves are far from constant over the course of an observing campaign; optimizing the regression based on the median (or mean) validation scatter can still, in principle, lead to overfitting in some segments. While splitting the light curves into segments with similar noise properties (§6.3.3) helps with this, I also modify the cross-validation process to prevent localized overfitting. For each PLD order n and for each value of λ_n, I split each light curve segment into three roughly equal sections. For each pair of sections, I train the model on them and compute the model prediction in the third section (the validation set). I then compute the scatter σ_v as the median absolute deviation of the de-trended validation set after removing the GP prediction.

There are now three $\sigma_v(\lambda_n)$ curves, one for each section. In general, the minima of these curves will occur at different values of λ_n, so determining the optimum value $\hat{\lambda}_n$ requires a compromise

between overfitting and underfitting in the different segments. For each segment, I compute the minimum scatter, find the set of all λ_n for which $\sigma_v(\lambda_n)$ is within 5% of the minimum, and keep the largest λ_n. I then pick $\hat{\lambda}_n$ to be the *smallest* of these values, provided it is smaller than the largest value of λ_n at the minima of the three σ_v curves. This process ensures that $\hat{\lambda}_n$ falls between the minima of the σ_v curves with the smallest and largest value of λ_n, and that it leads to no more than 5% overfitting in one of the segments. I illustrate this procedure in Figure 49, where I show $\sigma_v(\lambda_1)$ for each of the three light curve sections for EPIC 206103150. The red arrows indicate the minimum of each of the curves, and the dashed vertical line indicates the adopted $\hat{\lambda}_n$ based on a compromise between slight underfitting in the first two segments and slight overfitting in the third. This results in a more conservative cross-validation process than in Chapter 5.

6.4 Results

I de-trended all campaigns 0–8 stars with `nPLD` to produce the `EVEREST 2.0` catalog. Below I report my results, starting with injection/recovery tests and a comparison of `rPLD` and `nPLD`, followed by comparisons with other pipelines and the original *Kepler* light curves. I report most of my results in terms of the proxy 6 hr CDPP of the de-trended light curves, which I calculate in the same way as I did in Chapter 5: I smooth the light curves with a Savitsky-Golay filter, clip outliers at 5σ, and compute the median standard deviation in 13-cadence segments, normalized by $\sqrt{13}$.

6.4.1 Injection Tests

As in Chapter 5, I perform simple transit injection/recovery tests to ensure my model is not overfitting. For the same sample of 2,700 campaign 6 stars as before, I injected synthetic transits of varying depths at the raw pixel level and attempt to recover them after de-trending with `nPLD`. I follow the exact same procedure as in §4.1 of Chapter 5 and plot the results in Figure 51 (compare to Figure 6 in Chapter 5).

Each panel displays two histograms: a blue one, showing the number of transits recovered with a certain depth after de-trending with `nPLD`, and a red one, corresponding to a control run in which the transits were injected into the already de-trended light curve. Each row corresponds to a different injection depth D_0 (10^{-2}, 10^{-3}, and 10^{-4}, from top to bottom), and the x axis in each histogram is the recovered depth D scaled to this value (D/D_0). The left column corresponds to

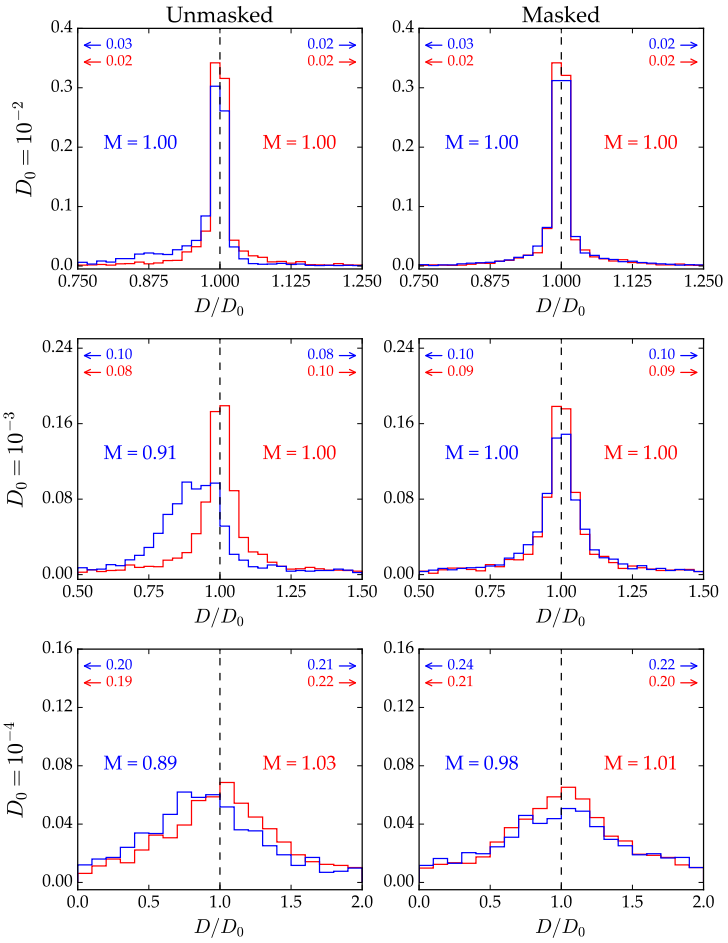

Figure 51: Transit injection/recovery statistics based on 2,700 randomly selected stars from campaign 6. Each panel shows histograms of the number of transits recovered with a certain depth ratio D/D_0 (recovered depth divided by true depth). Blue histograms correspond to the actual injection and recovery process, in which transits are injected into the raw light curves at the pixel level and recovered after de-trending with `EVEREST`; red histograms correspond to control runs in which the transits were injected into the *de-trended* data. The values to the left and right of each histogram are the median D/D_0 for my pipeline and for the control run, respectively. The smaller values at the top indicate the fraction of transits recovered with depths lower and higher than the bounds of the plots. Finally, the two columns distinguish between runs in which the transits were not explicitly masked prior to de-trending (left) and runs in which they were (right), while the three rows correspond to different injected depths: 10^2, 10^3, and 10^4. `EVEREST` preserves transit depths if the transits are properly masked; otherwise, a $\sim$10% bias toward smaller depths is introduced for transits with low SNR.

runs in which the transits were not explicitly masked during de-trending; the right column shows runs in which they were.

As with the previous version of the pipeline, I find a $\sim$10% bias toward smaller depths for low SNR transits when the transits are not explicitly masked. This is because a small decrease in the transit depth can greatly improve the CDPP of the light curve. Because the PLD regressors are noisy, the method is capable of partially fitting out transits by exploiting linear combinations of white noise in the regressors. This overfitting does not occur for high SNR transits because these are masked during the outlier clipping step.

Conversely, when transits are explicitly masked, there is no bias in the recovered depth; the median D/D_0 is consistent with unity for all three values of the injected depth. This is the same result I obtained with `EVEREST 1.0`, and I conclude that my new cross-validation scheme is robust in preventing overfitting when transits are masked. As before, I urge those using `EVEREST` light curves containing transits or eclipses to re-compute the model with those features masked. This process is quick and straightforward — refer to the `EVEREST 2.0` documentation for details. [15]

6.4.2 `rPLD`

In Figure 52 I plot a comparison of the CDPP values obtained with `rPLD` and `EVEREST 1.0` for my sample set of 2,700 campaign 6 stars. As before, the y axis corresponds to the normalized relative CDPP of each model, with negative values corresponding to lower CDPP for `rPLD`. Each star is plotted as a blue dot and the median relative CDPP is indicated as a black line. `rPLD` outperforms `EVEREST 1.0` at all Kepler magnitudes by $\sim 1-6\%$ on average. However, the scatter at any value of Kp is quite large, and the two models are roughly comparable for bright stars. As I argued in §6.3.7, the most important feature of `rPLD` is its increased robustness to local overfitting (see §6.5).

6.4.3 `nPLD`

The greatest improvement in the CDPP comes when neighboring stars' PLD vectors are included in the design matrix. In Figure 53 I plot the CDPP comparison between `nPLD` and `EVEREST 1.0`. `nPLD` outperforms regular PLD by $\sim 10-20\%$ on average, with the largest improvement occurring for fainter stars. Faint stars have the noisiest PLD vectors and benefit the most from the inclusion of higher SNR regressors. As I showed in Chapter 5, regular PLD already comes close to recovering

[15] `https://github.com/rodluger/everest`

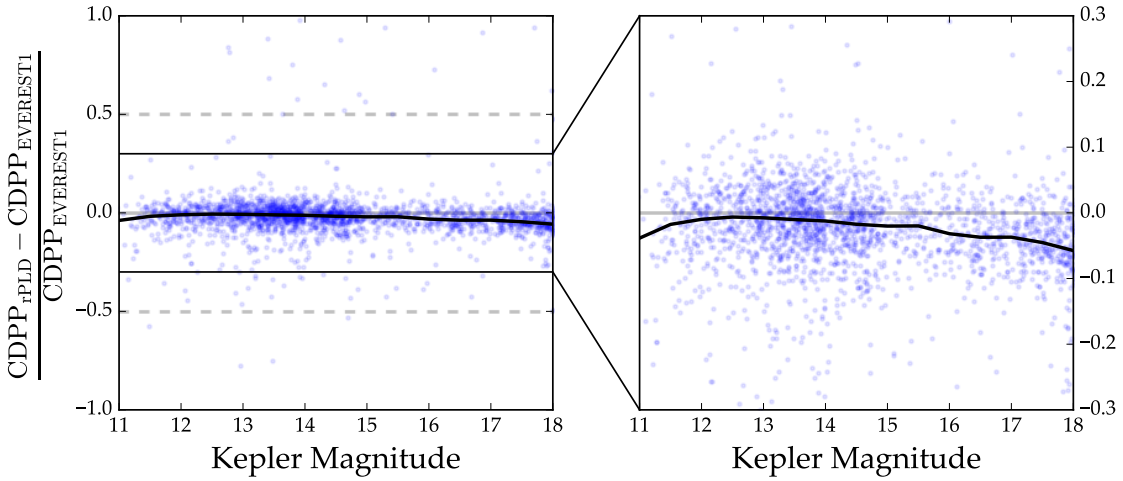

Figure 52: 6 hr CDPP comparison between de-trending with regularized regression (`rPLD`, this chapter) and de-trending with PCA (Chapter 5) for a sample of 2,700 randomly selected campaign 6 stars, as in Figure 50. Regularized regression leads to a small CDPP improvement of ∼1 to 5%.

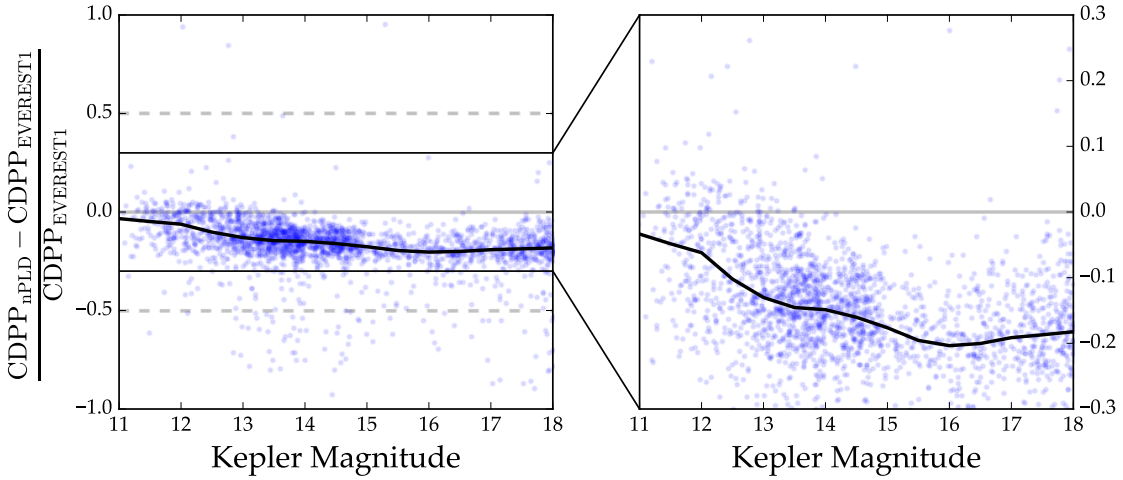

Figure 53: 6 hr CDPP comparison between PLD de-trending with regularized regression + neighboring targets (this chapter) and standard PLD de-trending (Chapter 5) for the same sample of stars as in Figure 52. Each target was de-trended with its own PLD vectors plus those of ten random bright stars on the same module. This method leads to a robust CDPP improvement of ∼10% for bright ($Kp \lesssim 13$) stars and ∼20% for fainter stars.

Kepler photometric precision for bright ($Kp \lesssim 13$) stars, so the improvement for these stars is naturally smaller.

6.4.4 Comparison to Other Pipelines

In this section I compare the `EVEREST 2.0` catalog to those produced by other pipelines, beginning with the previous version of `EVEREST`.

Figure 54 shows the CDPP comparison between `EVEREST 2.0` and `EVEREST 1.0` for all campaigns 0–8 stars. As in the example shown in Figure 53, `EVEREST 2.0` outperforms `EVEREST 1.0` by ∼20% for the faintest stars and by ∼10% for $Kp \gtrsim 12$. For $11 \lesssim Kp \lesssim 12$ the two pipelines yield comparable results, though regularized regression gives `EVEREST 2.0` a slight edge. Below $Kp \approx 11$, $K2$ stars become saturated; these are plotted as red dots, and the median relative CDPP for saturated stars is indicated by the dashed red line. For these stars, `EVEREST 1.0` yields much lower CDPP — for $Kp \lesssim 10$, the CDPP is over a factor of two smaller than that of `EVEREST 2.0`. As I discussed in Chapter 5, the increased performance of `EVEREST 1.0` for saturated stars is spurious, since the astrophysical information content of the pixels is highly variable across the aperture, leading regular PLD to overfit. As I showed in §6.3.5, `EVEREST 2.0` does not overfit saturated stars. In §6.4.9 below, I show that I approximately recover the *Kepler* photometric precision for these stars.

In Figure 55 I show the CDPP comparison between `EVEREST 2.0` and `K2SFF` (Vanderburg, 2014; Vanderburg and Johnson, 2014). My pipeline yields lower average CDPP at all magnitudes, by 40% for faint stars and 10–20% for bright stars. For saturated stars, `EVEREST 2.0` outperforms `K2SFF` by 5–10%.

In Figure 56 I show the comparison to the `K2SC` PDC light curves (Aigrain et al., 2015, 2016). `EVEREST 2.0` yields lower average CDPP at all magnitudes $K_p \gtrsim 9$, with a 20–25% improvement for all unsaturated stars. For $K_p \approx 9$, `K2SC` slightly overperforms `EVEREST 2.0`, but the scatter in the plot is quite large.

I also computed the relative CDPP for each of the campaigns individually; these are plotted in Figures 57–59. The improvement over `EVEREST 1.0` is approximately the same for all campaigns despite significant differences in the noise properties and stellar populations across the nine campaigns, showcasing the robustness of the `nPLD` method. The same is true when compared to `K2SFF`

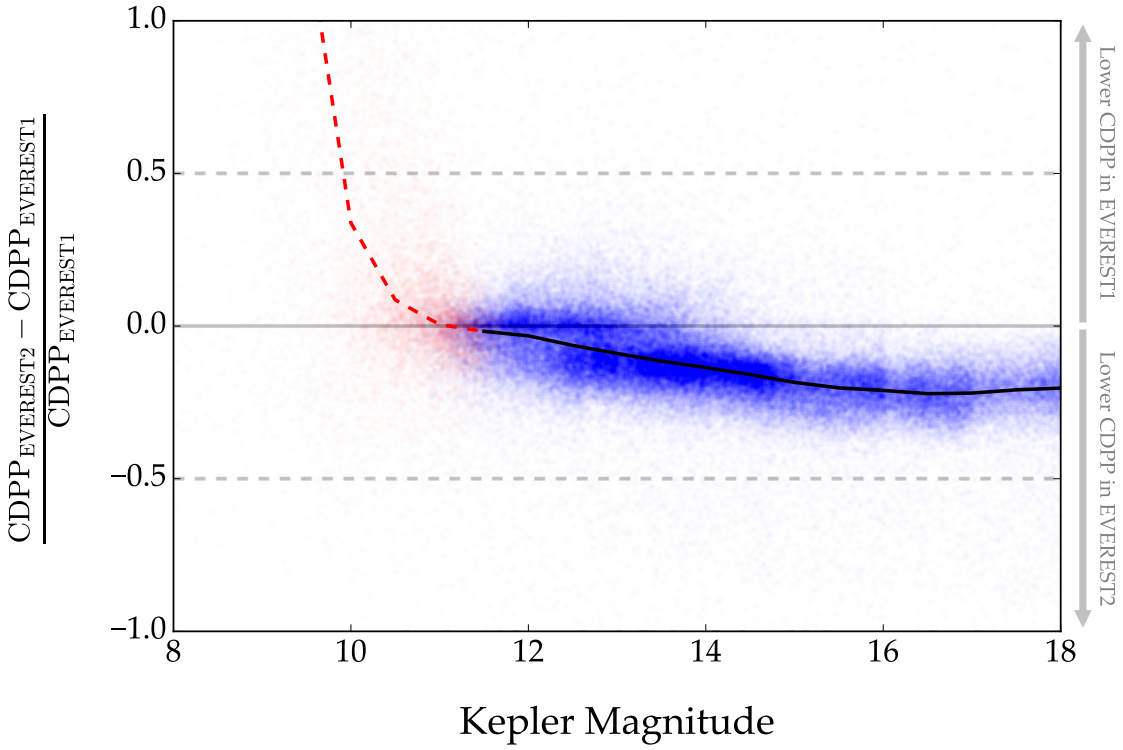

Figure 54: CDPP comparison between `EVEREST 2.0` (using `nPLD`) and `EVEREST 1.0` for all stars in campaigns 0–8. As before, individual stars are plotted as blue points and the median CDPP is indicated by a black line; note the $\sim 10 - 20\%$ improvement over the previous version of the pipeline. Saturated stars are plotted as red points, with their median CDPP indicated by a dashed red line. The apparently better performance of `EVEREST 1.0` for these stars is spurious, since traditional PLD typically leads to strong overfitting of saturated stars (see Figure 48).

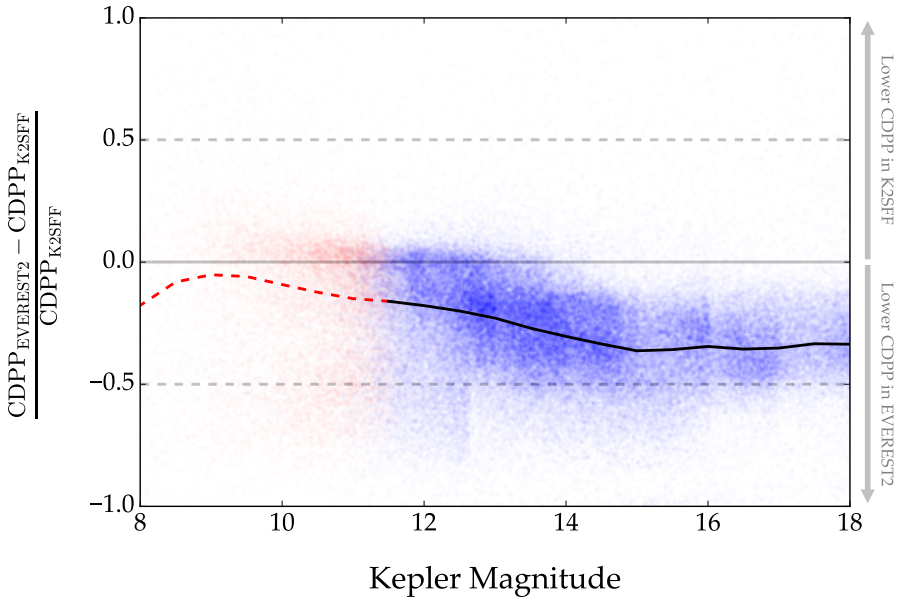

Figure 55: Similar to Figure 54, but showing a comparison between `EVEREST 2.0` and `K2SFF`. `EVEREST 2.0` outperforms `K2SFF` at all magnitudes, including $Kp \lesssim 11$, for which stars are saturated.

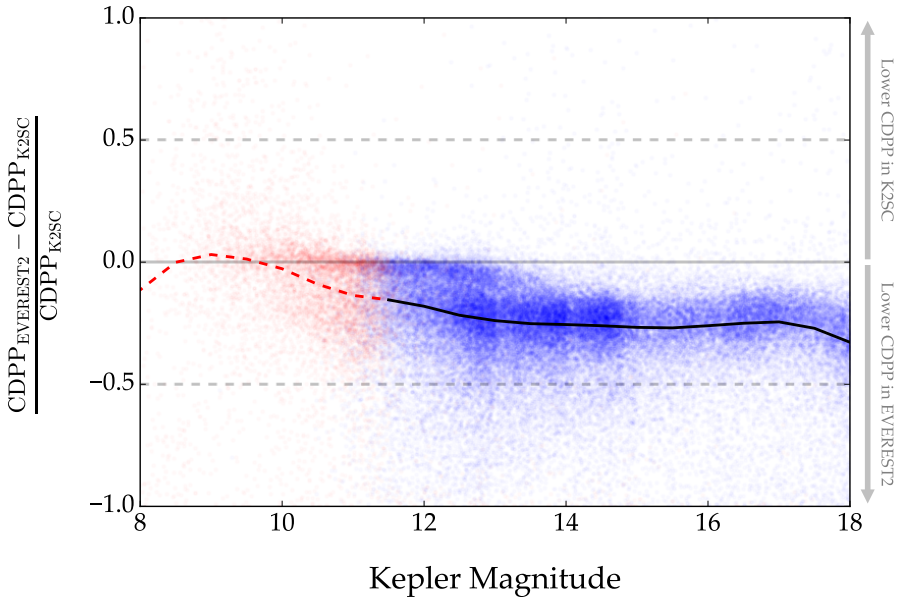

Figure 56: Similar to the previous figure, but showing a comparison between `EVEREST 2.0` and `K2SC`. `EVEREST 2.0` light curves have lower average CDPP at all magnitudes except around $Kp \approx 9$, for which the precision is comparable.

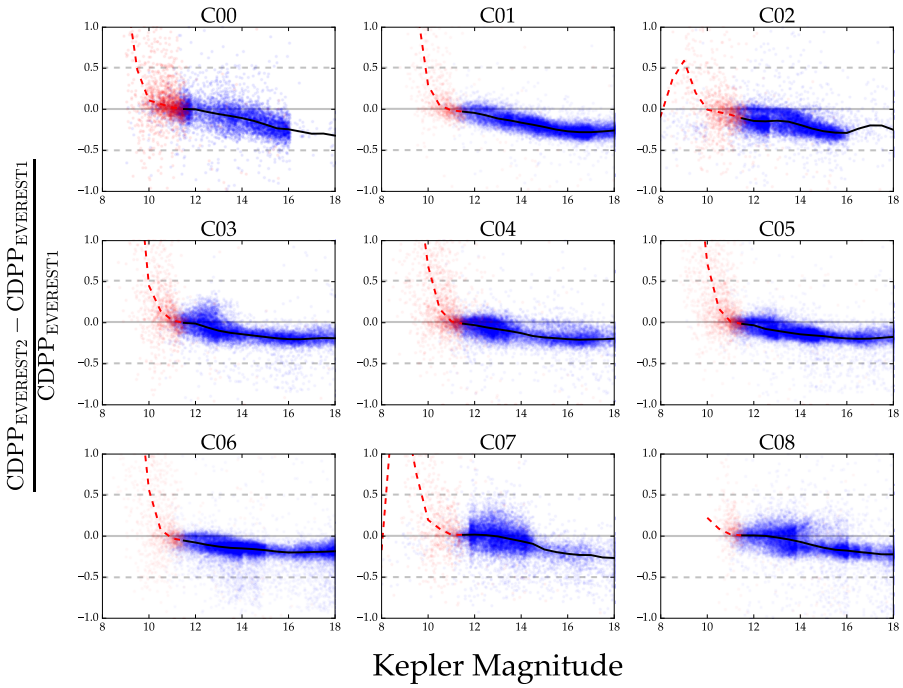

Figure 57: Similar to Figure 54, but showing a CDPP comparison between `EVEREST 2.0` and `EVEREST 1.0` for each of the first 9 $K2$ campaigns.

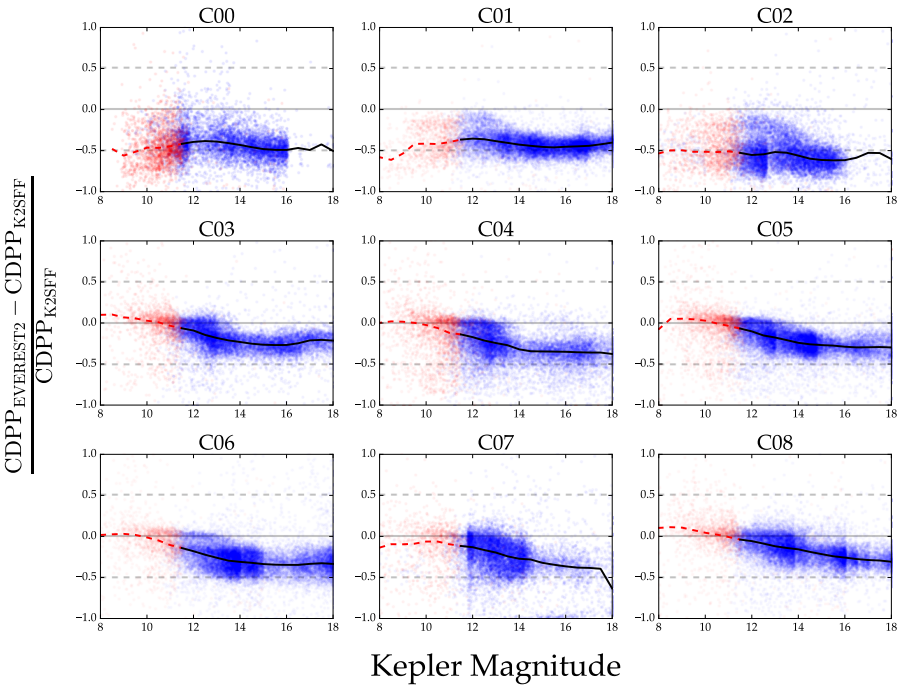

Figure 58: Similar to Figure 55, but showing a CDPP comparison between `EVEREST 2.0` and `K2SFF` for each of the first 9 $K2$ campaigns.

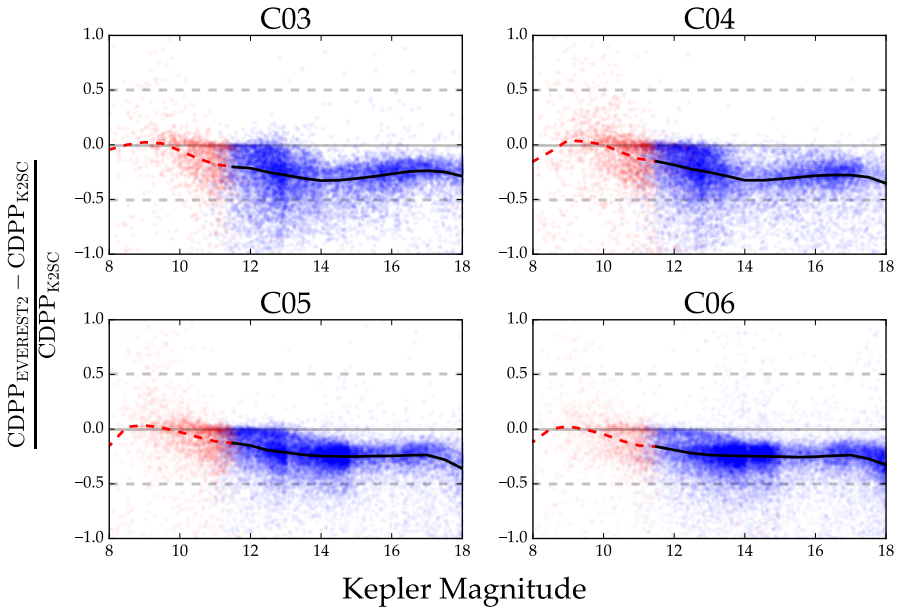

Figure 59: Similar to Figure 56, but showing a CDPP comparison between `EVEREST 2.0` and `K2SC` for each of the first 9 $K2$ campaigns.

and K2SC, with the exception of campaigns 0–2, for which EVEREST 2.0 outperforms K2SFF by nearly 50% (i.e., a factor of 2) at all magnitudes.

6.4.5 Outliers

In addition to yielding lower average CDPP than any other publicly available pipeline, EVEREST yields the largest number of *usable* data points for any of the $K2$ campaigns to date. These data points generally appear as outliers even in the raw data, since they are highly sensitive to inter- and intra-pixel sensitivity variations. Pipelines that regress on functions of the spacecraft motion alone therefore have trouble de-trending them, resulting in $\sim 5-10\%$ of the data points being discarded as outliers. In contrast, since PLD uses regressors containing both inter- and intra-pixel sensitivity information, it naturally de-trends data collected during thruster firing events (see Chapter 5).

To show this, I calculated the number of non-outlier data points per campaign for each of the four pipelines (K2SFF, K2SC, EVEREST 1.0, and EVEREST 2.0). After removing all data points with flagged QUALITY bits 1-9, 11-14, and 16-17, I smoothed each light curve with a Savitsky-Golay filter and performed iterative sigma clipping to remove all 5σ outliers. In Figure 60 I plot histograms of the number of remaining data points for all stars in each of the first 9 campaigns. As expected, both EVEREST 1.0 and EVEREST 2.0 have, on average, 100–300 more usable data points than the other two pipelines. This roughly corresponds to the number of thruster firings per campaign, as these happen every 6–12 hours on average.

6.4.6 Short Cadence

In Figure 61 I plot the relative CDPP distribution of 671 light curves that were observed in both short and long cadence. In order to compute the 6 hr CDPP of short cadence light curves, I first mask outliers and then down-bin to long cadence by taking the mean of every 30 cadences; I then compute the CDPP as usual. I achieve higher average precision in the short cadence light curves by 5-10% for unsaturated stars and by up to 25% for saturated stars. The higher information content in the short cadence light curves — particularly at sub-30 minute timescales — allows EVEREST to better de-trend those stars. I show an example in Figure 62, where I plot the light curves for EPIC 201601162 (raw short cadence at the top, de-trended short cadence in the center). In the bottom panel, I show both the down-binned short cadence de-trended light curve (black)

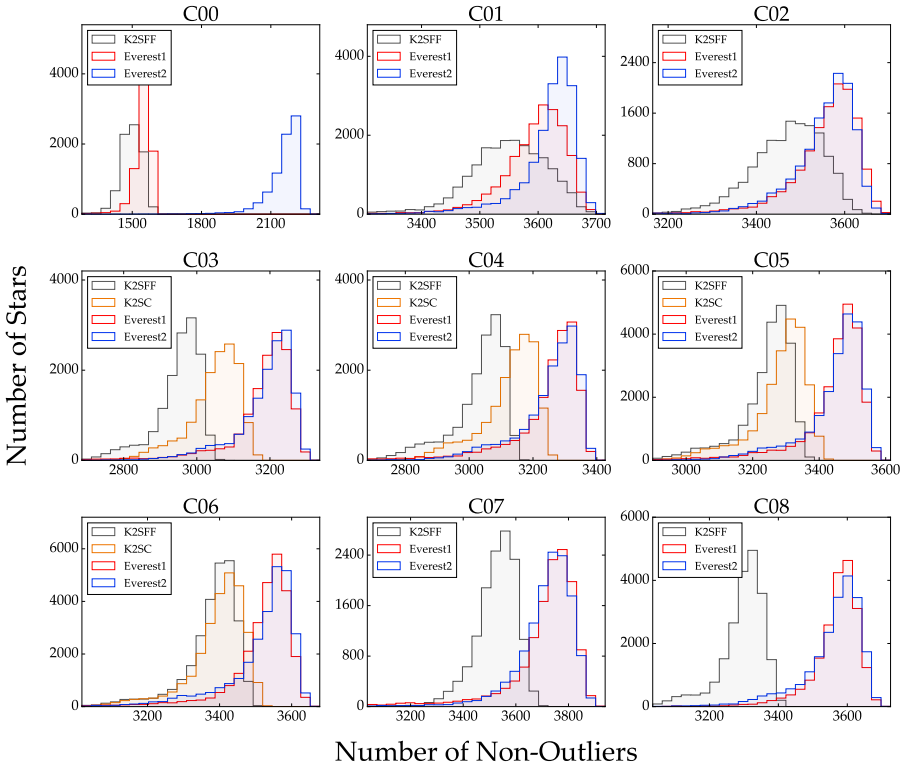

Figure 60: Histograms showing the number of non-outlier data points per campaign for each of four pipelines: `K2SFF` (gray), `K2SC` (orange; campaigns 3–6 only), `EVEREST 1.0` (red), and `EVEREST 2.0` (blue). To compute these, I remove all cadences with flagged `QUALITY` bits (excepting thruster fires) from all light curves, then smooth each light curve with a second order, 2-day Savitsky-Golay filter and perform iterative sigma clipping at 5σ to remove the outliers. The number of remaining cadences in each light curve is then used to plot the histograms. Both versions of `EVEREST` have more usable data points per campaign than the other pipelines. On average, `EVEREST` light curves have $\sim 200 - 300$ more non-outlier data points than `K2SFF` and ~ 100 more than `K2SC`.

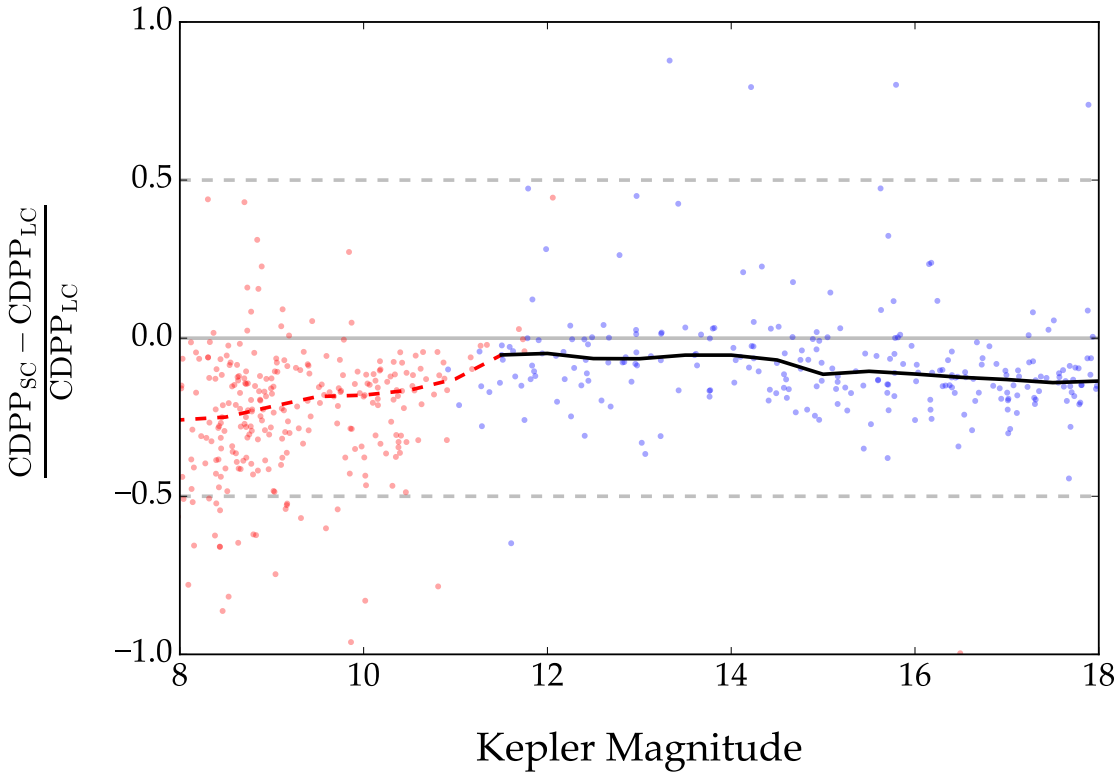

Figure 61: CDPP comparison between `EVEREST 2.0` light curves observed in short cadence and long cadence modes. Short cadence light curves have 5-10% lower CDPP on average; for saturated stars, short cadence light curves have up to 25% lower CDPP.

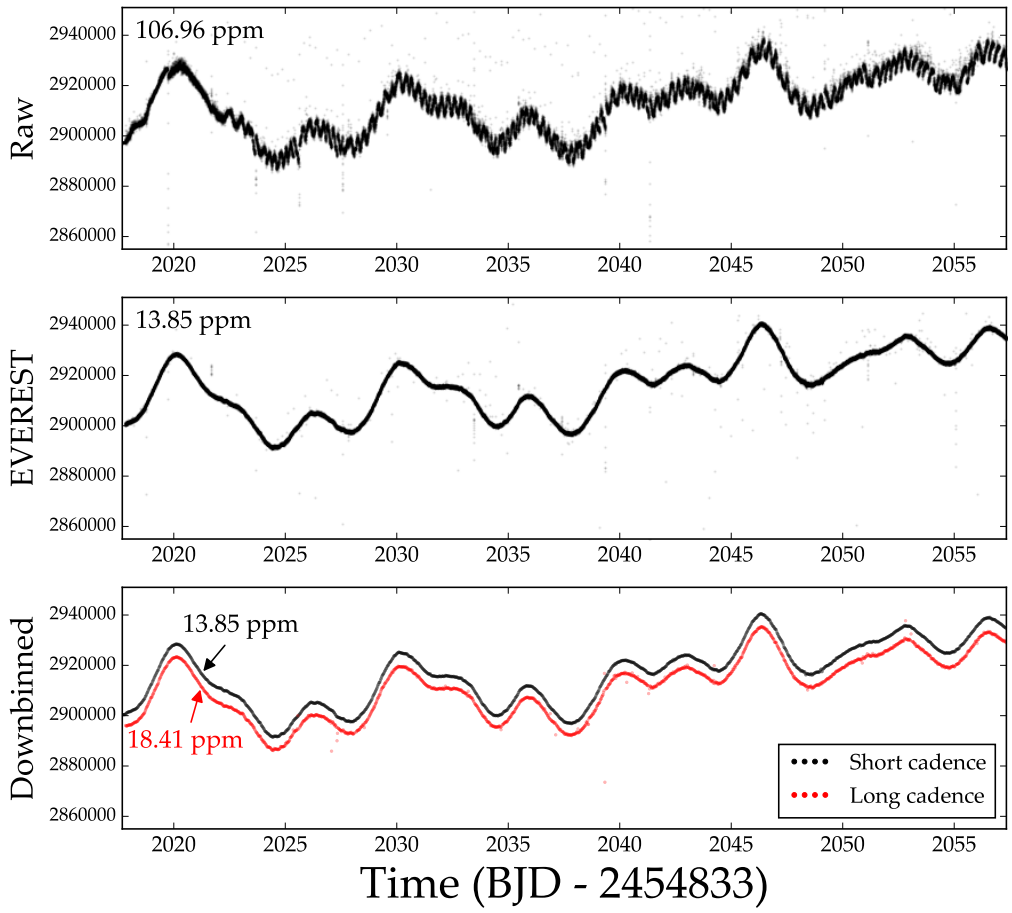

Figure 62: EPIC 201601162, a campaign 1 star observed in both long cadence and short cadence modes. A portion of the raw light curve is displayed at the top, and the de-trended light curve is shown in the center. In the bottom panel, I plot the down-binned de-trended short cadence light curve (black) and the de-trended long cadence light curve (red). The y axis in each panel is the flux in e^-/s. Short cadence `EVEREST 2.0` light curves have lower CDPP than their long cadence counterparts.

and the long cadence de-trended light curve (red). The CDPP in the short cadence light curve is ∼30% smaller than that of the long cadence light curve.

As I discussed in §6.3.6, one pitfall of my method for de-trending short cadence targets is the large number of breakpoints (∼30) I introduce in the light curve when computing the model. Overcomputing the model into adjacent segments and aligning the models at the breakpoints works very well for high SNR light curves but can often fail for light curves that are dominated by photon noise. This is the case for very faint stars ($Kp \gtrsim 17$) observed in short cadence mode, such as EPIC 201831393, which displays visible discontinuities at many of the breakpoints. Low SNR short cadence light curves may also have segments with varying white noise amplitudes due to differences in the value of the $\mathbf{\Lambda}$ prior on the PLD weights. PLD is known to perform poorly when the white noise dominates (Deming et al., 2015); in these cases, it is often desirable to down-bin the light curve and compute the model on a higher SNR signal, then predict onto the original short cadence data. Given the relatively small number of light curves for which this is an issue, I do not do this here.

6.4.7 Co-trending Basis Vectors

One downside of the algorithm employed by `EVEREST` is that GP regression has trouble distinguishing between low frequency stellar variability and low frequency instrumental systematics. Many of the `EVEREST` light curves display a steady rise over the course of the campaign, with hook-like features at the beginning, end, or both. While this does not affect transits or other high-frequency astrophysical signals, it could potentially lead to biases in stellar rotation studies that rely on low frequency modulation in the light curves.

I therefore run a post-processing step on all de-trended light curves to remove these residual instrumental signals. After some experimentation, I decided to use the `SysRem` algorithm (Tamuz et al., 2005), which identifies and removes signals shared by many light curves. `SysRem` is similar to PCA but allows for weighting of the input signals and is therefore better suited to dealing with light curves of varying noise properties. My approach is similar to the presearch data conditioning (PDC) algorithm of the *Kepler* pipeline, which uses co-trending basis vectors from many stars on the detector to remove common instrumental signals from light curves (Stumpe et al., 2012; Smith et al., 2012).

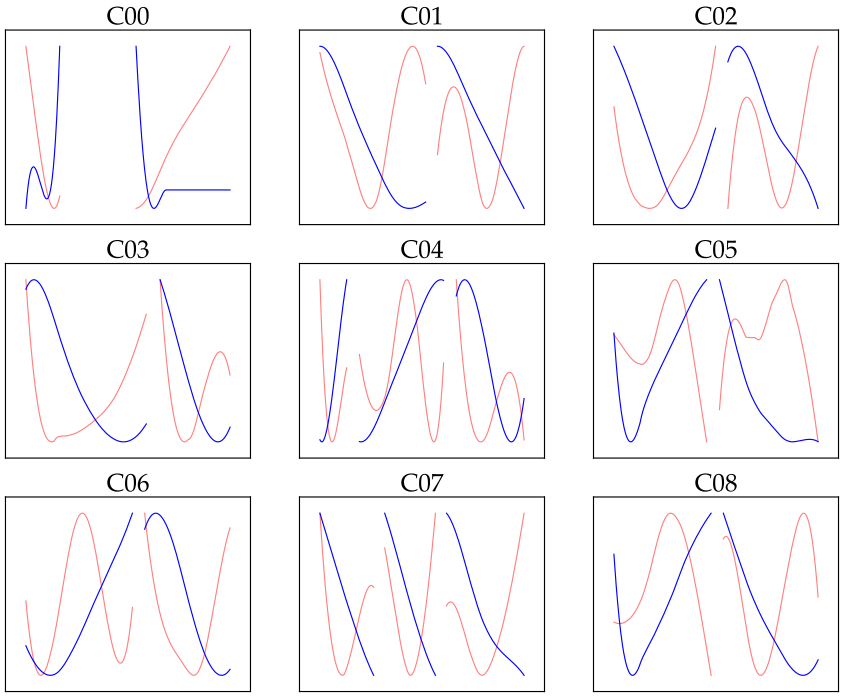

Figure 63: Co-trending Basis Vectors (CBVs) for each of the first 9 campaigns. I apply `SysRem` to all de-trended light curves in each campaign to obtain the first (blue) and second (red) CBVs; I do this independently for each of the segments in each campaign. The first set of CBVs contain primarily linear trends with hook-like features at the beginning or end of the segments; the second set of CBVs are dominated by quadratic or cubic trends. I correct all light curves by simple linear regression with the first two CBVs.

I apply `SysRem` to all of the light curves in each campaign, weighting each one by the quadrature sum of the flux measurement errors and the white noise component of its GP. Given known correlations between instrumental signals and spatial position on the detector (e.g., Petigura and Marcy, 2012; Wang et al., 2016), I attempted to apply `SysRem` individually on each CCD module, but found that the recovered signals were often dominated by astrophysical variability originating from the brightest star(s) in each module. I found that computing the `SysRem` signals for the entire detector alleviated this issue without compromising the de-trending power of the method.

I separately apply `SysRem` to each light curve segment, obtaining one co-trending basis vector (CBV) for each segment of each campaign. I then subtract a linear fit of this CBV from each light curve and repeat the procedure to obtain additional CBVs for each segment of each campaign. In order to prevent the CBVs from fitting out or introducing high frequency signals in the light curves, I aggressively smooth them with a third order, 1000-cadence Savitsky-Golay filter. The results are shown in Figure 63. The first CBV for each segment of each campaign is plotted in blue, and the second in red. As expected, the first CBV in each segment is dominated by a linear trend with S-like hooks on either end. The second CBVs are predominantly quadratic or cubic. The remaining CBVs (not shown) are dominated by higher order trends.

For the purposes of generating the `EVEREST 2.0` catalog, I perform ordinary least squares regression to fit all de-trended light curves using only the first CBV (blue curves in the figure). I find that fitting with additional CBVs often helps to remove additional systematics — in particular the hook-like features mentioned above — but may lead to overfitting of true astrophysical variability in some light curves. I therefore include all 5 CBVs of each segment of each campaign in the `EVEREST 2.0 FITS` files so that targets may be further corrected by the user on a case-by-case basis (see §6.6).

6.4.8 Sample Light Curves

In Figures 64 and 65 I show two sample light curves de-trended with both `K2SFF` and `EVEREST 2.0`. Figure 64 shows EPIC 201345483, a faint $Kp = 15$ planet host. The CDPP of the `EVEREST` light curve (bottom) is a factor of 2.4 lower than that of `K2SFF` (top), and the light curve has visibly fewer outliers. The folded transit is shown at the right and is similarly less noisy in the `EVEREST` light curve.

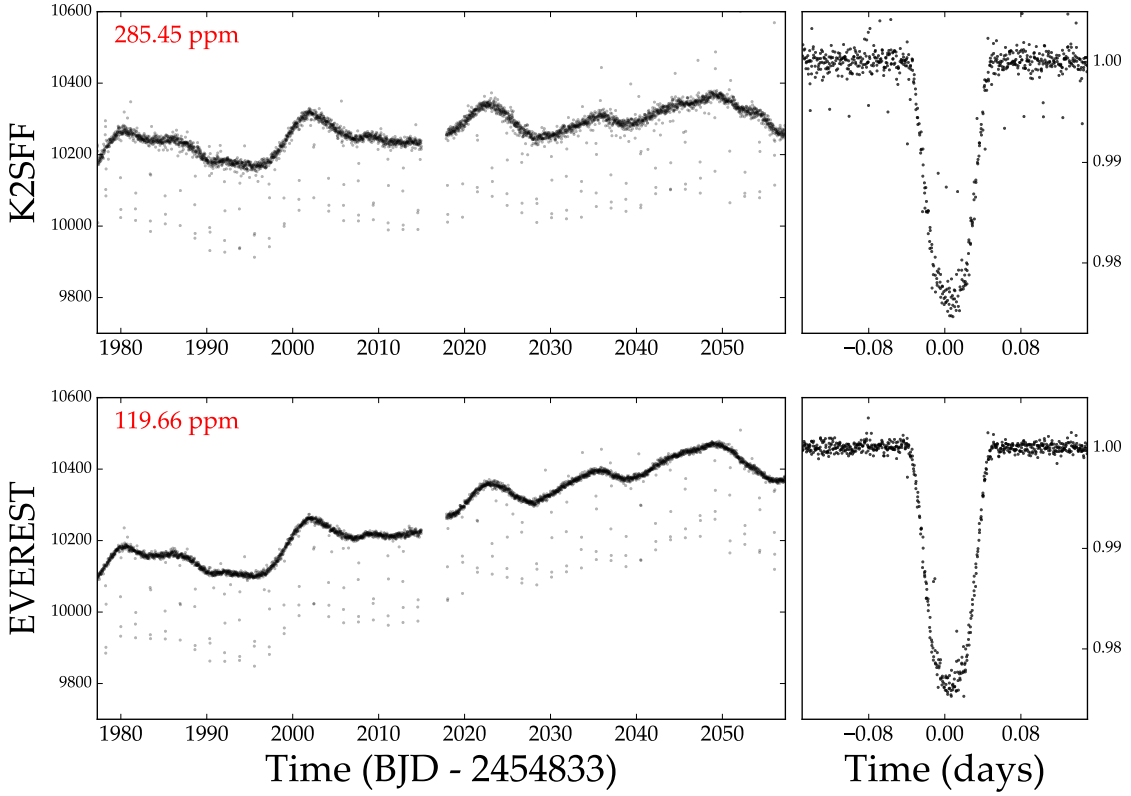

Figure 64: EPIC 201345483 (K2-45), a $Kp = 15$ campaign 1 planet host de-trended with `K2SFF` (top) and `EVEREST 2.0` (bottom). The CDPP of each light curve is indicated in the top left. The folded transit of K2-45b is shown at right. The `EVEREST 2.0` light curve has 2.4× higher photometric precision.

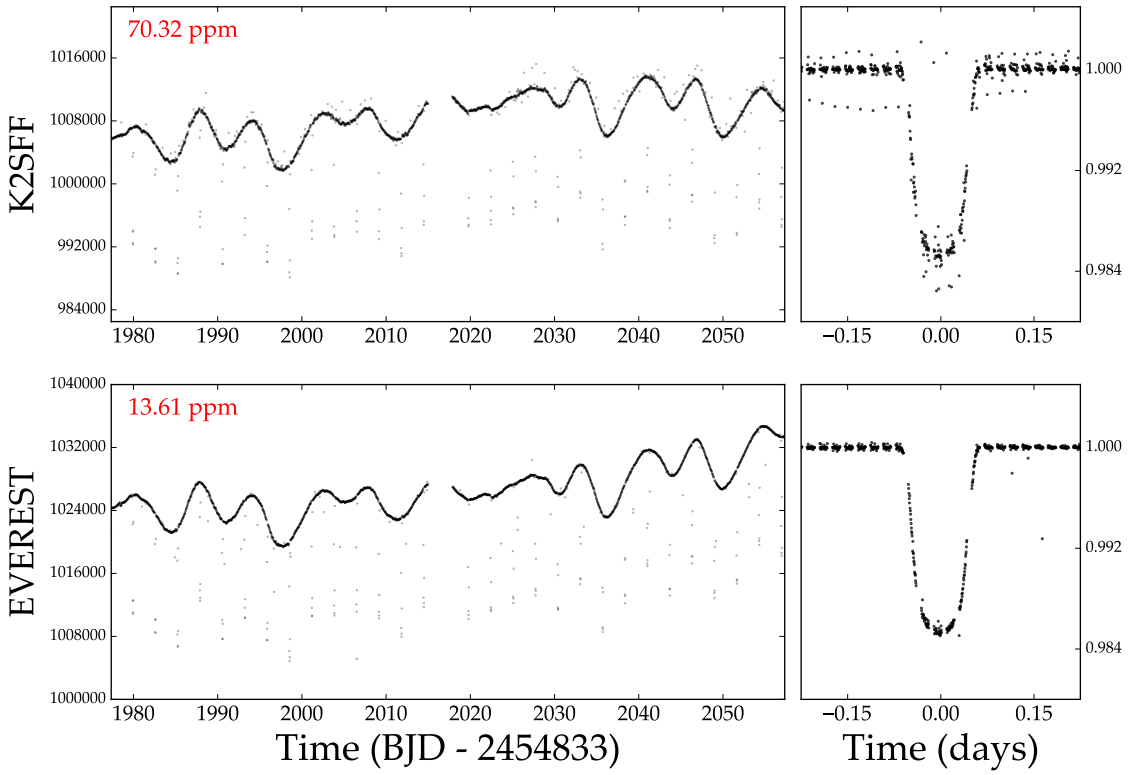

Figure 65: EPIC 201862715, a saturated campaign 1 planet candidate host. As in Figure 64, I show both the `K2SFF` and the `EVEREST 2.0` light curves. The CDPP of the `EVEREST 2.0` light curve is a factor of 5 lower.

In contrast, Figure 65 shows a very bright ($Kp = 10$) saturated planet candidate host, EPIC 201862715. With the column-collapsing scheme and the inclusion of neighboring PLD vectors, `EVEREST 2.0` is able to achieve a factor of 5 lower CDPP than `K2SFF`, as well as considerably fewer outliers.

6.4.9 Comparison to Kepler

In Figures 66 and 67 I compare the `EVEREST 2.0` photometric precision to that of the original *Kepler* mission. Figure 66 shows the CDPP as a function of Kp for each of the first 9 $K2$ campaigns and Figure 67 shows the comparison for all $K2$ stars. Because of differences in the raw photometric precision and in the stellar populations across the campaigns, the results are variable, but for all campaigns except 0, 2 and 7, I recover the original *Kepler* precision down to at least $Kp = 14$. For campaigns 1, 5, and 6, I recover the *Kepler* precision down to at least $Kp = 15$. This also applies to saturated stars, though in some campaigns the `EVEREST 2.0` CDPP is slightly higher for these stars. For stars dimmer than $Kp = 15$, the `EVEREST 2.0` CDPP is within a few tens of percent — or less — than that of *Kepler*.

6.5 Additional Remarks

6.5.1 Variable Stars

In Chapter 5 I discussed how `EVEREST` often fails to properly de-trend extremely variable stars, such as RR Lyrae variables and very short period eclipsing binaries, causing overfitting in many of these light curves. I attributed this to there being too much power in the GP model, which captured both the astrophysical and the instrumental variability, resulting in an improperly optimized PLD model. However, after considerable experimentation, I found that this behavior stemmed in large part from my cross-validation scheme. In Chapter 5, my cross-validation sets were 13 cadences (6.5 hours) long, and I sought to minimize the median (proxy) CDPP of all such sets. I chose this timescale because it is roughly the duration of a typical transit, and in practice it worked well to minimize overfitting of transits. However, the CDPP as defined in §6.4 — the normalized standard deviation in 13 cadence segments after the application of a high-pass filter — is not an adequate metric of the photometric precision for stars that are intrinsically variable on similarly short timescales. In other words, RR Lyrae and other extreme variables are dominated by astrophysical variability on

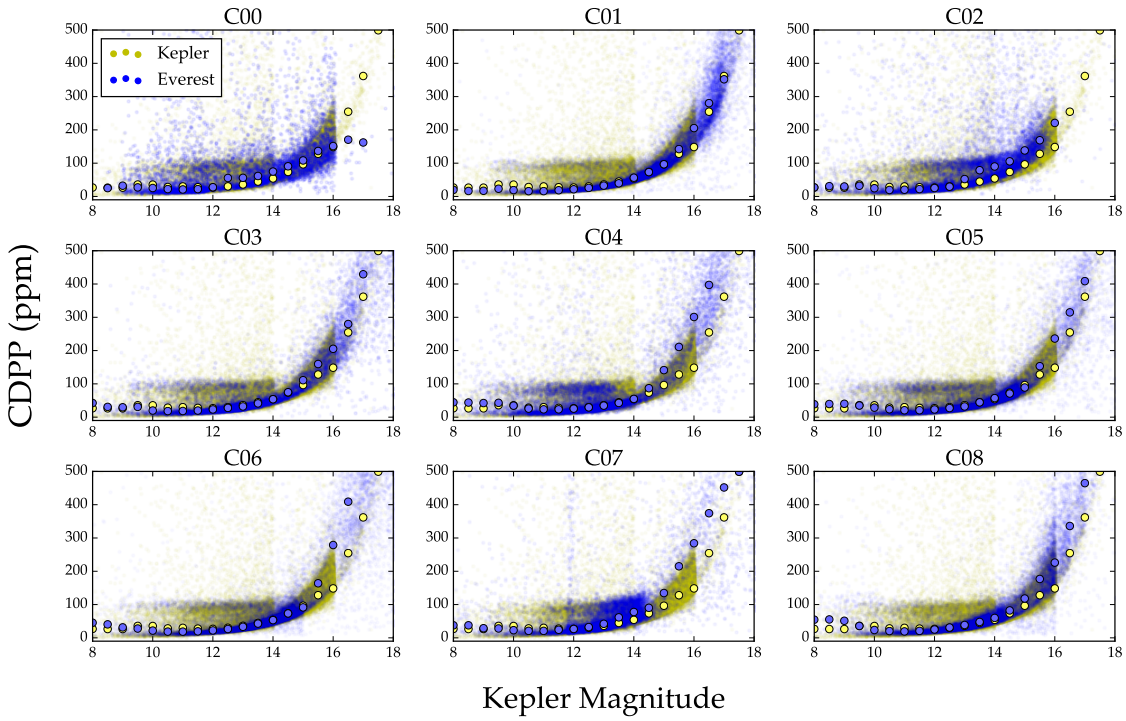

Figure 66: 6 hr photometric precision as a function of *Kepler* magnitude Kp for all stars observed by *Kepler* (yellow dots) and for all $K2$ targets in Campaigns 0-8 de-trended with `EVEREST` (blue). The median in 0.5 magnitude-wide bins is indicated by yellow circles for *Kepler* and by blue circles for `EVEREST`. For campaigns 1, 5, and 6, `EVEREST` recovers the raw *Kepler* photometric precision down to at least $Kp = 15$; for campaigns 3, 4, and 8, `EVEREST` recovers the *Kepler* precision down to $Kp = 14$. Campaigns 0 and 2 have a larger fraction of (variable) giant stars, leading to a higher average CDPP, while campaign 7 raw light curves have significantly worse precision due to a change in the orientation of the spacecraft and excess jitter.

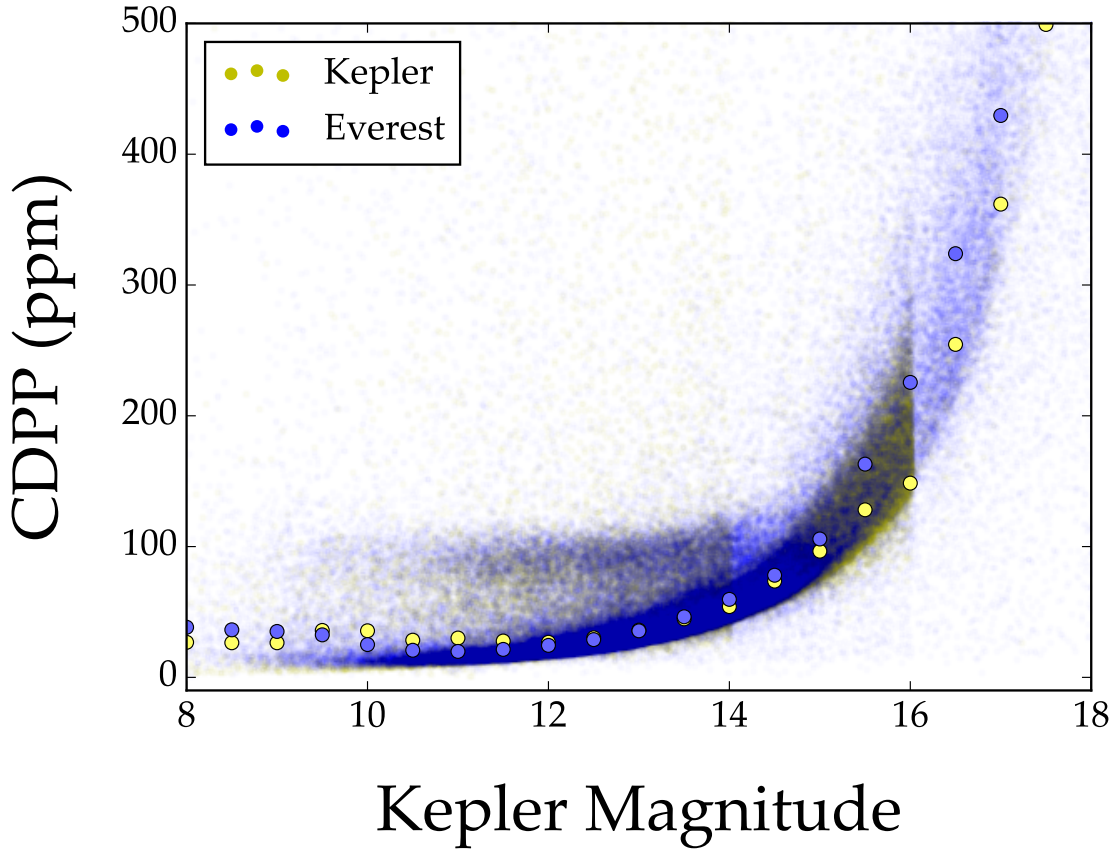

Figure 67: The same as Figure 66, but comparing the CDPP of *all* $K2$ stars to that of *Kepler*. `EVEREST 2.0` recovers the original *Kepler* photometric precision down to at least $Kp = 14$, and past $Kp = 15$ for some campaigns.

the timescale at which the CDPP is computed, and therefore minimizing the CDPP is a recipe for overfitting.

In `EVEREST 2.0`, I modified my cross-validation scheme (§6.3.7) in two important ways: I increased the average size of the validation sets to ∼500 cadences, and I minimized the validation scatter after the subtraction of a properly optimized GP model. This helps to ensure that I minimize *only* the instrumental component of the noise. In order to assess the performance of this new scheme, I visually inspected the de-trended light curves of 100 RR Lyrae stars in campaigns 0–4, chosen as the targets with the highest probability of being RRab stars according to the `K2VARCAT` catalog (Armstrong et al., 2016). Among the `EVEREST 1.0` light curves, 92/100 had visibly damped oscillation amplitudes or clearly overfitted stellar variability features. In contrast, only 44/100 `EVEREST 2.0` RR Lyrae stars showed any signs of overfitting. For 8 of these, the overfitting occurred only with the inclusion of the second or third order PLD models. While there are still issues with how the pipeline handles extremely variable stars, the improvement over version `1.0` is substantial.

For comparison, I visually inspected the same stars in the `K2SFF` and `K2SC` catalogs. 96/100 `K2SFF` light curves of RRab stars showed signs of overfitting (dampened oscillation amplitudes, distorted astrophysical signal, or significant de-trending artefacts), while only 12/65 (18%) of `K2SC` light curves appeared to be incorrectly de-trended (35 of the 100 stars were in campaigns 0–2 and are not present in the `K2SC` catalog). `K2SC` likely overperforms the other pipelines for these stars because of its robust GP optimization scheme. Since I optimize the GP and de-trend the light curve in separate steps, the covariance matrix I use is often an improper approximation to the true covariance of the astrophysical signal. Progressive optimization of the GP (§6.3.2) helps to maximize the amount of astrophysical information captured by the GP model, but a better procedure would be to simultaneously fit for the GP hyperparameters and the PLD coefficients. This would ensure that the PLD model captures only instrumental signals and the GP model captures only astrophysical signals. However, such a method is computationally intractable, since the problem would no longer be linear. I therefore settle for my linear method, which works extremely well for stars that do not exhibit extreme variability and has a ∼50% success rate for those that do. I encourage readers interested in RR Lyrae and other extreme variables to inspect the light curves of the different pipelines on a target-by-target basis.

6.5.2 Crowded Apertures

In Chapter 5 I also showed how PLD is not suited to de-trend stars in excessively crowded fields, since the algorithm will often use pixels from contaminant sources to fit out the target star's astrophysical signals. Directly addressing this issue is far less straightforward than mitigating the effects of saturation (§6.3.5) or extreme astrophysical variability (§6.5.1), as it likely requires accurate modeling and subtraction of the PSFs of contaminant sources. However, in practice I find that my new de-trending algorithm is far more robust to overfitting when contaminant sources are present. This is in part due to the inclusion of PLD regressors from neighboring stars, which reduce the contribution from potentially contaminated pixels in the target aperture. My new cross-validation scheme (§6.3.7) also helps guard against the effects of crowding, since it is better at preventing localized overfitting. Because of the large amplitude of the $K2$ spacecraft drift, the amount of contamination in the target aperture often varies considerably over the duration of the campaign, as nearby sources move in and out of the aperture. In `EVEREST 1.0` light curves, this resulted in a time-variable de-trending power — and time-variable overfitting — for crowded targets. As I discussed in §6.3.7, my new cross-validation method strongly disfavors this behavior.

In Figure 68 I show the light curves of five eclipsing binaries whose crowded apertures resulted in severe overfitting in the `EVEREST 1.0` catalog. Shown are the light curves (flux versus time) folded on the period of the binary and centered on the primary eclipse for `EVEREST 1.0` (left) and `EVEREST 2.0` (center). The right panel shows the Palomar Optical Sky Survey[16] (POSS II) red filter image of the target postage stamp, where contaminant sources are clearly visible in each of the apertures (indicated by the red contours).

The first three targets (EPIC 202072978, EPIC 202103762, and EPIC 218803648) have relatively bright contaminant sources centered on or near the edge of the aperture. In Chapter 5 I explained how this was the "worst case" scenario for crowding, as it results in the largest spatial variation of astrophysical information from the target across the aperture; as expected, `EVEREST 1.0` almost completely fits out their eclipses. `EVEREST 2.0`, on the other hand, achieves high de-trending power while preserving the eclipse shapes and depths seen in the raw light curves. The fourth target (EPIC 202733088) has a bright nearby source inside its aperture, separated by less than a pixel. I discussed in Chapter 5 how targets with co-located contaminants are typically

[16] `http://stdatu.stsci.edu/cgi-bin/dss_form`

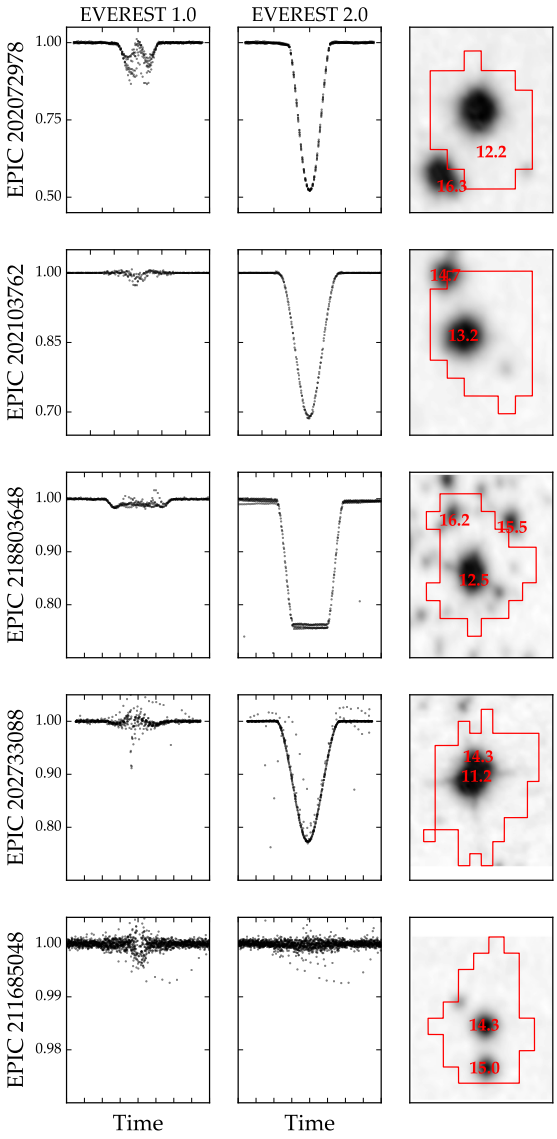

Figure 68: Five eclipsing binaries with significant contamination by bright nearby stars. The first two panels in each row show the folded `EVEREST 1.0` and `EVEREST 2.0` light curves, respectively, and the third shows the POSS high resolution image of the target postage stamp with my adopted aperture indicated in red. The magnitudes of the target and its bright neighbors are also indicated. While `EVEREST 1.0` severely overfits the eclipses of all five targets, `EVEREST 2.0` preserves the eclipse depths in all but the last one.

unaffected by PLD overfitting, since the ratio of the two PSFs is roughly constant everywhere. However, in this case the contaminant is sufficiently detached to result in severe overfitting in the `EVEREST 1.0` light curve. As before, `EVEREST 2.0` correctly de-trends the light curve with no overfitting. The final target (EPIC 211685048) is an eclipsing binary with deep transits that are completely fit out in both versions of the pipeline. In this case, two roughly equal magnitude stars are fully contained in the aperture and separated by $\sim$3 pixels, causing PLD to fail despite the modifications to the algorithm.

As in Chapter 5, I urge caution when using the `EVEREST` light curves of crowded targets, but note that in most of the cases known to us in which the previous version overfitted transits and eclipses, `EVEREST 2.0` succeeds in removing the instrumental noise without overfitting astrophysical information. Future updates to the pipeline will include a more careful aperture selection, which can mitigate the effects of crowding for targets like EPIC 211685048.

6.5.3 Aperture Losses

As I have demonstrated, column-collapsed PLD works extremely well for saturated stars, provided *all* of the target flux in the saturated columns is used. With this in mind, I used larger apertures for these stars and extended them by two pixels at the top and bottom of all saturated columns (§6.3.5). However, for some extremely bright ($Kp \lesssim 8.5$) targets, bleed trails in the saturated columns can extend past the edge of the target postage stamp. Since most of the astrophysical information in saturated columns is contained in the two pixels at the top and bottom, this can lead to substantial information loss. In this case, the column collapsing procedure fails to satisfy the PLD constraint that all pixels should have the same fractional astrophysical signal strength, and overfitting can occur. This is the case for EPIC 210703831, a campaign 4 $Kp = 8.1$ star, whose de-trended light curve displays several discontinuities. These occur because spacecraft drift results in aperture losses only during parts of the campaign. I therefore encourage those using the `EVEREST` catalog to inspect the postage stamps of extremely bright stars to ensure that aperture losses are not present.

6.6 Using EVEREST

All EVEREST 2.0 $K2$ light curves are available for download as FITS files from MAST.[17] As with the previous version, I urge users to interface with the catalog via the Python code, which can be installed following the instructions at https://github.com/rodluger/everest. A detailed description of the data products and how to use the code is available on the github page. Below, I provide a brief outline of the EVEREST resources available online.

6.6.1 FITS Files

Each FITS file contains six extensions. The primary (index 0) extension consists of a header with miscellaneous target information copied from the $K2$ target pixel file (TPF). The second (index 1) extension contains a header and a binary table. The header stores miscellaneous information about the target and the settings used in the de-trending, such as the GP hyperparameters and information on the neighboring targets used in the regression. The binary table stores arrays corresponding to the cadence number (CADN), the timestamp (TIME), the raw SAP flux (FRAW), the raw SAP flux errors (FRAW_ERR), the PLD-de-trended flux (FLUX), the five CBV regressors (CBV01 – CBV05), and the de-trended flux with the CBV correction (FCOR). This extension also includes the original $K2$ QUALITY bit array, with four additional bits that signal cadences that were masked when computing the model:

23	Data point is flagged in the $K2$ TPF
24	Data point is NaN in the $K2$ TPF
25	Data point is an outlier
26	*Not used*
27	Data point is in a transit (short cad. only)

When using the EVEREST light curves for science purposes, it is important to properly reject these outliers. Data flagged with bit 23 is typically affected by cosmic rays or detector anomalies, and can usually be ignored. In some cases, however, deep transits or eclipses can be mistaken for detector anomalies and are incorrectly flagged in the original TPF, so a visual inspection of the light curves is recommended. Data flagged with bit 24, on the other hand, is missing in the original TPF and can thus be safely ignored. Finally, bit 25 usually corresponds to instrumental outliers

[17] https://archive.stsci.edu/prepds/everest

that were not properly de-trended with PLD. However, transits, flares, and other short timescale astrophysical features will also be flagged with this bit, so this data should *not* be blindly excluded when performing transit searches.

The four remaining `FITS` extensions include the PLD regressors, the target aperture, and additional information used internally by the Python code.

6.6.2 Data Validation Summaries

Each target in the catalog also has an associated data validation summary (DVS), a `PDF` document showing the raw, de-trended, and CBV-corrected light curves, as well as cross-validation diagnostic plots such as those in Figure 49 and a high resolution POSS image of the target aperture like those shown in Figure 68.

6.6.3 Python Code

The `EVEREST` code can be installed using the package managing system `pip`:

```
pip install everest-pipeline
```

or directly from source by following the instructions on the `github` page. The primary way of interfacing with the catalog is to instantiate an `Everest` object:

```
import everest
star = everest.Everest(EPIC)
```

where `EPIC` is the EPIC number of the target. These lines download the target's `FITS` file and populate the `star` object with the de-trending information and light curve arrays (`time`, `flux`, `fcor`, etc.). The `QUALITY` flags `23`, `24`, `25`, and `27` are used to generate the `badmask`, `nanmask`, `outmask`, and `transitmask` arrays, respectively; these arrays contain the indices of all data points whose corresponding bit is flagged.

As I discussed in §6.4.1, it is important that transits are properly masked when computing the PLD model; otherwise, slight overfitting may occur. Users should therefore always mask known transits and re-compute the PLD model, as follows:

```
star.mask_planet(time, period, duration)
star.compute()
```

where `time` is the time of first transit (in units of BJD − 2454833), `period` is the transit period (in days), and `duration` is the full transit duration (also in days). Re-computing the model takes no more than a few seconds for long cadence light curves.

Users can also easily change the number of CBVs used to correct the light curve:

```
star.cbv_num = n
star.compute()
```

where `n` is the desired number of CBVs (0 – 5).

The code also implements various visualization routines, which are described in the documentation linked on the `github` page.

6.7 Conclusions

I have presented `EVEREST 2.0`, an update to the `EVEREST` pipeline (Luger et al., 2016) for removing instrumental noise from $K2$ photometry. In version `1.0`, I constructed a linear model from the principal components of products of the fractional pixel fluxes in the aperture of each star, a variant of a method known as pixel level decorrelation (PLD, Deming et al., 2015). Here, I regress on *all* PLD vectors, imposing Gaussian priors on the model weights to prevent overfitting. I additionally include the PLD vectors of bright neighboring stars to increase the signal-to-noise ratio of the regressors and enhance the predictive power of the model. I developed a fast gaussian process (GP) regression scheme to de-trend all stars in the $K2$ catalog, achieving lower combined differential photometric precision (CDPP) than in version `1.0`, by ∼10% for bright stars and ∼20% for faint stars. I also adapted PLD to work for saturated stars, yielding comparable de-trending power, and stars observed in short cadence mode, yielding higher photometric precision on 6 hr timescales than their long cadence counterparts. I further find that the inclusion of neighboring PLD vectors and a more conservative cross-validation scheme enhance the pipeline's robustness to overfitting, particularly for highly variable stars.

`EVEREST 2.0` light curves have higher photometric precision than the two other publicly available catalogs, `K2SFF` (Vanderburg and Johnson, 2014) and `K2SC` (Aigrain et al., 2016), at all Kp magnitudes. For faint stars, `EVEREST 2.0` has ∼40% lower CDPP than `K2SFF` and ∼25% lower CDPP than `K2SC`; for bright unsaturated stars, the CDPP improvement is ∼20% compared to both pipelines. For saturated stars, `EVEREST` outperforms both pipelines, but by a smaller margin.

I also find that EVEREST light curves have, on average, 100–300 fewer outliers than those of other pipelines, owing primarily to the ability of PLD to correct data collected during thruster firing events.

When compared to the original *Kepler* mission, EVEREST 2.0 recovers *Kepler* photometry on average to $Kp \approx 14.5$, and past $Kp = 15$ for some campaigns. For dimmer stars, the CDPP is within a few tens of percent of that of *Kepler*. EVEREST light curves should thus enable continued high-precision transiting exoplanet and stellar variability science for the vast majority of $K2$ stars as if they had been observed by the original four-wheeled mission.

The EVEREST 2.0 catalog of de-trended light curves is publicly available at https://archive.stsci.edu/prepds/everest. As with the previous version of the code, EVEREST 2.0 is open source under the MIT license and available at https://github.com/rodluger/everest. The reader is encouraged to use this code to interface with the EVEREST catalog and to customize the de-trending of targets of interest, particularly for masking transits to remove biases in the depth due to overfitting. I have implemented each of the PLD models discussed above (rPLD, nPLD, pPLD) in a general framework that can be adapted to different missions, including *Kepler* and the upcoming *TESS*. For more information, refer to the documentation linked on the github page.

7 Application to TRAPPIST-1

Portions of this chapter were originally published in collaboration with M., Sestovic, E. Kruse, S. L. Grimm, B.-O. Demory, E. Agol, E. Bolmont, D. Fabrycky, C. Fernandes, V. Van Grootel, A. Burgasser, M. Gillon, J. G. Ingalls, E. Jehin, S. N. Raymond, F. Selsis, A. H. M. J. Triaud, T. Barclay, G. Barentsen, S. B. Howell, L. Delrez, J. de Wit, D. Foreman-Mackey, D. L. Holdsworth, J. Leconte, S. Lederer, M. Turbet, Y. Almleaky, Z. Benkhaldoun, P. Magain, B. M. Morris, K. Heng, and D. Queloz in the June 2017 edition of Nature Astronomy (Luger et al., 2017, Nat Astron, Vol. 1, 0129; © 2017 Macmillan Publishers Limited), and are reproduced below with permission of the Nature Publishing Group.

7.1 Introduction

In the previous two chapters, I described the `EVEREST` pipeline for de-trending light curves taken by the *K2* mission. This pipeline is currently being used to detect new exoplanets, including ones in the habitable zone (Kruse et al., 2017, in prep., ; see also Chapter 4). In this chapter, I discuss in detail the application of `EVEREST` to TRAPPIST-1, the first planetary system found orbiting an ultra-cool dwarf star (Gillon et al., 2016). At least seven planets similar to Earth in radius are known to transit this star, at least three of which are in the habitable zone (Gillon et al., 2017).

Three of the planets were first discovered around TRAPPIST-1 with the *TRAPPIST* ground-based telescope (Gillon et al., 2016). The space-based infrared *Spitzer* telescope was used for intensive follow-up, detecting four additional planets and constraining the orbital properties of all but one, TRAPPIST-1h, the smallest and most distant of the seven. Only one transit of h had been seen with *Spitzer* and therefore its period was unknown. An estimate based on the single transit duration constrained the period between about 14 and 35 days, which could have placed it in the habitable zone were it at the shorter end of this range. However, because of the long (and unknown) period of TRAPPIST-1h, neither *Spitzer* nor any ground-based telescopes found additional transits of the planet.

Fortunately, the proximity of TRAPPIST-1 to the ecliptic meant that the star was potentially observable by *K2*. During Campaign 12 of the mission, TRAPPIST-1 would have been within one degree of the spacecraft's planned field of view. Following the discovery of the planetary

system (and after much debate!), the field of view was shifted to include the star (Steve Howell, personal communication). As I describe below, these observations provided the longest baseline light curve of TRAPPIST-1 to date, enabling a search for TRAPPIST-1h and for additional planets in the system. De-trending with `EVEREST` was instrumental in the search, detection, and ultimate confirmation of TRAPPIST-1h.

7.2 Observations with K2

K2 observed TRAPPIST-1 (EPIC 246199087) for 79 days as part of Campaign 12, starting on 2016 Dec 15 and ending on 2017 Mar 04. The spacecraft was in safe mode between 2017 Feb 1 and 2017 Feb 6, resulting in a 5-day data loss. Typically upon downlink from the spacecraft, the raw cadence data are calibrated with the *Kepler* pipeline (Quintana et al., 2010), a lengthy procedure that includes background subtraction, smear removal, and undershoot and nonlinearity corrections. However, given the unique science drivers in this dataset, the raw, uncalibrated data for Campaign 12 were made publicly available on 2017 Mar 8 shortly after downlink.

I used the package `kadenza` to generate a target pixel file (TPF) from the Campaign 12 raw data for TRAPPIST-1. In addition to EPIC ID 246199087, which corresponds to a standard-size postage stamp centered on TRAPPIST-1, the *Kepler* GO Office also made available a larger, 11×11 pixel custom mask with a different ID (200164267), which I used for the purposes of this study. I manually selected a rectangular 6×6 pixel (24" $\times$ 24") aperture centered on TRAPPIST-1 and used the median values of the remaining pixels to perform a column-by-column background subtraction. This process removes dark current and background sky signals while mitigating smear due to bright stars in the same CCD column as the target. Based on simple error function fits to the stellar image, I find that my aperture encloses > 0.9996 of the flux of TRAPPIST-1h throughout the entire timeseries. I performed this pre-processing step on both the long cadence ($t_{\rm exp} = 30\,{\rm min}$) and short cadence ($t_{\rm exp} = 1\,{\rm min}$) light curves.

Because of its two failed reaction wheels, rolling motion of the *Kepler* spacecraft due to torque imbalances introduces strong instrumental signals, leading to an increase in photometric noise by a factor of $\sim 3 - 5$ compared to the original mission (see Chapters 5 and 6). Since TRAPPIST-1 is a faint M8 dwarf with *Kepler* magnitude $Kp \sim 16 - 17$ (see below), these instrumental signals must be carefully removed to reach the $\sim$0.1% relative photometric precision required to detect

Earth-size transits (Demory et al., 2016b). To this end, I detrended the long cadence light curve for TRAPPIST-1 using both `EVEREST` and a Gaussian process-based pipeline, achieving an average 6-hr photometric precision of 281.3 ppm, a factor of 3 improvement over the raw light curve. After analysis of the long cadence light curve, I detrended the short cadence light curve in the vicinity of features of interest, achieving comparable or higher 6-hr precision.

7.2.1 Light curve detrending – EVEREST

The `EVEREST` *K2* pipeline (Chapters 5 and 6) uses a variant of pixel level decorrelation (Deming et al., 2015) (PLD) to remove instrumental systematics from stellar light curves. Given a stellar image spread out over a set of pixels $\{\vec{p_i}\}$, `EVEREST` regresses on polynomial functions of the fractional pixel fluxes, $\vec{p_i}/\sum_j \vec{p_j}$, identifying the linear combination of these that best fits instrumental signals present in the light curve. Because astrophysical signals (such as transits) are equally present in each of the pixels, whereas instrumental signals are spatially variable, PLD excels at removing instrumental noise while preserving astrophysical information. `EVEREST` uses a Gaussian process (GP) to model correlated astrophysical noise and employs an L2-regularized regression scheme to minimize overfitting.

Since PLD may overfit in the presence of bright contaminant sources in the target aperture, I manually inspected the TRAPPIST-1 *K2* postage stamp and high resolution images of TRAPPIST-1 taken with the APO ARC 3.5m telescope in the SDSS z band to verify that there were no other targets brighter than the background level in our adopted aperture. Given the faint magnitude of TRAPPIST-1 in the *Kepler* band, I used the PLD vectors of 14 nearby bright stars (EPIC IDs 246177238, 246165150, 246211745, 246171759, 246127507, 246228828, 206392586, 246121678, 246229336, 246196866, 246217553, 246239441, and 246144695) generated using the same method as above to improve the signal-to-noise ratio of the instrumental model (Chapter 6). I further masked the data in the vicinity of all transits of planets TRAPPIST-1b–g and the potential transits of TRAPPIST-1h when computing the model to prevent transit overfitting. Finally, I divided the light curve into three roughly equal segments and detrended each separately to improve the predictive power of the model. Following these steps, I obtained a detrended light curve for TRAPPIST-1 with a 6-hr photometric precision (Christiansen et al., 2012) of 281.3 ppm, a factor of 3 improvement on that of the raw light curve (884.4 ppm); see Figure 74. Before the *K2* observation, I estimated the

Kp magnitude of TRAPPIST-1 to be 17.2 ± 0.3 based on a fit to a corrected blackbody spectrum. However, the photometric precision I achieved with `EVEREST` is inconsistent with a target dimmer than $Kp \sim 17$. My detrending therefore suggests that the magnitude of TRAPPIST-1 in the *Kepler* band is $16 < Kp < 17$.

7.2.2 Light curve detrending – Gaussian process (GP) model

Independently, one of my collaborators (M. Sestovic) also detrended the data with a GP-based pipeline. To perform aperture photometry, he located the star using a centroid fit and apply a circular top-hat aperture following the star's centroid coordinates. He then used a GP model to remove the pointing drift systematics using an additive kernel with separate spatial, time and white noise components (Aigrain et al., 2015, 2016):

$$k_{xy}(x_i, y_i, x_j, y_j) = A_{xy} \exp\left[-\frac{(x_i - x_j)^2}{L_x^2} - \frac{(y_i - y_j)^2}{L_y^2}\right] \tag{150}$$

$$k_{xy}(t_i, t_j) = A_t \exp\left[-\frac{(t_i - t_j)^2}{L_t^2}\right] \tag{151}$$

$$K_{ij} = k_{xy}(x_i, y_i, x_j, y_j) + k_t(t_i, t_j) + \sigma^2 \delta_{ij} \tag{152}$$

where x and y are the pixel coordinates of the centroid, t is the time of the observation, and the other variables (A_{xy}, L_x, L_y, A_t, L_t, σ) are hyperparameters of the GP model. He used the `GEORGE` package (Ambikasaran et al., 2016) in `PYTHON` to implement the GP model. To find the maximum likelihood hyperparameters he used a differential evolution algorithm (Storn and Price, 1997), followed by a local optimization. This method was tested on magnitude $16 - 18$ stars observed in Campaign 10 of *K2*, and we used the results of those tests, and of previous GP applications to *K2* data (Aigrain et al., 2016), to inform the priors on the hyperparameters.

For the TRAPPIST-1 data, he used an iterative sigma-clipping method to remove outliers and prevent the time component from overfitting. This method has been previously used in the `k2sc` (Aigrain et al., 2016) pipeline. First, using fiducial hyperparameter values based on analysis of a Campaign 10 target, he removed all measurements with residuals greater than 3σ from the mean GP prediction. With the remaining measurements, he updated the hyperparameters by maximizing the GP likelihood. Using these parameters, he once again clipped all 3σ outliers and

maximized the GP likelihood using only the remaining measurements. The final detrending was calculated for all points, including outliers.

7.3 Laplace Resonances

A dynamical analysis made by my collaborators (B.-O. Demory, E. Agol, and D. Fabrycky) prior to the release of the *K2* data suggested certain values of the period of TRAPPIST-1h based on the presence of three-body resonances among the planets. In order to understand a three-body resonance, we must first define the mean longitude of a planet's orbit: this is an angular variable that progresses at a constant rate with respect to time t, which is measured from the time the planet passes a given reference direction:

$$\lambda = \frac{360^\circ}{P} t, \tag{153}$$

where P is the orbital period and λ is measured in degrees. For transiting planets, the reference direction is taken as the plane perpendicular to the observer's line of sight, as the planet is progressing towards the transiting configuration. If we assume orbits with negligible eccentricities (Mills et al., 2016), for which $\lambda = 90^\circ$ at transit mid-time, we may write

$$\lambda = 360^\circ \left(\frac{1}{4} + \frac{t - T_n}{P} \right) \tag{154}$$

for each planet, where T_n is the time of transit of the n^{th} planet.

Three-body resonances satisfy $pP_1^{-1} - (p+q)P_2^{-1} + qP_3^{-1} \approx 0$ and $p\lambda_1 - (p+q)\lambda_2 + q\lambda_3 = \phi$ for integers p, q and where P_i and λ_i are the period and mean longitude of the i^{th} planet (Peale, 2002; Rivera et al., 2010) and ϕ is the 3-body angle, which librates about a fixed value:

$$\phi = 360^\circ \left[-p\frac{T_1}{P_1} + (p+q)\frac{T_2}{P_2} - q\frac{T_3}{P_3} + p\frac{t}{P_1} - (p+q)\frac{t}{P_2} + q\frac{t}{P_3} \right] \tag{155}$$

for transit times T_1, T_2, and T_3 taken near each other. These resonances occur both in our Solar System — the archetypical case being the Laplace resonance among Jupiter's satellites, satisfying $(p, q) = (1, 2)$ — and in exoplanet systems, two of which were recently observed to have resonant chains among four planets: Kepler-223 (Mills et al., 2016) with $(p, q) = (1, 1)$ and Kepler-80

planets 1,2,3	p	q	$\frac{p}{P_1} - \frac{(p+q)}{P_2} + \frac{q}{P_3}$ [d^{-1}]	$\phi = p\lambda_1 - (p+q)\lambda_2 + q\lambda_3$
b,c,d	2	3	$[-4.6, -0.3] \times 10^{-5}$	[176°, 178°]
c,d,e	1	2	$[-5.2, 4.5] \times 10^{-5}$	[47°, 50°]
d,e,f	2	3	$[-1.9, +1.9] \times 10^{-4}$	[−154°, −142°]
e,f,g	1	2	$[-1.4, +1.1] \times 10^{-4}$	[−79°, −72°]
f,g,h	1	1	$[-6.0, +0.2] \times 10^{-5}$	[176.5°, 177.5°]

Table 2: The three-body resonances of TRAPPIST-1. The transit times are used to track the ϕ angles of each set of three adjacent planets over the dataset, assuming low eccentricities such that transits occur at a phase angle $\lambda = 90°$ (Mills et al., 2016). The ranges of three-body frequency and angle given encompass the changes — most likely librations — seen during the observations.

(MacDonald et al., 2016) with $(p, q) = (2, 3)$.

Among the inner six planets in TRAPPIST-1, there are *four* adjacent sets of three planets that satisfy this relation for $1 \leq p \leq 2$ and $1 \leq q \leq 3$ (Table 2). This suggested that the period of planet TRAPPIST-1h may also satisfy a three-body resonance with TRAPPIST-1f and g. The six potential periods of TRAPPIST-1h that satisfy three-body relations with $1 \leq p, q \leq 3$ are 18.766 d ($p = q = 1$), 14.899 d ($p = 1, q = 2$), 39.026 d ($p = 2, q = 1$), 15.998 d ($p = 2, q = 3$), 13.941 d ($p = 1, q = 3$), and 25.345 d ($p = 3, q = 2$). My collaborators examined $\sim$ 1000 hours of ground-based data taken prior to the Spitzer dataset (Gillon et al., 2017) and found a lack of obvious additional transits at the expected times for all of these periods save 18.766 d. The period of 18.766 d corresponds to prior transit times in windows that were missed by the previous ground-based campaigns; hence, this was the only period that could not be ruled out. Furthermore, as this value is consistent with the period estimate of 20^{+15}_{-6} d based on the duration of the *Spitzer* transit, we had reason to believe it was the correct period for TRAPPIST-1h.

7.4 Search for TRAPPIST-1h

I conducted three separate transit searches on the long cadence light curve, aiming to constrain the period of TRAPPIST-1h, which had only been observed to transit once (Gillon et al., 2017), as well as to detect additional planets in the system.

7.4.1 Fold-and-inspect

To test the hypothesis based on the Laplace resonance argument, I folded the long cadence data on the dynamically predicted orbital period of 18.765 d for TRAPPIST-1h with a time of first transit

constrained by the *Spitzer* observation, revealing a feature consistent with a transit in both the `EVEREST` and GP-based light curves. In order to confirm the planetary nature of this signal, I next analyzed the short cadence data. To this end, I used `kadenza` to generate a short cadence TPF of TRAPPIST-1 and detrended it in windows of 1.5–2 d centered on each of the four features using `EVEREST`. I used the PLD vectors of five bright stars observed in short cadence mode (EPIC IDs 245919787, 246011640, 246329409, 246331757, and 246375295) to aid in the detrending. When generating these light curves, I explicitly masked large flares so that these did not inform the fit. Following this procedure, I obtained binned 6-hr photometric precision of 266.6, 176.1, 243.4, and 219.3 ppm in each of the four windows. The short and long cadence data in these windows is shown in Figure 69. For transits 3 and 4, additional correction of the light curve was necessary, since the transit of TRAPPIST-1h coincides with a transit of TRAPPIST-1b (panel 3a) and a small flare (panel 4b). The transit of TRAPPIST-1b was subtracted out using a transit model (Mandel and Agol, 2002) with the *Spitzer* parameters and a mid-transit time determined from the data, yielding the light curve in panel 3b. The flare was modeled using a 3-parameter flare model (Davenport et al., 2014) for stars observed with *Kepler*, yielding the light curve in panel 4b (see also Figure 70). In both cases, the transit of TRAPPIST-1h is visible in the residuals. In Figure 71 I show the folded short cadence data after accounting for transit timing variations (TTVs), with a transit model based solely on the *Spitzer* parameters.

7.4.2 Formal Transit Search

Independently, my collaborators (B.-O. Demory, L. Delrez, and M. Gillon) used the long-cadence detrended light curve to perform a transit search with a box least-squares fitting algorithm (BLS) (Kovács et al., 2002). They set the BLS to orbital periods ranging from 10 d to 50 d. The ratio of the transit duration over the planet orbital period was set between 0.0007 and 0.06 to include a wide range of orbital periods, eccentricities, and impact parameters to search for additional planets in the system. The highest peak in the periodogram corresponded to a signal with a 15.44-day period. Its origin stemmed from residuals in blended transits with TRAPPIST-1c and d and from one outlier in the data. No signal was seen at the two other epochs where a transit should have appeared, which confirms the signal being spurious. The next highest peak in the periodogram corresponded to a ∼18.77-day period and a transit center at BJD = 2, 457, 756.39 that is consistent

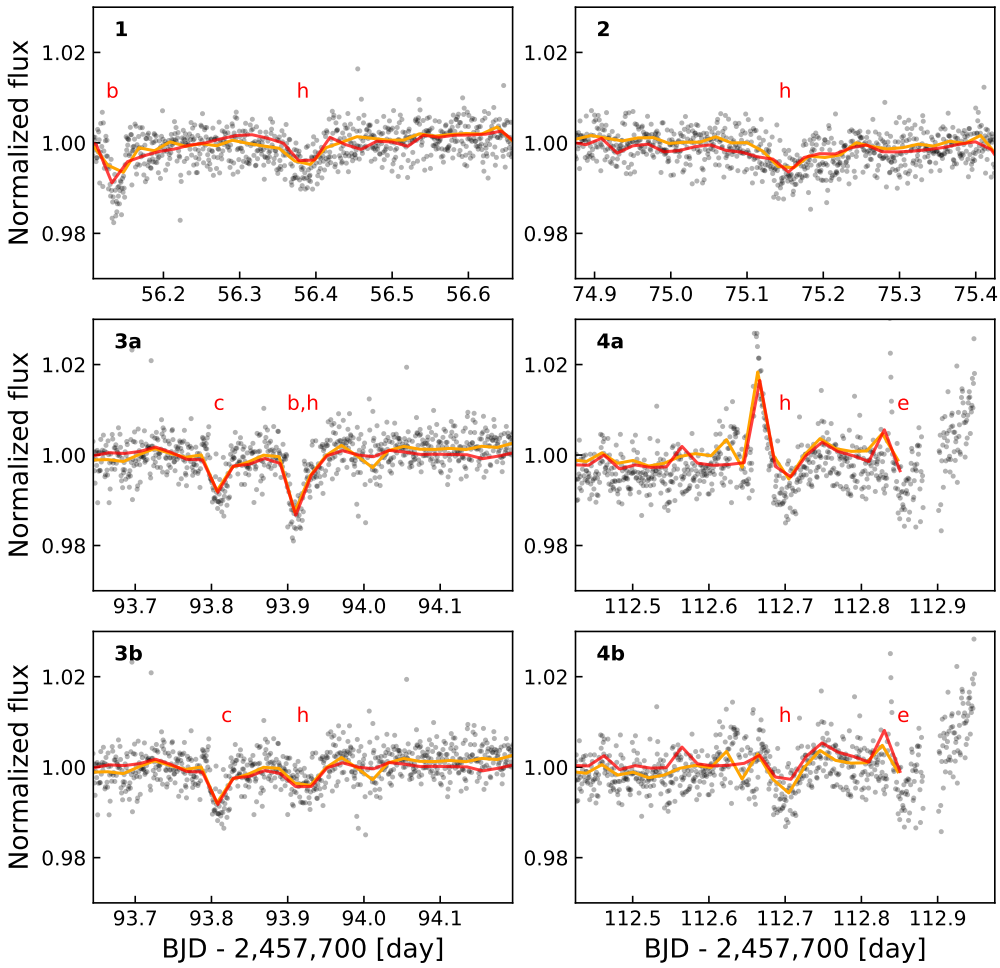

Figure 69: The four transits of TRAPPIST-1h. The detrended short cadence data is shown as black dots. Orange curves are these data binned to 30-minute cadence, and red curves are the detrended long cadence data. The transits of TRAPPIST-1h and other planets are indicated with red letters. Corrections had to be made to remove the simultaneous transit of *b* in transit 3 and a near-simultaneous flare in transit 4. The uncorrected data for these transits is shown in the middle row, and the data with these features removed is shown in the bottom row.

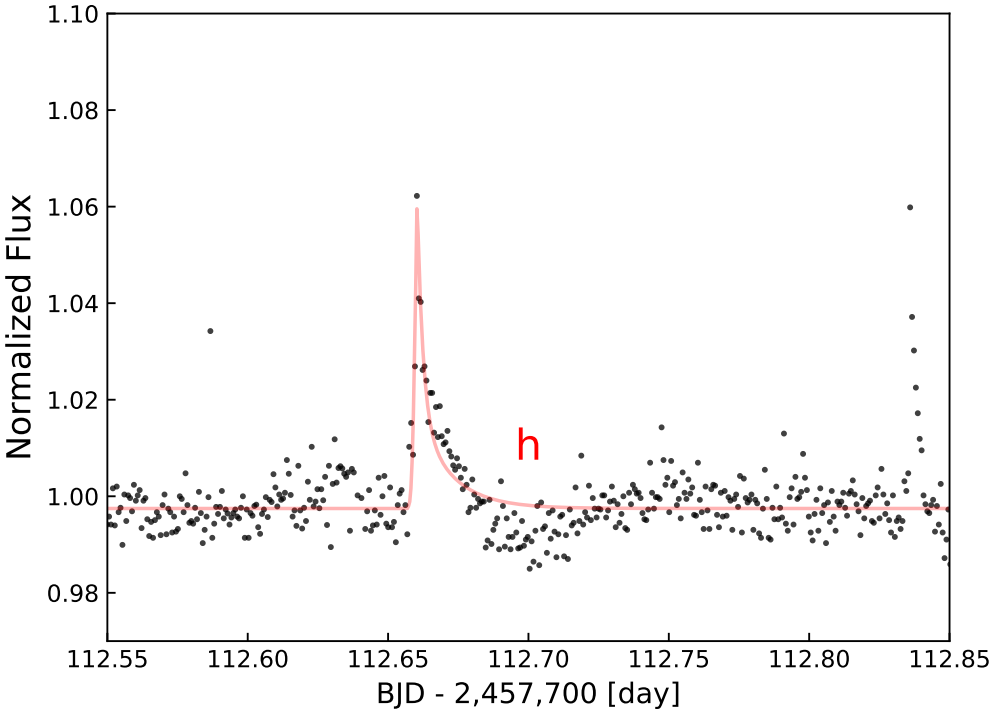

Figure 70: The short cadence data in the vicinity of transit 4 of planet TRAPPIST-1h, where a small flare is visible. The transit of TRAPPIST-1h occurs ∼60 short cadences (1 hr) after the peak of the flare. A least-squares fit to the flare is shown in red; the data during the transit of TRAPPIST-1h is clearly lower than the baseline.

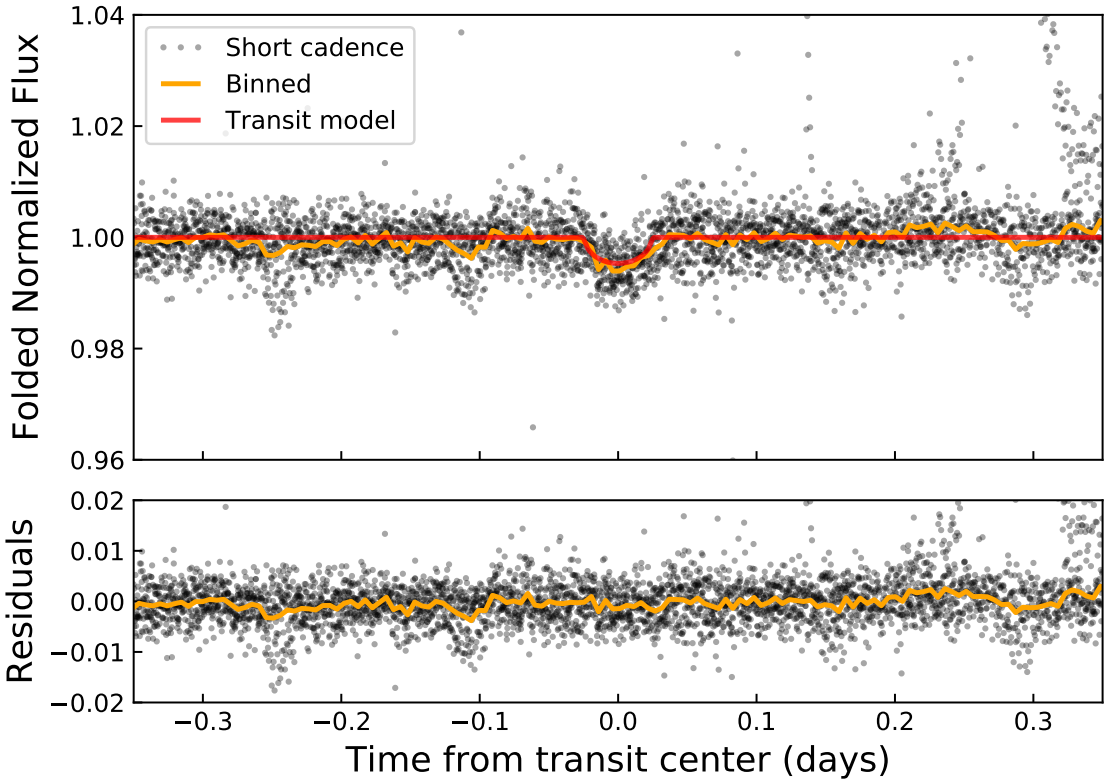

Figure 71: The short cadence data folded on the four transits of TRAPPIST-1h after correcting for TTVs and subtracting a simultaneous transit of TRAPPIST-1b and a near-simultaneous flare. Other transits of TRAPPIST-1b to g have not been removed and are visible in parts of the data. The data downbinned by a factor of 30 is shown as the orange line, and a transit model based *solely* on the *Spitzer* parameters is shown in red. The residuals (data minus this model) are shown at the bottom.

with the single transit seen with *Spitzer*.

They then used a Markov Chain Monte Carlo (MCMC) algorithm previously described in the literature (Gillon et al., 2012) to derive the transit parameters of TRAPPIST-1h from the detrended light curve. Each photometric data point was attached to a conservative error bar that accounted for the uncertainties in the detrending process presented in the previous section. My collaborators imposed normal priors on the orbital period, transit mid-time center and impact parameter for planets TRAPPIST-1b to g based on the parameters in the original discovery paper (Gillon et al., 2017). They further assumed circular orbits for all planets (Gillon et al., 2016, 2017). They also included normal priors for the stellar properties, which are $\mathcal{N}(0.082,\, 0.011)\,\mathrm{M}_\odot$ for the mass, $\mathcal{N}(0.114,\, 0.006)\,\mathrm{R}_\odot$ for the radius, $\mathcal{N}(2{,}555,\, 85)\,\mathrm{K}$ for the effective temperature and $\mathcal{N}(0.04,\, 0.08)$ dex for the metallicity (Gillon et al., 2017). They used these stellar parameters to compute the quadratic limb-darkening coefficients u_1 and u_2 in the *Kepler* bandpass from theoretical tables (Claret and Bloemen, 2011).

In a first MCMC fit, my collaborators used a 7-planet model that included all known planets, with no prior information on the orbital period or t_0 of TRAPPIST-1h. In the second fit they employed a 6-planet model that excluded TRAPPIST-1h. They used the results from both MCMC fits to compute the Bayesian and Akaike Information Criteria (BIC and AIC, respectively) to determine which model was favored. They found BIC values of 2888 and 2897 for the 7 and 6-planet models, respectively. This corresponds to a Bayes Factor $e^{(BIC_1 - BIC_2)/2} = 90$ in favor of the 7-planet model. Similarly, they found AIC values of 2691 and 2725 for the 7 and 6-planet models, respectively.

Finally, they performed a third MCMC fit to refine the transit parameters of TRAPPIST-1h. For this fit, they used as input data the *K2* short-cadence data centered on the 4 transits of TRAPPIST-1h (Figure 74) and the single transit light curve previously obtained with *Spitzer*. This fit included a model for TRAPPIST-1b and a flare that both affect the transit shape of TRAPPIST-1h in the K2 short-cadence data. This fit also allowed for TTVs in the individual transit timings. They found photometric precisions of 365 ppm and ∼1100 ppm per 10 min for *Spitzer* and *K2* respectively. The median and 1-sigma credible intervals of the posterior distribution functions for the transit parameters of TRAPPIST-1h are shown in Table 3, along with the individual transit times.

7.4.3 Joint Instrumental/Transit Model

In order to prevent the overfitting of transit features, I masked all transits of planets b–h when detrending with `EVEREST`. However, this inevitably results in a lower detrending power during transits. A powerful alternative to the detrend-then-search method employed above is to simultaneously fit the instrumental and transit signals, without masking those features (Foreman-Mackey et al., 2015). I therefore conducted a second, separate blind search on the `EVEREST` light curve specifically for TRAPPIST-1h. Given a raw light curve $\vec{y}$, a data covariance matrix $\mathbf{\Sigma}$, and a single-transit model $\vec{m}_{t_0}$ centered at $t = t_0$, the log likelihood of the transit fit is

$$\log \mathcal{L} = -\frac{1}{2}(\vec{y} - \vec{m}_{t_0})^\top \mathbf{\Sigma^{-1}} (\vec{y} - \vec{m}_{t_0}) + C \tag{156}$$

where C is a constant. The data covariance matrix, $\mathbf{\Sigma}$, is the sum of the astrophysical covariance and the L2-regularized PLD covariance and is given by

$$\mathbf{\Sigma} = \mathbf{X \Lambda X}^\top + \mathbf{K} \tag{157}$$

where $\mathbf{K}$ is the astrophysical covariance given by the `EVEREST` GP model, $\mathbf{X}$ is the matrix of PLD regressors (the design matrix), and $\mathbf{\Lambda}$ is the prior covariance of the PLD weights (the regularization matrix) which we obtain by cross-validation (Chapter 6). Since the transit shape and duration of TRAPPIST-1h are known (Gillon et al., 2017), the only free parameter in the search was t_0, the time of transit. I therefore evaluated $\vec{m}_{t_0}$ multiple times, centering the transit model at each long cadence and computing the likelihood of the transit model fit as a function of cadence number. I then subtracted these values from the log likelihood of the data with no transit model ($\vec{m}_{t_0} = 0$) and multiplied by 2 to get the delta-chi squared ($\Delta\chi^2$) metric, which measures the decrease in the χ^2 value of the fitted light curve for a transit of TRAPPIST-1h centered at each cadence. Finally, I also computed $\Delta\chi^2$ conditioned on the known "true" transit depth of TRAPPIST-1h, $d_0 = 0.00352 \pm 0.000326$:

$$\Delta\chi^2_{\text{cond}} = \Delta\chi^2 - \left(\frac{d - d_0}{\sigma_d}\right)^2 \tag{158}$$

where

$$d = \sigma_d^2 \vec{m}_{t_0}^\top \mathbf{\Sigma^{-1}} \vec{y} \tag{159}$$

is the maximum likelihood depth of the transit model and

$$\sigma_d^2 = (\vec{m}_{t_0}^\top \mathbf{\Sigma^{-1}} \vec{m}_{t_0})^{-1} \tag{160}$$

is the variance of the depth estimate. Positive peaks in $\Delta\chi^2$ indicate features that are well described by the transit model, while positive peaks in $\Delta\chi^2_{\rm cond}$ reveal features that are well described by the transit model with depth $d = d_0$. In Figure 72 I show the two $\Delta\chi^2$ metrics across the full TRAPPIST-1 light curve after subtracting a transit model for planets $b-g$ based on their *Spitzer* parameters. The strongest features in the $\Delta\chi^2$ plot (top) are flares, as these can be fitted out with an inverted transit model. When conditioning on the true depth of TRAPPIST-1h (bottom), the significance of most of the flare features decreases, revealing the four peaks of TRAPPIST-1h (red arrows).

In order to assess the robustness of my detection, I computed the total $\Delta\chi^2_{\rm cond}$ as a function of orbital period. Starting from the time of transit of TRAPPIST-1h in the *Spitzer* dataset, I computed the times of transit in the *K2* light curve for 500,000 values of the orbital period evenly spaced between 1 and 50 d and summed the values of $\Delta\chi^2_{\rm cond}$ at each transit time to produce the total $\Delta\chi^2_{\rm P}$. I summed these in two different ways. First, I linearly interpolated the grid of $\Delta\chi^2_{\rm cond}$ to each transit time to obtain $\Delta\chi^2_{\rm P}$ for perfectly periodic transits. Next, to allow for TTVs of up to one hour, I took the largest value of $\Delta\chi^2_{\rm cond}$ in the vicinity of each transit time and summed them for each period. I adopted a tolerance of 2 cadences, corresponding to maximum TTVs of 1 hr. My results are shown in Figure 73, where I plot the periodic $\Delta\chi^2_{\rm P}$ (top) and the $\Delta\chi^2_{\rm P}$ allowing for TTVs (bottom). In both cases, the period of TRAPPIST-1h (18.766 d) and its 1/2 and 1/3 period aliases emerge as the three strongest peaks. The peak at 18.766 d is the strongest signal in the period range constrained by the *Spitzer* transit (Gillon et al., 2017) and confirms my detection of TRAPPIST-1h.

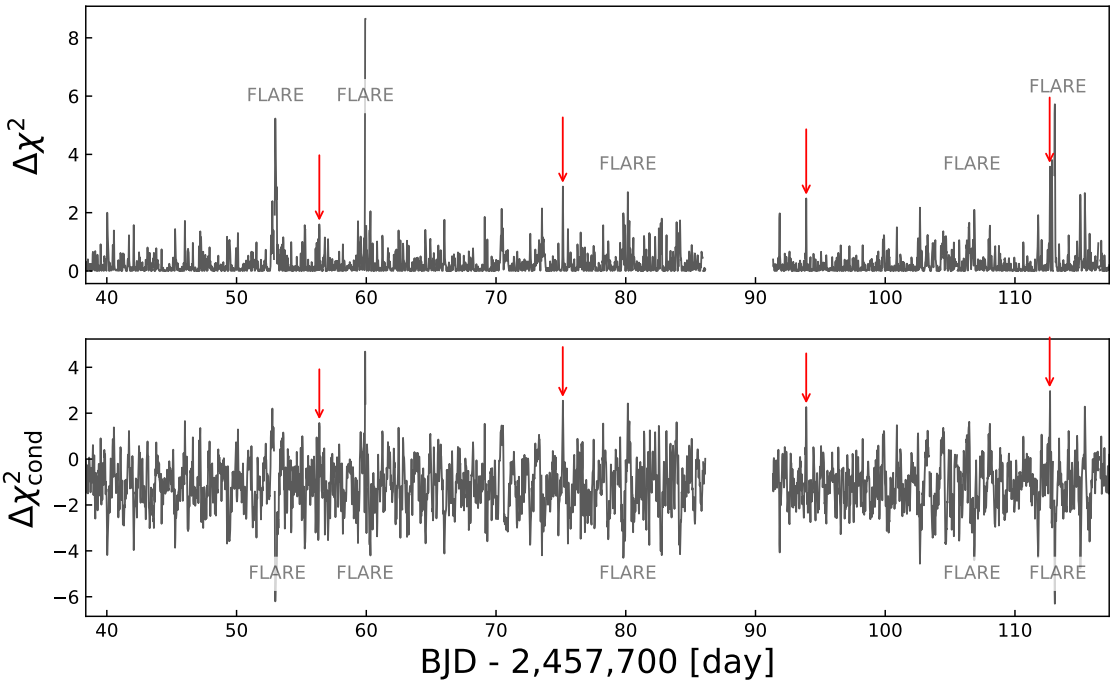

Figure 72: Delta-chi squared (top) and delta chi-squared conditioned on the true depth (bottom) for each long cadence in the TRAPPIST-1 `EVEREST` light curve. The four transits of TRAPPIST-1h are indicated with red arrows. In the top plot, spikes appear at the location of several flares, as these can be fit with inverted transits. Conditioning on the true depth (bottom) removes many of these features by penalizing transit depths that are inconsistent with the observed *Spitzer* value.

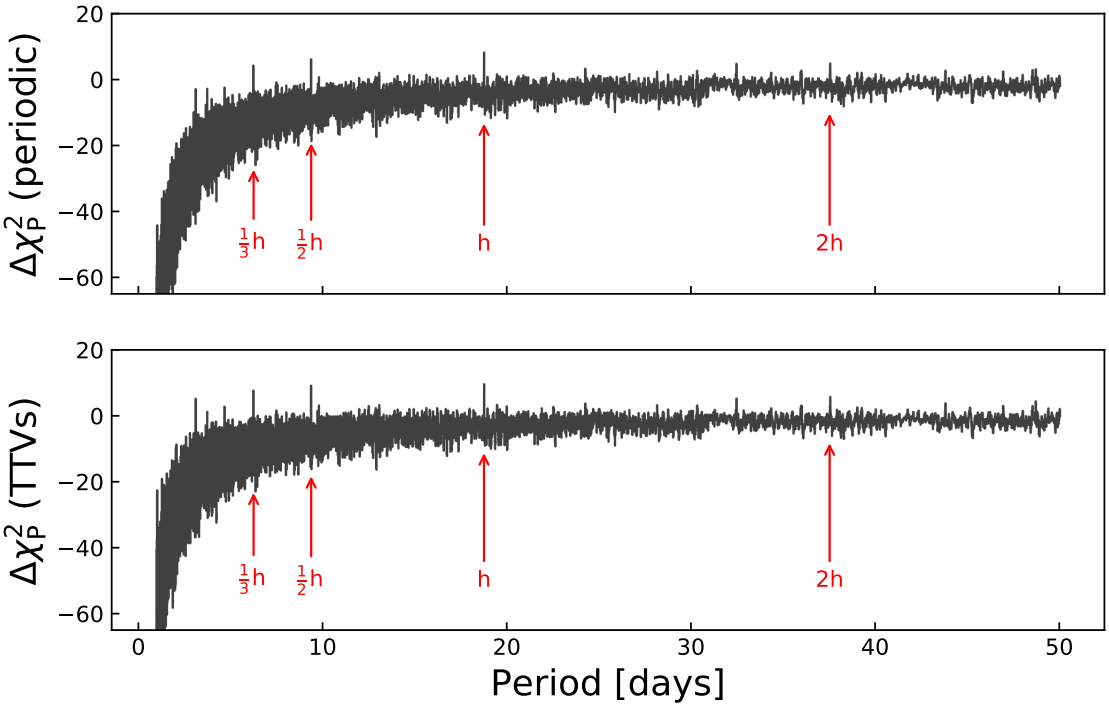

Figure 73: Delta-chi squared values as a function of orbital period for TRAPPIST-1h, taking the *Spitzer* transit time and assuming perfectly periodic orbits (top) and allowing for TTVs of up to 1 hr (bottom). The period of TRAPPIST-1h (18.766 d) and its aliases emerge as the strongest peaks in both plots.

7.4.4 Transit Search Summary

The detection of TRAPPIST-1h is thus supported by 1) the three transit search analyses that recovered both the orbital phase from the Spitzer ephemeris (Gillon et al., 2017) and the period of 18.766 d, 2) the Bayes Factor in favour of the 7-planet model and 3) the orbital period that is the exact value predicted by Laplace relations. Figure 74 shows the full light curve, the newfound transits of TRAPPIST-1h, as well as an update to the geometry of the orbits given the new orbital period. In Table 3 I report the properties of the planet derived in this study.

7.5 Discussion

7.5.1 Three-body resonances

To characterize the three-body resonances in TRAPPIST-1, my collaborators and I used the transit timing data to identify ϕ for each set of three planets. A full transit-timing cycle has not elapsed within the data, an estimate of the libration centers for each ϕ is not possible. However, I report in Table 2 the values of ϕ represented in the dataset. In the case of Jupiter's satellites, $\phi = 180°$, but due to the complexity of the multi-planetary system, I make no predictions for TRAPPIST-1 at this time. Migration and damping simulations applied to Kepler-80 (MacDonald et al., 2016) naturally predicted the values of the libration centers in that system; the measured values for TRAPPIST-1b–h call for future theoretical work interpreting them.

The resonant structure of the system suggests that orbital migration may have played a role in its formation. Embedded in gaseous planet-forming disks, planets growing above $\sim 1\,M_{Mars}$ create density perturbations that torque the planets' orbits and trigger radial migration (Baruteau et al., 2014). One model for the origin of low-mass planets found very close to their stars proposes that Mars- to Earth-sized planetary embryos form far from their stars and migrate inward (Raymond et al., 2008). The inner edge of the disk provides a migration barrier (Masset et al., 2006) such that planets pile up into chains of mean motion resonances (Terquem and Papaloizou, 2007; Ogihara and Ida, 2009; Cossou et al., 2014). This model matches the observed period ratio distribution of adjacent super-Earths (Fabrycky et al., 2014) if the vast majority ($\sim 90\%$) of resonant chains become unstable and undergo a phase of giant impacts (Izidoro et al., 2017). Some resonant chains do survive, and a handful of multiple-resonant super-Earth systems have indeed been characterized

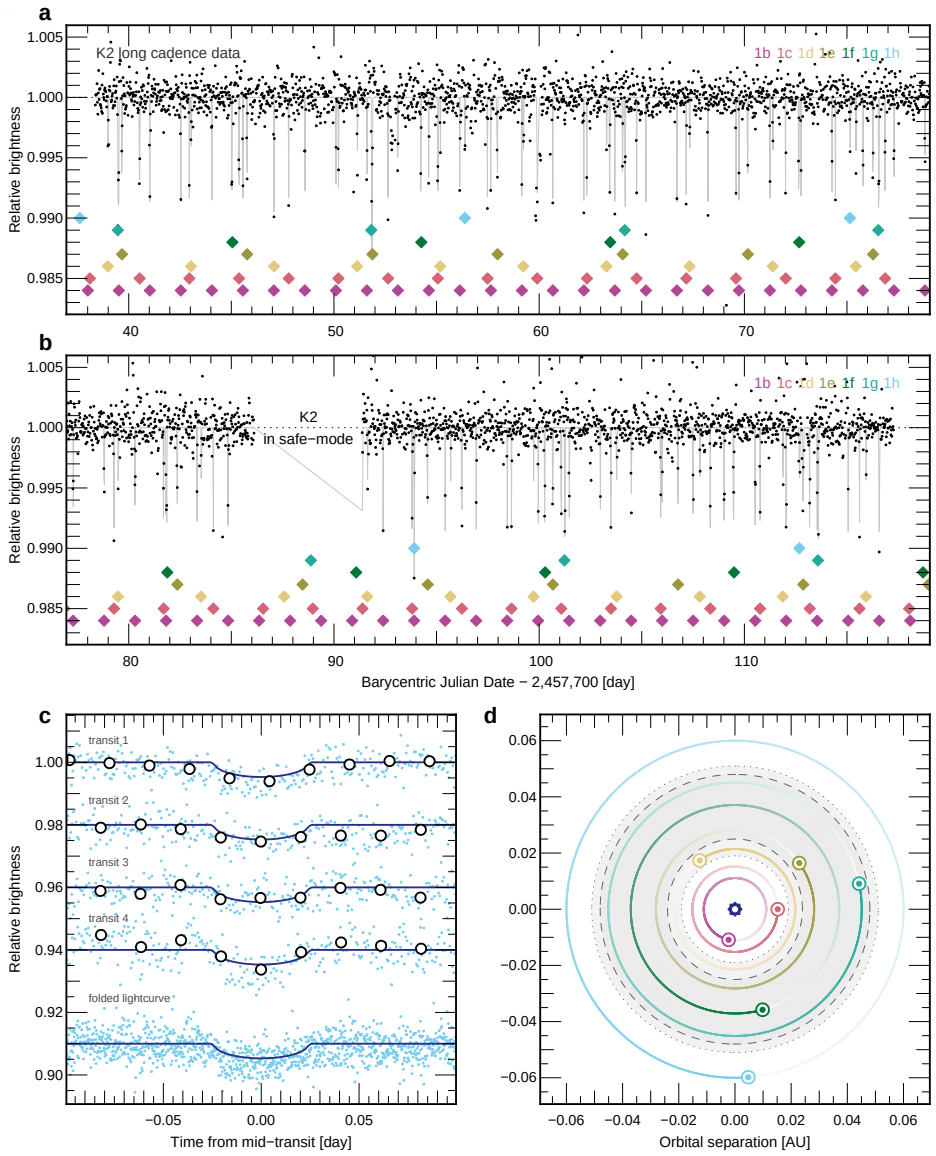

Figure 74: The long cadence *K2* light curve of TRAPPIST-1 detrended with `EVEREST`. **a, b**: The full detrended *K2* light curve with stellar variability removed via LOESS regression (order = 1; width = 0.15 d). Data points are in black, and our highest likelihood transit model for all seven planets is plotted in thin grey. Coloured diamonds indicate which transit belongs to which planet. Four transits of TRAPPIST-1h are observed (light blue diamonds). **c** The top four curves show the detrended and whitened short-cadence in light blue, with a transit model based on the *Spitzer* parameters in dark blue. Binned data is over-plotted in white for clarity. The folded light curve is displayed at the bottom. **d** View from above (observer to the right) of the TRAPPIST-1 system, at the date when the first transit was obtained for this system. The grey region is the surface liquid-water habitable zone.

(Goździewski et al., 2016; Mills et al., 2016). The TRAPPIST-1 system may thus represent a pristine surviving chain of mean motion resonances. Given that TRAPPIST-1's planet-forming disk was likely low in mass (Pascucci et al., 2016) and the planets themselves are low-mass, their migration was likely relatively slow. This may explain why TRAPPIST-1's resonant chain is modestly less compact than chains in systems with more massive planets (Goździewski et al., 2016; Mills et al., 2016; Izidoro et al., 2017), which may have protected it from instability (Matsumoto et al., 2012).

7.5.2 Tides

Given how close the planets orbit TRAPPIST-1, tidal interactions are likely to be important in the planets' orbital evolution (Ferraz-Mello et al., 2008). My collaborator (E. Bolmont) performed N-body simulations of the system including an equilibrium tidal dissipation formalism (Eggleton et al., 1998; Leconte et al., 2010) using the Mercury-T code (Bolmont et al., 2015). She used orbital parameters from the discovery paper (Gillon et al., 2017) and the period for TRAPPIST-1h determined above. She considered the planets' spins to be tidally synchronized with small obliquities. This is justified because even if the age of the system is 400 Myr (a lower estimate for the age of TRAPPIST-1), planetary tides would have had time to synchronize the spins (Leconte et al., 2015) (even when atmospheric tides are accounted for). She tested different initial eccentricities and different values for the planets dissipation factors (from 0.01 to 10 times the Earth's value (Neron de Surgy and Laskar, 1997)).

The simulations show that the planets' orbital eccentricities are likely to be low. In just a few Myr all eccentricities decrease to below 0.01 for dissipation factors ≥ 0.1 times the Earth's value. Due to planet-planet interactions, the eccentricities do not decrease to zero but instead reach an equilibrium value determined by the competition between tidal damping and planet-planet eccentricity excitation (Bolmont et al., 2013). All planets stay in resonance during the evolution towards tidal equilibrium. The small equilibrium eccentricities are sufficient to generate significant tidal heating. Assuming the TRAPPIST-1 planets have a tidal dissipation equal to the Earth's, TRAPPIST-1b might have tidal heat flux similar to Io's (Spencer et al., 2000) ($\sim 3\,\mathrm{W\,m^{-2}}$), with peaks at more than $10\,\mathrm{W\,m^{-2}}$ (corresponding to $\sim 10^4$ TW) when the eccentricity oscillation is at its maximum. Planets c through e have tidal heat fluxes higher than Earth's internal (primarily

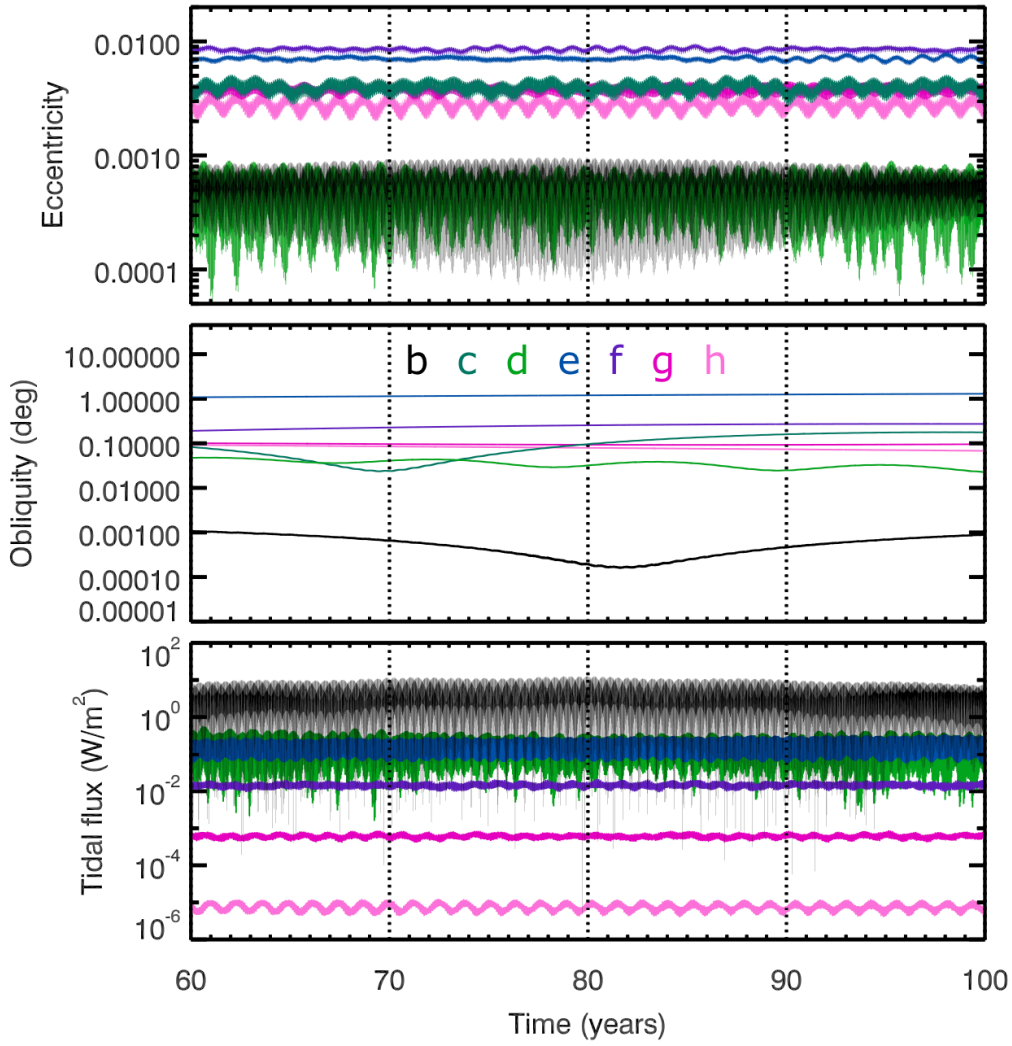

Figure 75: Possible short term evolution of the eccentricity, obliquity and tidal heat flux of the TRAPPIST-1 planets. The different planets are represented by different colors from black (planet b) to light pink (planet h).

radioactive) heat flux ($\sim$0.08 W m^{-2}; Davies and Davies, 2010) but lower than Io's. TRAPPIST-1f, g and h have a tidal heat flux inferior to Earth's. Figure 75 shows a possible snapshot of the system's evolution over the course of 40 yr. This very high flux could plausibly generate intense volcanism on the surfaces of the inner planets, with potential consequences for their internal structures.

7.5.3 Habitability of TRAPPIST-1h

My collaborators (F. Selsis, J. Leconte, and M. Turbet) calculated the minimum stellar flux required for liquid water with the LMD 1D/3D Global Climate Model (Turbet et al., 2016) using a synthetic spectrum of TRAPPIST-1 based on its reported $T_{\rm eff}$, $\log g$, metallicity, and bolometric luminosity (Gillon et al., 2017), obtaining a value of 300 W/m^2, which is 100 W/m^2 higher than the planet's present-day instellation. These results are in agreement with habitable zone boundaries computed for a 3000 K star (Kopparapu et al., 2013). Assuming zero albedo, my collaborators found that the equilibrium temperature of TRAPPIST-1h is 173$\pm$4 K.

Whether or not TRAPPIST-1h presently hosts an atmosphere is unclear. Given its radius measurement and a range of possible compositions (from pure water ice to pure iron), the mass of TRAPPIST-1h is likely in the range 0.067–0.863 $M_\oplus$; if TRAPPIST-1h has an Earth-like composition, this value is 0.33 $M_\oplus$ (Sotin et al., 2007). Assuming TRAPPIST-1h migrated to its current location quickly, the planet's low surface gravity could have led to vigorous hydrodynamic escape of a primordial atmosphere in the first few 100 Myr after its formation, since at that time TRAPPIST-1 was significantly brighter and TRAPPIST-1h would have been interior to the habitable zone (Chapters 2 and 3). The presence of a subsequently outgassed atmosphere, however, cannot be ruled out.

In theory, the surface of TRAPPIST-1h could harbor liquid water under such an outgassed atmosphere if it is H_2-rich. Atmospheres made of H_2 (Stevenson, 1999; Pierrehumbert and Gaidos, 2011), N_2-H_2 (Wordsworth and Pierrehumbert, 2013a) or CO_2-H_2 (Ramirez and Kaltenegger, 2017; Wordsworth et al., 2016) have been shown to provide a sufficient greenhouse effect and internal heat blanketing for even lower instellation levels. Unless prevented by high altitude clouds, transit spectroscopy with the *Hubble* Space Telescope may be able to reveal or rule out such H_2-rich atmospheres.

Alternatively, a liquid water ocean is possible under a layer of ice. The minimum thickness h

of this layer depends on the internal heat flux Φ_{int}:

$$h \sim 250 \left(\frac{\phi_{int}}{1\,\mathrm{W\,m^{-2}}} \right)^{-1} \mathrm{m}$$

Assuming the Earth's current geothermal flux, a layer of 2.7 km (the mean depth of Earth's oceans) would be necessary.

7.5.4 Stellar variability

While the long spin-down times of ultra-cool dwarfs prevent derivation of a robust gyrochronology relation (Reid and Hawley, 2013), the rotational period of TRAPPIST-1 can be used to derive a provisional age estimate for the system. Fourier analysis of the detrended *K2* data (Figure 76), which is visibly modulated by star spots, led to the determination of a rotational period of $\sim 3.3 \pm 0.14$ d for the host star. This rotation corresponds to an angular momentum of about 1% that of the Sun. It is roughly in the middle of the period distribution of nearby late M dwarfs (Newton et al., 2016), suggesting an age in the range $3 - 8$ Gyr based on a star formation history that declines slightly with time (Aumer and Binney, 2009). Such an age is consistent with the star's solar metallicity (Gillon et al., 2016) and borderline thin disk/thick disk kinematics (Burgasser et al.). The amplitude of the modulation due to star spots and infrequent weak optical flares ($0.26\,\mathrm{d^{-1}}$ for peak fluxes above 1% of the continuum, 40 times less frequent than active M6-M9 dwarfs (Hilton et al., 2011)) are consistent with a low-activity M8 star, also arguing in favor of a relatively old system. Near the end of the *K2* campaign, a very energetic flare erupted, and was observed by *Kepler*. Full modeling of flares will be presented in a forthcoming paper.

7.6 Conclusions

In this chapter, I showed how the *K2* observations of TRAPPIST-1 star enabled the detection of four additional transits of TRAPPIST-1h, allowing the determination of its orbital period (18.77 days). I used `EVEREST` to de-trend the light curve and enable the detection of the transits, which were marginally above the noise level. In order to assess the robustness of the detection, I developed a joint instrumental/transit model and conditioned the search on the known depth and time of transit from the single *Spitzer* observation, finding strong evidence for the planet. Moreover, the detected period matches the predicted period to *one part in ten thousand*, confirming the existence

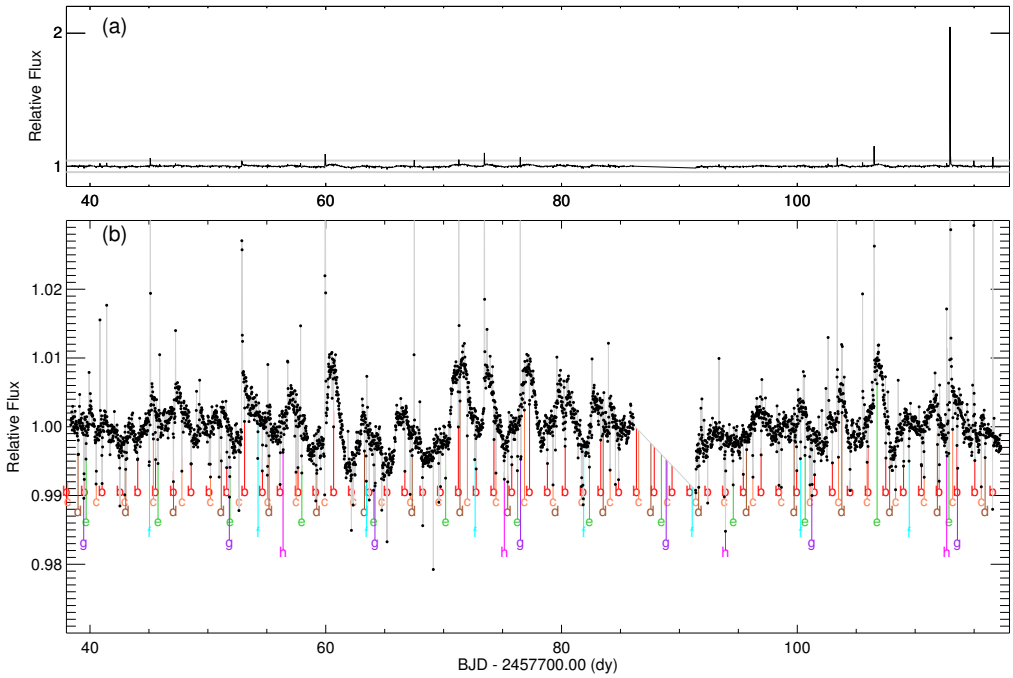

Figure 76: The entire systematics-corrected *K2* dataset with low frequency trends removed. The stellar rotation is apparent in the peaks and troughs of the variability, as are the flares which in some cases appear as single spikes. The planet transits are marked. **a**: Full dynamic range of the curve, including an extreme event at approximately day 113. **b**: Zoomed view of the region outlined in gray in **a**.

of a chain of Laplace resonances linking every planet in the system and its neighbors. This is to date the longest known chain of resonances in a planetary system.

At 18.77 days, TRAPPIST-1h is not in the habitable zone and would require a hydrogen greenhouse to maintain habitable surface conditions. Given the probable history of high XUV activity of its ultracool dwarf host star, it is unlikely that TRAPPIST-1h has a significant hydrogen envelope today (Chapter 3). Alternatively, a thick layer of ice could support a sub-surface liquid water ocean, particularly if tidal heating is significant; in this case, TRAPPIST-1h may be analogous to Europa or Enceladus.

I searched for but did not detect additional planets in the system. Given the extreme coplanarity of the TRAPPIST-1 planetary disk (see Chapter 10), additional planets out to very large periods would be expected to transit. However, signals due to planets smaller than TRAPPIST-1h could have easily been missed due to their small signal-to-noise ratios.

Finally, the *K2* observations of TRAPPIST-1 enabled the determination of the rotation period of the host star (3.3 days) and the detection of numerous flares consistent with a middle-aged, moderately active star. Future work on the short cadence dataset will help to further constrain the flaring activity of TRAPPIST-1.

The code I used to generate and analyze the `EVEREST` light curves for TRAPPIST-1 is openly available at `https://github.com/rodluger/trappist1`. A static version of the repository has been archived at `https://doi.org/10.5281/zenodo.437548`. The $K2$ raw cadence data used in this study is available for download at `https://archive.stsci.edu/missions/k2/c12_raw_cadence_data/`. The pseudo-target pixel files for TRAPPIST-1 and its neighboring stars generated with `kadenza` are archived at `https://doi.org/10.5281/zenodo.437876`. The detrended TRAPPIST-1 long cadence light curve and segments of the short cadence light curve in the vicinity of the transits of TRAPPIST-1h are available at `https://doi.org/10.5281/zenodo.437548`.

Parameter	Value
Transit depth $(R_p/R_*)^2$ [%]	0.346±0.018
Transit duration [d]	0.0525±0.0008
Impact parameter b [R_*]	$0.45^{+0.06}_{-0.08}$
Mid-transit time, T_0 ($\mathrm{BJD_{TDB}}$)	2,457,662.55284±0.00037
Period, P [days]	$18.767^{+0.004}_{-0.003}$
Radius ratio, R_p/R_*	0.0588±0.0016
Radius, R_p [$R_\oplus$]	$0.752^{+0.032}_{-0.031}$
Inclination, $i(^\circ)$	$89.76^{+0.05}_{-0.04}$
Scale parameter, a/R_*	109±4
Equilibrium temperature [K]	173±4
Irradiation, S_p [S_{Earth}]	0.165±0.025
Limb-darkening parameters (*Kepler* bandpass)	
u_1	1.00±0.02
u_2	-0.04±0.04
Individual transit timings from K2 ($\mathrm{BJD_{TDB}}$)	
Transit 1	$2{,}457{,}756.3874^{+0.0013}_{-0.0013}$
Transit 2	$2{,}457{,}775.1539^{+0.0016}_{-0.0016}$
Transit 3	$2{,}457{,}793.9230^{+0.0024}_{-0.0025}$
Transit 4	$2{,}457{,}812.6987^{+0.0045}_{-0.0042}$

Table 3: Properties of TRAPPIST-1h, limb-darkening parameters and transit timings derived using a joint *Spitzer* and *K2* dataset. Parameter values are the medians of the posterior distributions from the MCMC and the associated error bars are the 1-sigma credible intervals.

Part III:

Novel Techniques to Detect and Characterize Planets in the Habitable Zones of M Dwarfs

8 Overview

The past two decades have seen a revolution in the detection of extrasolar planets. Today, we live at the brink of a revolution in our ability to characterize them. While our ability to detect these planets continues to improve with new generations of planet-hunting telescopes such as the Transiting Exoplanet Survey Satellite (TESS; Ricker et al., 2015) and the PLAnetary Transits and Oscillations telescope (PLATO; Ragazzoni et al., 2016), the soon-to-be-launched James Webb Space Telescope (JWST) and concept flagship missions such as LUVOIR and HABEX will revolutionize the field of exoplanet characterization. These telescopes will offer unprecedented views of the atmospheres of planets via transmission and secondary eclipse spectroscopy, potentially enabling the detection of biosignatures in habitable zone planets (e.g., Johnson, 2014; Barstow and Irwin, 2016; Domagal-Goldman et al., 2017). From the ground, a new generation of extremely large telescopes (ELTs) will enable the direct imaging of small planets in temperate orbits around nearby stars (Guyon and Martinache, 2013).

However, such exponential improvements in the sensitivity of our instrumentation enable more than just the discovery of smaller planets or the detection of smaller signals in their atmospheres: in some cases these improvements can open up doors to new detection and characterization methods altogether. In this part of my dissertation, I discuss two novel methods for the detection and characterization of terrestrial exoplanets: exoauroral searches (Chapter 9) and searches for planet-planet occultations (Chapter 10).

In Chapter 9, I show how the detection of auroral emission from Proxima Centauri b could confirm its existence and yield information about its atmospheric composition and its three-dimensional orbit. While auroral emission from Earth is relatively weak, if Proxima Centauri b is a terrestrial planet with a magnetic field and an atmosphere containing oxygen-bearing compounds, auroral emission could be 100 to 10,000 times stronger than on Earth due to the planet's extremely active host star. While challenging, detection of exo-aurorae could be performed with future space or ground-based telescopes, providing a novel method to characterize non-transiting exoplanets.

In Chapter 10, I discuss the detectability of planet-planet occultations (PPOs), events in which an exoplanet occults another planet in the same system as seen by an observer on Earth. Because of

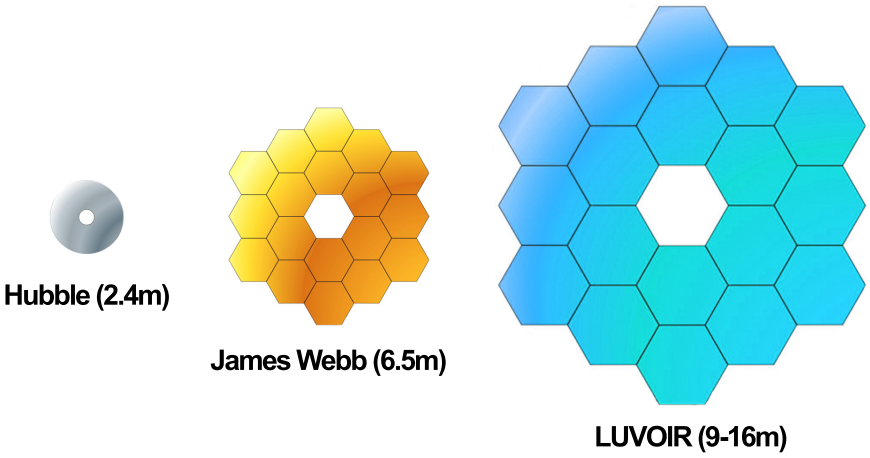

Figure 77: Progression of NASA flagship space telescope mirror sizes, 1990–2020+. The diameter of *LUVOIR* has not yet been decided; the mission will be up for selection in the 2020 NASA Decadal Survey. *Figure adapted from NASA images.*

the extreme faintness of the signal, PPOs have never been detected (or searched for). However, the discovery of extremely coplanar systems of closely packed transiting planets, such as TRAPPIST-1, and the advent of large space-based telescopes such as JWST and the proposed Origins Space Telescope (OST), could make the detection of PPOs possible in the near future. I discuss how such detections could offer unprecedented views of the three-dimensional architecture of multiplanet systems and enable crude surface maps of close-in exoplanets.

9 Exoaurora

Portions of this chapter were originally published in collaboration with J. Lustig-Yaeger, D. P. Fleming, M. Tilley, E. Agol, V. S. Meadows, R. Deitrick, and R. Barnes in the March 2017 edition of the Astronomical Journal (Luger et al., 2017, ApJ, Vol. 837, 63; © 2017 American Astronomical Society), and are reproduced below with permission of the American Astronomical Society.

9.1 Introduction

The discovery of Proxima Centauri b (henceforth 'Proxima Cen b'), only 1.3pc distant from the Sun (Anglada-Escudé et al., 2016), ushers in a new era of characterization of nearby potentially habitable exoplanets. Although Proxima Cen b is not known to transit—making transmission spectroscopy impossible—it is an ideal candidate for high-contrast direct spectroscopy using an extremely large coronagraph-equipped telescope. However, even with the enhancement in angular resolution provided by the proximity of its host star, Proxima Cen b's close-in orbit ($a = 0.0485$ AU; Anglada-Escudé et al., 2016) precludes imaging with current coronagraphs, such as the Gemini Planet Imager (GPI; Macintosh et al., 2014) and the Very Large Telescope's Spectro-Polarimetric High-contrast Exoplanet REsearch facility (VLT-SPHERE; Beuzit et al., 2008), which operate primarily in the near-infrared. This is in part due to the poorer Strehl ratios currently achievable at visible wavelengths with ground-based adaptive optics (AO) systems[18]. Consequently, in advance of larger diameter ground- and space-based telescopes, and improvements in visible AO systems, we must initially consider observations that do not rely on transits or current coronagraphy to search for and characterize the atmosphere of Proxima Cen b.

Phase curves may offer one of the first means to study the atmosphere of Proxima Cen b (Turbet et al., 2016; Kreidberg and Loeb, 2016; Meadows et al., 2016) by potentially showing the reduction in day-night thermal emission contrast associated with an atmosphere. Phase curves have proven to be a successful means to characterize the atmospheres of planets larger and hotter than Proxima Cen b (Cowan et al., 2007; Knutson et al., 2007; Knutson et al., 2008; Crossfield et al., 2010; Brogi et al., 2012; Zellem et al., 2014; Stevenson et al., 2014), including ones that

[18]See, e.g., `https://www.eso.org/sci/facilities/paranal/instruments/sphere/overview.html`

do not transit (Selsis et al., 2011; Faigler and Mazeh, 2011; Maurin et al., 2012; Brogi et al., 2014). However, the expected planet-star contrast ratio in the visible and NIR due to reflected stellar radiation is likely to be below the anticipated systematic noise floor for JWST/NIRSpec (Meadows et al., 2016), and although the planet-star contrast ratio becomes quite favorable beyond 10μm, where the planetary thermal emission peaks, mid-IR phase curves will require JWST/MIRI.

To complement the anticipated JWST thermal phase curve measurements, in this chapter I explore the possibility of directly detecting optical auroral emission from the atmosphere of Proxima Cen b using high-resolution optical spectroscopy. Numerous studies have investigated exoplanet aurorae in the radio due to cyclotron and synchrotron emission to constrain the planetary magnetic field (e.g. Bastian et al., 2000; Grießmeier et al., 2007; Zarka, 2007; Hess and Zarka, 2011; Driscoll and Olson, 2011; Grießmeier, 2015). Others have explored the detectability of optical and UV auroral emission from hot Jupiters, including France et al. (2010), who searched for far-UV auroral and dayglow H_2 emission from the hot Jupiter HD 209458b and placed upper limits on its magnetic field strength, and Menager et al. (2013), who studied the detectability of Lyman α auroral emission from HD 209458b and HD 189733b. Some studies have also investigated auroral emission from terrestrial planets, including Smith et al. (2004), who modeled the role of aurorae in redistributing high energy incident stellar flux to the surface of rocky exoplanets, and Bernard et al. (2014), who investigated how the detection of the green oxygen airglow line could be used to infer the presence of planetary hydrogen coronae on CO_2-dominated planets. Finally, Sparks and Ford (2002) suggested that exoplanet airglow and/or aurorae could be detected with a combination of high contrast imaging and high dispersion spectroscopy. However, a detailed calculation of the expected auroral signal strength on a nearby terrestrial exoplanet and the feasibility of its detection has not yet been fully performed.

Detecting optical auroral emission from the possible atmosphere of Proxima Cen b is likely much more favorable for this system than for an Earth-Sun analog. This is due to both planetary and stellar characteristics that favor auroral production and improve detectability (see Table 4). In particular, Proxima Cen b's intrinsic planetary properties may favor production of aurorae from oxygen atoms. If Proxima Cen b is Earth-like in composition, recent dynamical/planetary interior modeling results by Barnes et al. (2016) and Zuluaga and Bustamante (2016) suggest that the planet may have a magnetic field, potentially increasing the likelihood of atmospheric retention

Table 4: Proxima Centauri b properties

Property	Value	1σ Interval
Distance from Earth (pc)	1.295	
Host spectral type	M5.5V	
Host mass, $M_\star$ ($M_\odot$)	0.120	[0.105 – 0.135]
Period, P (days)	11.186	[11.184 – 11.187]
Semi-major axis, a (AU)	0.0485	[0.0434 – 0.0526]
Minimum mass, $m_p \sin i$ ($M_\oplus$)	1.27	[1.10 – 1.46]
Radius, R_p ($R_\oplus$)	Unknown	[0.94 – 1.40]
Eccentricity, e	< 0.35	
Mean longitude, λ (°)	110	[102 – 118]
Inclination, i (°)	Unknown	[0 – 90]

Note. — All values are from Anglada-Escudé et al. (2016), except for the planet minimum mass, which is a plausible range from Brugger et al. (2016), assuming $m_p = 1.27 M_\oplus$.

and of auroral emission. Atmospheres rich in oxygen-bearing molecules, including O_2 and CO_2, have been predicted for Proxima Cen b (Meadows et al., 2016) as a result of the evolutionary processes for terrestrial planets orbiting M dwarfs (Luger and Barnes, 2015; Barnes et al., 2016). On Earth, the oxygen (OI) auroral line at 5577Å provides the distinctive green glow observed in both the Aurora Borealis and the Aurora Australis, and is the brightest (i.e., highest photon emission rate) auroral feature (Chamberlain, 1961; Dempsey et al., 2005). For emissions from the upper atmosphere, only the 1.27μm O_2 airglow and combined near-infrared OH night glow features are brighter (Hunten et al., 1967). The oxygen green line is seen in both the Earth's O_2-rich atmosphere (Chamberlain, 1961) and Venus' CO_2-dominated atmosphere (Slanger et al., 2001), where it has been observed to increase in brightness after CME events (Gray et al., 2014).

The stellar properties and the planet-star separation are also likely to enhance the auroral power on Proxima Cen b relative to an Earth-Sun analog. Proxima Centauri is an active flare star with a magnetic field ∼600× stronger than that of the Sun (Reiners and Basri, 2008; Davenport et al., 2016). Since stellar activity drives auroral emission for an Earth-like magnetosphere, such features may be much stronger on planets orbiting active M dwarfs. Additionally, with a close-in orbit of 0.0485 AU, Proxima Cen b is about 20× closer to Proxima Centauri than the Earth is to the Sun (Anglada-Escudé et al., 2016). This proximity further increases particle fluxes incident on

the planetary atmosphere that drive ionization and the subsequent recombination radiation.

In addition to increasing the likelihood and strength of the aurora, the characteristics of the Proxima Centauri system may also enhance its detectability. Since the Proxima system is only 1.3pc away, it is perhaps the best-case scenario for the detection of the faint auroral signal from a terrestrial exoplanet. Even though the planet-star contrast ratio in reflected visible light is poor ($\lesssim 10^{-7}$; see Turbet et al., 2016; Kreidberg and Loeb, 2016; Meadows et al., 2016), if Proxima Cen b exhibits auroral emission, this will brighten the planet and potentially boost the planet-star contrast by one or more orders of magnitude at the wavelengths of the auroral emission features. The short wavelength of the oxygen green line also improves the contrast of the planet relative to the star due to the star's cool temperature and TiO absorption, which strongly suppresses the brightness of the star in the visible. This improvement in contrast is significantly less for the near-infrared O_2 1.27μm and OH airglow lines. In addition to increasing the contrast, the small semi-major axis of Proxima Cen b results in an orbital velocity of $\sim$50 km/s, which will cause its auroral emission to be Doppler-shifted by as much as 1Å over the course of its orbit, making it easier to disentangle it from stellar features via high resolution spectroscopy. An additional advantage of the short wavelength of the OI feature is the smaller inner working angle and point-spread function that may be achieved with a coronagraph at that wavelength (Agol, 2007). These factors all improve the chance of detection with ground-based telescopes.

The detection of the oxygen auroral line at 5577Å would provide an important diagnostic for planetary properties. Its detection would not only confirm the existence of the planet, but would point to the presence of an atmosphere with abundant oxygen atoms, which is more likely to indicate a terrestrial body. Additionally, the detection of the line would yield a measurement of the radial velocity (RV) of the planet, which combined with the RV measurements of the star (Anglada-Escudé et al., 2016) would enable the measurement of the eccentricity and inclination of the orbit, ultimately yielding the mass of the planet (see, e.g., Lovis and Fischer, 2010). Detection of the oxygen auroral line would therefore provide several key planetary parameters that could be used to constrain Proxima Cen b's potential habitability (Barnes et al., 2016; Meadows et al., 2016).

This chapter is organized as follows: in §9.2 I calculate the expected auroral emission strength of Proxima Cen b under different assumptions of stellar and planetary properties. In §9.3 I model

the planet-star contrast ratio in a narrow band centered on the OI 5577Å line and calculate the integration times required to detect the feature with different instruments. In §9.4 I conduct a preliminary search for auroral emission in the HARPS high-resolution, ground-based spectroscopy used by Anglada-Escudé et al. (2016) for the RV detection of Proxima Cen b. Finally, in §9.5 I discuss my results and present my conclusions.

9.2 Auroral Signal Strength

Below, I provide quantitative estimates of the auroral intensity for steady-state stellar input. I assume the planet to be terrestrial with the orbital characteristics of Proxima Cen b (see Table 4) and calculate the auroral emission via two different methods. Method 1 (§9.2.3) involves a simple estimation of the emitted electromagnetic auroral power driven by the stellar wind power delivered at the magnetopause of the planet. Method 2 (§9.2.4) uses the prediction of a magnetohydrodynamical (MHD) model that was tuned to calculate the auroral response at Earth, with modifications to the relevant inputs of the stellar wind of Proxima Centauri and assumed planetary parameters for Proxima Cen b (Anglada-Escudé et al., 2016).

The quantities calculated here include only the estimated, localized emissions caused by magnetospheric particle precipitation into a discrete auroral oval — not the diffuse, global phenomenon of airglow. On Earth, the 5577Å airglow can be visible to the naked eye and could be significant on Proxima Cen b, but is commonly driven by different physical processes (e.g., nightside recombination due to dayside photoionization) that are outside the scope of this analysis. Similarly, the 5577Å airglow has been observed at Venus (e.g. Slanger et al., 2001) and Mars (e.g. Seth et al., 2002) — both having no present-day global magnetic field. For these reasons I cannot suggest basing the existence of or placing constraints on Proxima Cen b's planetary magnetic field based on the detection of this auroral line (see, for instance, Grießmeier, 2015). A search for radio emission from Proxima Cen b — which may be correlated with optical auroral emission, as it is on Earth and on Saturn (Kurth et al., 2005) — would likely be necessary to constrain the planetary magnetic field. However, it is worth noting that an Earth-like 1 kR (1 R = 1 Rayleigh $\equiv 10^6$ photons s^{-1} cm^{-2}) airglow across the entire planet would still emit ∼2 orders of magnitude less energy than the discrete polar aurora — see §9.2.6 below.

9.2.1 Stellar winds at Proxima Cen b

M dwarf mass-loss rates, and therefore stellar winds, are not well constrained due to observational sparsity and difficulty (e.g. Wood et al., 2004). To model the M dwarf winds for Proxima Centauri, I adopt the predictions from the modeling efforts of Cohen et al. (2014), who generated an MHD stellar wind model for the M3.5 star EV Lacertae based on available observations. There are two primary differences between EV Lac and Proxima Centauri that should be taken into account when considering the stellar wind at the planet's location of interest: 1) the relative mass-loss rates, 2) the difference in rotation rates.

The first of these factors has been estimated by Wood et al. (2005), who find that the mass-loss per unit surface area for Proxima Centauri and EV Lacertae are quite similar. This suggests comparable wind conditions at equal distances in units of their respective stellar radii.

The second factor, the rotation rate, affects the morphology of the stellar wind magnetic field by changing the Alfvén radius. The Alfvén radius, R_A, is defined as the point where the Alfvén Mach number is equal to unity — i.e., $M_A \equiv u_{sw}/v_A$=1, where u_{sw} is the stellar wind speed and v_A is the Alfvén speed. Interior to R_A (the sub-Alfvénic wind) the magnetic field of the star is mostly radial, and corotates at the angular rate of the star; exterior to R_A (the super-Alfvénic wind) the field begins to lag behind corotation as the magnetic tension is overcome by the flow of the wind. In the super-Alfvénic regime, the interplanetary magnetic field (IMF) exhibits the well-known Parker-spiral (Parker, 1958). The Alfvén point is an important boundary that modifies the energy transfer between the stellar wind and the planetary magnetosphere.

To correctly estimate the interactions, it is important to consider Proxima Cen b's orbital distance from its host star, for both the dynamic parameters (mass density, velocity) and the magnetic structure — i.e., we must consider where Proxima Cen b orbits relative to its Alfvén radius, R_A. Note that the rotational period of Proxima Centauri (82.6 days; Collins et al., 2016) is ∼19 times lower than EV Lacertae (4.376 days; Testa et al., 2004). For my purposes, I estimate an average R_A for a simple stellar dipole moment:

$$R_A = \left(\frac{4\pi {\mathcal{M}_\star}^2}{\dot{M}_\star \omega_\star \mu_0} \right)^{\frac{1}{5}}, \tag{161}$$

where $\mathcal{M}_\star$ is the magnetic dipole moment for the star, $\dot{M}_\star$ is the mass-loss rate, $\omega_\star$ is the angular

frequency of stellar rotation, and μ_0 is the vacuum permeability. For EV Lacertae and Proxima Centauri, R_A are $\sim$65.4 $R_\star$ (0.075 AU) and 115 $R_\star$ (0.192 AU), respectively. This is the average value for a simple dipole moment, as I am not including magnetic topology, but nonetheless the value obtained for EV Lac agrees well with the approximate average for the more complicated magnetic treatment simulated in Cohen et al. (2014). The relative orbit for Proxima Cen b is therefore $\sim$0.76 R_A. Coincidentally, this corresponds well to the simulated Planet B at EV Lac in Cohen et al. (2014), which orbits at $\sim$0.79 R_A.

Recently, Garraffo et al. (2016) applied an MHD model of stellar winds based on the Zeeman-Doppler Imaging (ZDI) of GJ51, and scaled the magnitude of the surface field to match the anticipated value of 600 G for Proxima Centauri. Their results from the assumed magnetic environment are in line with the values I adopt from Table 5, and my value calculated for the magnetopause distance using Equation (163) below is within the range of their calculations for magnetopause distance for Proxima Cen b. However, the structure of the magnetic topology in the simulation of Garraffo et al. (2016) places Proxima Cen b primarily in the super-Alfvénic wind, contrary to both the simple method above and the bulk of the structure found by Cohen et al. (2014).

The estimate of R_A above does not take into account the complicated magnetic topology of a realistic stellar magnetic field, which could indicate the planet likely orbits primarily through sub-Alfvénic conditions (e.g., Fig. 1 of Cohen et al., 2014) or through primarily super-Alfvénic conditions (Fig. 2 from Garraffo et al., 2016). Therefore, I consider both super- and sub-Alfvénic conditions for the steady-state stellar wind, using the reported parameters at Planet B from Cohen et al. (2014); see Table 5.

9.2.2 Magnetic dipole moment of Proxima Cen b

Tidal locking is likely for the expected orbital parameters of Proxima Cen b and the age of the system. One would therefore expect a rotational period equal to the orbital period, 11.186 days, or 8.94% of the Earth's rotational frequency. Following the magnetic moment scaling of Stevenson (1983) and Mizutani et al. (1992), I assume the upper limit of the rotationally-driven planetary dynamo as $\mathcal{M} \propto \omega^{1/2} r_c^3$, where $\mathcal{M}$ is the magnetic moment, ω is the rotation rate of the planet, and r_c is the core radius (which I assume to be proportional to the planetary radius). This suggests a magnetic moment for an Earth-radius Proxima Cen b of $\sim$0.3$\mathcal{M}_\oplus$. Taking the upper limit of

Table 5: Proxima Centauri stellar wind conditions

Quantity	Sub-Alfvénic	Super-Alfvénic
n (cm^{-3})	433	12895
T (10^5 K)	3.42	4.77
$\mathbf{u}$ (km s^{-1})	(-630, -1, 30)	(-202, 102, 22)
$\mathbf{B}$ (nT)	(-804, 173, 63)	(-57, 223, 92)
M_A	0.73	4.76

Note. — Stellar wind conditions from Cohen et al. (2014), at EV Lacertae for a∼51.98 R$_*$ (0.073 AU). n is the stellar wind number density, T is the ion temperature, $\mathbf{u}$ is the velocity, $\mathbf{B}$ is the interplanetary magnetic field (IMF), and M_A is the Alfvén mach number.

the expected radius of Proxima Cen b, 1.4 R$_\oplus$, this gives a magnetic moment of $\sim 0.8\mathcal{M}_\oplus$, which agrees with the upper limit of Zuluaga and Bustamante (2016). However, Driscoll and Barnes (2015) showed that for an Earth-like terrestrial planet orbiting a star of 0.1 M$_\odot$ with high initial eccentricity ($e \geq$0.1) within 0.07 AU, the planet will circularize before 10 Gyr. On this timescale, the orbital energy dissipated as tidal heating is sufficient to drive a strong convective flow in the planetary interior that could generate a magnetic moment in the range of $\sim 0.8 - 2.0\mathcal{M}_\oplus$ during the process of circularization. Given the above, I consider the situation of an Earth magnitude magnetic field for Proxima Cen b, but discuss how each of the methods below can be scaled to various magnetic dipole moments.

9.2.3 Auroral stellar wind power scaling

Desch and Kaiser (1984) suggested a correlation between incident stellar wind power and the power of planetary radio emissions in the solar system, a so-called "radiometric Bode's law." Zarka (2006, 2007) extended the work to modern solar system measurements as well as potential exoplanetary systems, and further suggested that a similar "auroral UV-magnetic Bode's law" could exist, though the author notes such scaling would be less generally applicable than the radio case across planetary systems due to the complexities of UV auroral generation for differing planetary atmospheres and magnetospheric dynamics. The calculations in this section can be

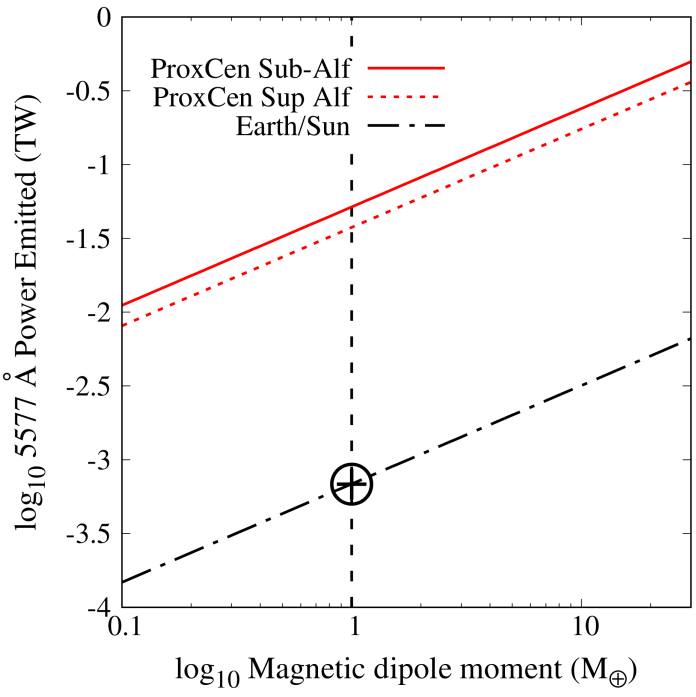

Figure 78: Predicted 5577Å auroral power as a function of planetary magnetic dipole moment calculated using the stellar wind scaling method from §9.2.3. The solid (dotted) red line corresponds to the sub-(super-) Alfvénic stellar wind conditions at Proxima Cen b. The black dash-dotted line corresponds to Earth in its natural orbit around the Sun, and the black Earth symbol corresponds to the method's calculation for Earth. The dashed vertical black line indicates an Earth-equivalent magnetic dipole moment.

thought of similarly as a "visible-kinetic Bode's law" for the specific case of exoplanets with an Earth-like atmosphere. A similar relation may also be derived for the magnetic stellar wind interaction (e.g., a "visible-magnetic Bode's law"; details below).

The stellar wind kinetic power delivered to the magnetosphere of the planet can be expressed as:

$$P_U = \rho\, v^3\, \pi\, R_{MP}^2, \tag{162}$$

where ρ and v are the stellar wind mass density and velocity relative to planetary motion (∼48 km/s), respectively, and R_{MP} is the magnetopause distance along the line connecting the star and planet (sub-stellar point). The latter can be estimated through magnetospheric pressure balance

with the stellar wind dynamic ram pressure (e.g. Schield, 1969):

$$\frac{2\,\mathcal{M}^2}{K_{SW}\mu_0\,R_{MP}^6} = p_{ram}, \tag{163}$$

where K_{SW} is related to particle reflection at the magnetopause (herein the interaction is assumed to consist of inelastic collisions, or K_{SW}=1) and R_{MP} is the distance from the planet at which the magnetic pressure of the planet balances the pressure of the stellar wind. The RHS of Equation (163) represents the dynamic ram ($p_{ram} = \rho v^2$) pressure of the stellar wind, calculated from the values in Table 5.

I consider an Earth-strength magnetic dipole moment of $\mathcal{M} = 8.0{\times}10^{15}$ Tesla m^3. Solving for R_{MP} in Equation (163) and inserting into Equation (162) provides an estimate of the stellar wind power incident on the planetary magnetopause. Externally-driven planetary auroral systems are not typically 100% efficient at converting the incident stellar wind power into electromagnetic auroral emission, and range from ∼0.3% at Neptune, ∼1% at Earth, and up to almost 100% at Jupiter (e.g. Cheng, 1990; Bhardwaj and Gladstone, 2000). For reference, at Earth, this method gives us a reasonable estimate of the total emitted electromagnetic auroral power of ∼30 GW for nominal solar wind conditions (4 cm^{-3}, 400 km s^{-1}), which is consistent with the anticipated power of 1-100 GW, depending on solar and magnetospheric activity. While the intensities of various emissions vary widely with activity and atmospheric conditions, I assume an averaged auroral emission. In order to estimate the emitted power of the OI 5577Å line, I assume it represents 2% of all emitted electromagnetic power (Chamberlain, 1961; Kivelson and Russell, 1995), as calculated by Equation (162) I note that this assumes an Earth-like atmosphere for the planet; I briefly discuss the effect of different atmospheric compositions in §9.5.

Fig. 78 shows the predicted emitted power of the 5577Å line based on Equation (162) and multiplied by the 2% factor mentioned above and by the conversion efficiency of 1%. For the Earth, this method predicts a power of $\phi_\oplus$ ∼0.68 GW in the 5577Å line. Assuming a 5° latitudinal width starting at ∼18° co-latitude and extending equatorward, this corresponds to a photon flux of ∼13.7 kR. This is in agreement with moderate auroral activity (IBC II[19], 10 kR 5577Å emission; see Table II.1 in Chamberlain, 1961), and within a factor of 2–5 of observations during moderate

[19]IBC = International Brightness Coefficients, a standardized scale for quantifying auroral intensities (see, e.g., Hunten, 1955).

geomagnetic disturbance (2.5–6 kR 5577 Å emission, e.g. Steele and McEwen, 1990).

Power estimates for the 5577Å line for a 0.05 AU orbit around Proxima Centauri are shown in Table 6. The calculated power is ∼75.3 (54.7) times $\phi_{\oplus}$, in the sub-(super-) Alfvénic stellar wind. These are the estimates for a steady-state stellar wind, for a terrestrial planet with an Earth-like magnetic dipole moment. Note that by inspection of Equations (162) and (163), one can see that the expected power scales as $\mathcal{M}^{2/3}$, and so can easily be extended to different planetary dipole moments.

The method above has a weakness in that it completely ignores the incident Poynting flux from the IMF, and potential direct magnetic interactions between the stellar wind and planetary magnetic field, e.g. flux merging or reconnection. These interactions can produce a significant amount of magnetospheric energy input, and so they are important to consider. Similar to Equation (162), a scaling relation between power emitted at the 5577 Å line and incident magnetic flux in the stellar wind, akin to a visible-magnetic Bode's law, can be given as (e.g., Zarka, 2006, 2007; Grießmeier et al., 2007):

$$P_B = \epsilon\, K \left(\frac{v\, B_{\perp}^2}{\mu_0} \right) \pi\, R_{MP}^2 \tag{164}$$

where ϵ is the efficiency of reconnection (typically of order 0.1–0.2), K is related to the "openness" of the magnetosphere, and for an Earth-like dipole is $K = \sin^4(\theta/2)$ where θ is the angle between the perpendicular IMF and planetary dipole field, $B_{\perp}$ is the perpendicular IMF ($\sqrt{B_Y^2 + B_Z^2}$), μ_0 is the vacuum permeability, and R_{MP} is the magnetopause sub-stellar point discussed above. We can estimate the magnetic interaction at Proxima Cen b using the stellar wind conditions above by taking the ratio of Equations (164) and (162):

$$\frac{P_B}{P_U} = \frac{\epsilon\, K\, B_{\perp}^2}{\mu_0\, \rho\, v^2}, \tag{165}$$

which is essentially the ratio of the perpendicular IMF magnetic pressure to the ram pressure, modulated by magnetic field orientation and reconnection efficiency. In the best case scenario, K is equal to 1 (indicating θ=π, driving strong reconnection at the magnetopause), and ϵ is of order 0.2. Assuming this best case, and inserting the values from Table 5 for the sub- and super-Alfvénic cases, one obtains a ratio of ∼0.019 and 0.0084 for the sub- and super-Alfvénic cases,

Table 6: Calculated 5577Å auroral power, by method

Case	Meth. 1 [TW]	Meth. 2 (quiet) [TW]	Meth. 2 (SS) [TW]	Meth. 2 (CME) [TW]	Meth. 2 (CME+SS) [TW]
PCb (Sub)	0.051	0.09	0.24	8.103	21.42
PCb (Sup)	0.038	0.049	0.14	4.41	12.10
Earth/Sun	6.7×10^{-4}	7.5×10^{-4}	1.5×10^{-3}	0.068	0.1317

Note. — Power emitted for the OI 5577Å line in terawatts (TW) for an Earth-strength magnetic dipole on Proxima Cen b in the sub-Alfvénic (Sub) and super-Alfvénic (Sup) stellar winds. For method 2: column 2 assumes no significant stellar activity and a quiet magnetosphere; column 3 assumes geomagnetic substorm (SS) activity; column 4 assumes CME conditions in the stellar winds, but no magnetospheric disturbance; column 5 assumes both CME conditions and substorm activity.

respectively. For the particular stellar wind parameters I have chosen, the kinetic power dominates the anticipated auroral output for Proxima Cen b. It is worth noting, however, that the magnetic environment (both planet and star) is largely unconstrained, and highly dynamic—particularly near active M dwarfs.

9.2.4 3D MHD empirical energy coupling

Equations (162) & (164) above are decent first approximations, but involve significant uncertainties concerning the energy dissipation in physical phenomena throughout the magnetosphere (i.e., auroral activity) (Perreault and Akasofu, 1978; Akasofu, 1981). Wang et al. (2014) developed a global, 3D MHD model to obtain a fit for the energy coupling between the solar wind and Earth's magnetosphere to estimate the energy transferred directly from the wind into the magnetosphere and auroral precipitation (see their Equation 13, and below). The simulations were performed over 240 iterations across their solar wind parameter space, and the resulting nonlinear fit for the energy transfer to the terrestrial magnetosphere was found to be

$$P_{trans} = K_1\, n_{sw}^{0.24}\, v_{sw}^{1.47}\, B_T^{0.86}\, \left[\sin^{2.7}(\theta/2) + 0.25\right], \tag{166}$$

where $K_1 = 3.78 \times 10^7$ is a coupling constant, n_{sw} and v_{sw} are the stellar wind number density (in cm^{-3}) and velocity relative to planetary motion (in km s^{-1}), respectively, B_T is the magnitude of the transverse component of the Sun's IMF ($B_T = \sqrt{B_X^2 + B_Y^2}$) in nT, and θ is the so-called IMF

clock angle ($\tan\theta = B_Y/B_Z$). The coordinate system used is the geocentric solar magnetospheric (GSM) system, with $\hat{X}$ pointing from the planet to the star, $\hat{Z}$ aligned with the magnetic dipole axis of the planet (here assumed to be perpendicular to the ecliptic), and $\hat{Y}$ completing a right-handed coordinate system.

Wang et al. (2014) were focused solely on the Earth's magnetosphere, but one can scale to any dipole moment by noting that Equation (166) scales just as in §9.2.3: the dipole moment term is implicitly included in the coupling constant K_1 above and scales with the planetary magnetic dipole magnitude as $\mathcal{M}_P^{2/3}$ (Vasyliunas et al., 1982, also Equations 162 & 163 above).

Equation (166) is the total power delivered by the stellar wind to the magnetosphere, which Wang et al. (2014) estimate is ∼13% of the total incident stellar wind energy. They further estimate that 12% of that energy is dissipated by particle precipitation into the auroral regions, yielding a total solar wind/auroral coupling efficiency of ∼1.56% – very similar to the efficiency value of 1% assumed for Earth and Proxima Cen b in §9.2.3. As simple validation for my purposes, I use this method to predict a maximum coupling of auroral particle precipitation (with IMF clock angle $\theta = \pi$, driving reconnection and likely substorm activity) at Earth of ∼0.17 TW. This is in agreement with terrestrial plasma observations during periods of geomagnetic disturbance (e.g. Hubert et al., 2002). This method is useful in that it provides a direct relationship between the power delivered as auroral particle precipitation and incident stellar wind conditions.

For Proxima Cen b subjected to the stellar winds from Table 5, this method predicts a total power of auroral particle precipitation of ∼10.7 (5.8) TW for the sub-(super-)Alfvénic stellar wind. The stellar wind parameters in Table 5, however, are a snapshot and not indicative of the highly variable conditions likely experienced at Proxima Cen b.

Magnetospheric substorms, related to transient populations of energized particles driven by magnetic reconnection in the magnetotail, can drive strong increases in auroral particle precipitation. Though not a one-to-one indicator, substorm activity can be associated with periods of strong reconnection at the magnetopause—correlated with a significant negative B_Z component in the IMF. In this chapter, I assume $\theta=\pi$, or $B_Y=0$, to obtain an upper limit to substorm influence under my model. Although this is not a strict definition, Wang et al. (2014) calibrated the model used here to include periods of substorm activity and high hemispheric energy input. Assuming with this strong negative B_Z that a substorm is driven at Proxima Cen b, I predict an energy

input of ∼28.3 (15.9) TW for the sub-(super-)Alfvénic wind. To compare directly to the 5577Å line auroral power output such as that calculated in §9.2.3, we must link these values to the aurora by including the efficiency of precipitating charged particles in the production of auroral emission for the 5577Å line, which will be done below.

To calculate the auroral 5577Å photon flux, I use the precipitating auroral particle powers above obtained from Equation (166), and combine with the anticipated size of the auroral oval and an observed conversion efficiency for electron precipitation to 5577Å emission. This gives the photon flux in kR, ϕ_{5577}:

$$\phi_{5577} = P_{in}\, A_{mag}^{-1}\, \epsilon_e, \tag{167}$$

where P_{in} is 12% (discussed above) of P_{trans} from Equation (166), A_{mag} is the summed area of both the northern and southern auroral ovals (I assume N-S symmetry), and ϵ_e is the efficiency with which magnetospheric electrons are converted to auroral emission of the 5577Å oxygen line. I use the reported values from Steele and McEwen (1990) (noted below), who used ground-based observations of auroral line intensities and the related satellite observations of energetic electron flux to draw a relation between electron precipitation and auroral photon flux. I then integrate the resulting flux over a nominal 5° auroral oval (for each hemisphere), the colatitude of which is dependent on the sub-stellar magnetopause distance (discussed below).

Steele and McEwen (1990) reported the conversion efficiency for the 5577Å OI line as 1.73±0.51 (1.23±0.44) kR/(erg cm^{-2} s^{-1}) for a magnetospheric Maxwellian electron population of characteristic temperature 1.8 (3.1) keV. In the present work, I take the average values for these populations, ∼1.48 kR/(erg cm^{-2} s^{-1}). I assume the fraction of total hemispheric power (P_{in}) delivered by electrons to be 0.8 (Hubert et al., 2002), so this factor is included in the P_{in} factor.

The magnetopause distance I calculate via Equation (163) for the Earth-like magnetic dipole moment is ∼4.2 (3.3) R_P for the (sub-)super-Alfvénic conditions. From these values, I can provide a simple estimate of the total auroral oval coverage. The magnetic co-latitude of the boundary between open and closed flux for my assumed ideal planetary dipole geometry (i.e., the co-latitude where the field structure no longer intersects the planetary surface) is $\sin^{-1}(1/\sqrt{R_{MP}})$ (Kivelson and Russell, 1995). Discrete auroral activity occurs primarily due to energized plasma originating from closed field structure stretched out behind the planet in the stellar wind, i.e., the magnetotail.

This field structure intersects the planet equatorward of the open/closed boundary co-latitude. Assuming a nominal 5° auroral oval width beginning at the co-latitude obtained, and extending equatorward, I calculate a single-hemisphere coverage of $\sim 1.17 \times 10^{17}$ cm^2 for the auroral oval under sub-Alfvénic conditions, and $\sim 1.30 \times 10^{17}$ cm^2 under super-Alfvénic conditions.

Following the above, I obtain a photon flux value of ϕ_{5577} ∼2.26 (1.16) MR for the sub-(super-)Alfvénic wind conditions. This corresponds to the predicted emission power in Table 6, method 2 (quiet) of ∼0.090 (0.049) TW under steady-state sub-(super-)Alfvénic conditions. For the maximum emission during a magnetospheric substorm, I obtain values of ∼0.24 (0.14) TW for sub-(super-)Alfvénic winds.

There is another case of interactions that we should consider that involves stellar activity — flaring and coronal mass ejections (CME). During these events, stellar wind densities could increase by a factor of ∼10, velocities by a factor of ∼3, and IMF magnitude by a factor of $\sim 10 - 20$ (Khodachenko et al., 2007; Gopalswamy et al., 2009). Inserting such ratios in the 3D MHD-fit predicted power in from Equation (166), I predict transient maximum 5577Å emissions of ∼8.10 (4.40) TW for the sub-(super-)Alfvénic CME conditions. For the maximum emission during a magnetospheric substorm under CME conditions, I obtain values of ∼21.42 (12.10) TW for sub-(super-)Alfvénic winds. These transient CME conditions can have timescales of $\sim 10 - 10^3$ minutes per event, with multiple, consecutive events possible. Given that Davenport et al. (2016) report such high stellar activity for Proxima Centauri, Proxima Cen b could experience CME impacts for a large percentage of its orbital phase (e.g., Khodachenko et al., 2007).

9.2.5 Unmagnetized planet

The above results all assume a large, Earth-like planetary magnetic dipole moment for Proxima Cen b. If, in fact, the planet does not sustain a global dynamo, it will only be protected by a relatively thin spherical shell (∼1000 km) of plasma in the upper atmosphere - similar to Earth's ionosphere.

For the sub-Alfvénic case, the interaction is a unipolar interaction similar to the Jupiter-Io interaction (Zarka, 2007). In this case, the power dissipated by the wind is similar to the form of Equation (164), where ϵ and K are replaced by a single parameter indicating the fraction of magnetic flux convected onto the "obstacle" (the ionosphere), and R_{MP} becomes the size of the

"obstacle" — for an Earth-like ionosphere, ∼1.16 planetary radii. While there is no dipolar focusing mechanism for the particle precipitation in this case, it is worth considering the energized particles flowing on the flux tube connecting the unmagnetized planet with the star, producing maxima on the unmagnetized body in the plane perpendicular to the IMF (see, e.g., Saur et al., 2000, 2004).

Assuming 100% of incident magnetic energy flux is convected onto the planet and ionosphere, the expected 5577Å auroral power becomes 5.9×10^{-4} TW for the sub-Alfvénic stellar wind conditions in Table 5. This interaction is likely insignificant in the context of remote sensing.

For the super-Alfvénic flow, this could be considered as analogous to Venus' situation, which sustains no global magnetic field. In this case, the discrete aurora would obviously not be expected due to a lack of magnetic structure, though induced airglow is still a consideration. Lacking a planetary magnetosphere, the magnetic structure fails to focus precipitating particles into the upper atmosphere of such a planet, though there is still magnetic interaction at the planet. The ionosphere is a spherical, conducting shell, and so interacts with the magnetic flux from the IMF as it drapes over and around the planet. Energized particles in the impacting stellar wind magnetic flux could still dip down into the upper atmosphere, depositing sufficient energy to produce airglow - this is especially true for the strong flows from CME activity, or fast stellar wind flow.

A study of the intensity of 5577Å oxygen emission at Venus, relative to Earth, for CME/flare events from the Sun was performed by Gray et al. (2014). The results indicated that the airglow was relatively on par with that of Earth's upper atmosphere, varying between 10 and a few hundred Rayleigh, which, if integrated over an entire hemisphere of an Earth-like planet, gives a value less than 1% of the discrete values given in Table 6. For the super-Alfvénic flow in Table 5, the number density is ∼1500 times greater than the average at Venus, and the velocity is a factor of ∼0.5 that at Venus or Earth. Given the power scales as $\rho\, v^3$, this is a factor of ∼200 greater power delivered to the planet. Assuming that airglow at an unmagnetized Proxima Cen b scales linearly with the incident power, this would give a brightness of ∼2-60 kR, which is at most a factor of ∼5 times the value for Earth using Method 1 and 2 in Table 6, or on the order of 10^{-3} TW. Given the discussion of auroral detectability below (§9.3), I do not expect that this signal could be observed with either current or upcoming missions.

9.2.6 Signal Summary

The preceding estimates are mostly conservative. It is possible that all the auroral numbers reported for the sub-Alfvénic cases above could be a factor of 4–5 (or more) larger. I am assuming a simple dipolar interaction with the stellar wind, which isn't specifically the case for a planetary dipole in the sub-Alfvénic stellar wind; these interactions are more akin to the interactions of Ganymede and Io with the corotating magnetosphere of Jupiter, with the formation of Alfvén wings. Modeling efforts by Preusse et al. (2007) showed that for a giant planet with a dipole magnetic moment, field-aligned currents (which are associated with auroral activity) are significantly stronger for planets orbiting inside the Alfvén radius of their stellar host. My estimates, therefore, could be viewed as lower limits. It is also worth noting that Cohen et al. (2014) suggested that a transition between the sub- and super-Alfvénic conditions would likely produce enhanced magnetospheric activity and therefore could lead to a periodicity in the auroral activity depending on combined planetary orbital and stellar rotational phases.

For planets in the solar system, only Mars and Earth exhibit significant 5577Å emission for both diffuse airglow and discrete aurora. Mars does not presently have a global magnetic field, but there are crustal regions containing the remnants of previous magnetization that exist and focus particles into the upper atmosphere to produce a relatively weak (inferred ∼30 R at 5577Å) discrete aurora that is ∼10 times the strength of the nominal airglow (e.g. Acuña et al., 2001; Bertaux et al., 2005; Lilensten et al., 2015). On Earth, the airglow and aurora are typically in the range 0.01–1 kR and 1–1000 kR, respectively. During transient periods of minimal auroral activity and maximum airglow emission, emissions can be roughly equivalent, but the average ratio of airglow emission to auroral emission is ≤1% (e.g. Chamberlain, 1961; Greer et al., 1986). Even for a constant, planet-wide 1 kR airglow on an Earth-sized planet at Proxima Cen b (R_P ∼6371 km), the total signal from the observer-facing hemisphere would be ∼4.54×10^8W, which is ∼1% of the lowest signal from Table 6 and would not be detectable.

Nevertheless, it is important to note that the FUV flux from Proxima Centauri is nearly two orders of magnitude higher than that of the Sun (Meadows et al., 2016). Airglow stemming from recombination of photodissociated O_2 and CO_2 could thus be significantly stronger on Proxima Cen b than on Earth. Barthelemy and Cessateur (2014) stress the importance of stellar UV/FUV emissions on the production of UV and visible aurorae, and note that, e.g., Lyman-α flux can

contribute up to 25% to the production of the O(^1D) red-line. However, even if Proxima Cen b had a sustained 100 kR airglow—one hundred times the maximum Earth airglow—its emission would be comparable to the lowest estimate of auroral emission in Table 6, which is still unlikely to be detectable (see §9.3). I therefore ignore this potential contribution in the present work, noting that a detailed photochemical treatment would be required to pin down the expected airglow emission at Proxima Cen b.

In summary, I predict a steady-state auroral emission at 5577Å from Proxima Cen b that is of order 100 times stronger than seen on Earth for a quiet magnetosphere, corresponding to an emitted auroral power for the OI line on the order of ∼0.090 (0.049) TW for the sub-(super-)Alfvénic winds using method 2 (§9.2.4). I believe that this method yields more realistic results than the purely kinetic power estimate in method 1 (§9.2.3), due to the inclusion of magnetic interactions in method 2 — though the magnitudes are similar to within a factor of 2. Assuming Proxima Cen b is an Earth-like terrestrial planet, my maximum transient power estimate for the 5577Å line for CME conditions that drive a magnetospheric substorm is ∼21.42 TW, or ∼30,000 times stronger than on Earth under nominal solar wind conditions. The actual values for Proxima Cen b will naturally change based on planetary parameters (e.g., magnetic dipole moment, magnetospheric particle energy distributions, substorm onset, atmospheric Joule heating) and stellar activity. By my analysis, a $\sim 10^3$ (or higher) enhancement compared to Earth as suggested by O'Malley-James and Kaltenegger (2016) is only possible due to one or more of the following: transient magnetospheric conditions driven by either CME or substorm activity, a magnetic dipole significantly stronger than Earth's, or higher stellar mass-loss than predicted (Wood et al., 2005; Cohen et al., 2014).

9.3 Auroral Detectability

In this section I assess the detectability of the 5577Å OI auroral emission line from the atmosphere of Proxima Cen b. Below, I investigate the line profile shape and then calculate planet-star contrast ratios and integration times required for auroral detection.

9.3.1 OI Auroral Line Profile

To estimate the signal-to-noise as a function of spectral resolution, I need to estimate the auroral spectral line width. The OI 5577Å green line has no hyperfine structure and because it is a

forbidden line, it has negligible ($\lesssim 10^{-15}$Å) natural width (Hunten et al., 1967). Spectroscopic observations of the OI airglow by Keck/HIRES (Slanger et al., 2001) and by HARPS (Anglada-Escudé et al., 2016, see §9.4) are unresolved, revealing the resolution element width of the instrument used for the observation at the wavelength of the line (∼0.1Å Keck/HIRES; ∼0.05Å HARPS) rather than the full width at half maximum (FWHM) of the line.

To determine the width of the line, I examine several line broadening mechanisms that play a key role in terrestrial atmospheres. The planet's rotation will broaden the OI line, but calculations by Barnes et al. (2016) and Ribas et al. (2016) show that Proxima Cen b is likely tidally locked with a rotation period of 11.2 days, resulting in negligible instantaneous rotational broadening ($FWHM = 0.002$Å). Pressure broadening can also be safely neglected since OI auroral emission occurs in terrestrial atmospheres at an elevation of ∼100 km where the atmosphere is thin (Slanger et al., 2001). Similarly, broadening due to atmospheric turbulence can also safely be neglected due to the stratospheric origin of the line. Thermal Doppler broadening should therefore be the dominant line broadening mechanism, resulting in a Gaussian line profile. For the 5577Å OI line, Doppler broadening gives the following scaling relation:

$$FWHM = 2\Delta\lambda = 0.014 \left(\frac{\mathrm{T}}{200 \text{ K}} \right)^{1/2} \text{ Å}, \tag{168}$$

where T is the temperature of the emitting layer, for which I adopt the value of 200 K (c.f. Slanger et al., 2001). A FWHM of 0.014Å is in good agreement with the Fabry-Perot interferometric line width measurements of Wark (1960).

Given the relatively short period of Proxima Cen b and the long exposure times expected for high resolution spectroscopy, I must also consider the possibility of broadening due to the orbital motion of the planet over the course of an observation. One could take a series of shorter exposures, but this strategy will introduce significant instrumental noise, which is likely to overwhelm any planetary signals. In Fig. 79 I plot the orbital broadening of the 5577Å line as a function of the exposure time, calculated from the maximum change in the radial velocity of the planet over the course of the observation and assuming an inclination of 90°. The effect is strongest at full and new phases (dotted line), where the time derivative of the radial velocity is highest, and weakest at quadrature (solid line), where the derivative is smallest. Two intermediate phases are also shown. The FWHM given by Equation (168) is indicated by a horizontal red line; orbital broadening

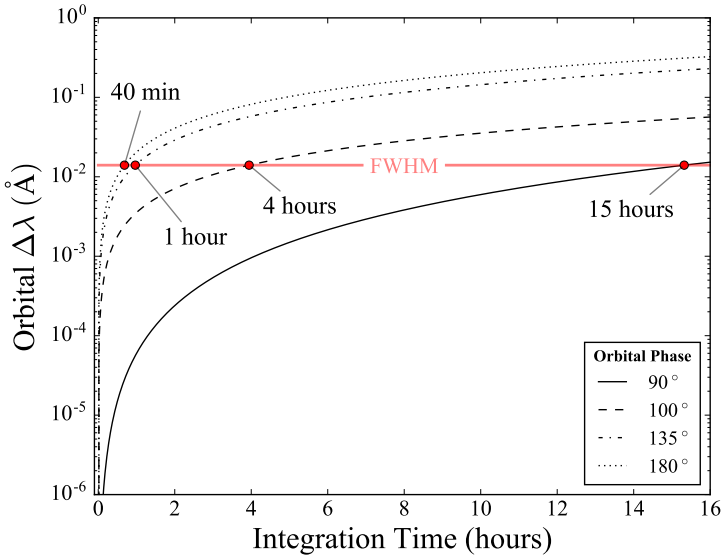

Figure 79: Orbital broadening of the 5577Å OI line as a function of the exposure time for observations made at different orbital phases: 90° (quadrature), 100°, 135°, and 180° (full phase). The FWHM given by Equation (168) (0.014Å) is indicated by a horizontal red line; the intersection of this line with the black curves corresponds to the integration time for which the FWHM doubles. At quadrature, exposures up to ∼6 hours long have a negligible effect ($\Delta\lambda \lesssim 10^{-3}$Å) on the width of the line. At all other phases, the broadening is larger and can cause a significant increase in the FWHM in ∼1 hour.

becomes significant as the curves approach this line. In general, observations made at quadrature with exposure times up to ∼6 hours cause negligible broadening. At all other phases, however, broadening becomes significant in a matter of one or a few hours. At full and new phase, the line width doubles after an exposure of only 40 minutes. However, at these phases the radial velocity of the planet relative to the star is zero, and as I argue in §9.4 below, disentangling stellar and planetary emission becomes difficult. In the discussion that follows, I therefore focus on observations made close to quadrature.

Fig. 80 shows a high-resolution model spectrum of Proxima Cen b at quadrature, illustrating an auroral emission feature that could be expected from the planet. I injected a Gaussian line at 5577Å with FWHM = 0.014Å, normalized to a steady-state Proxima Cen b OI auroral power of $L_{OI} = 0.1$ TW. This OI auroral power yields an equivalent width of $\sim 3.63(L_{OI}/1\text{ TW})$Å relative to my model of the reflected planetary spectrum at quadrature.

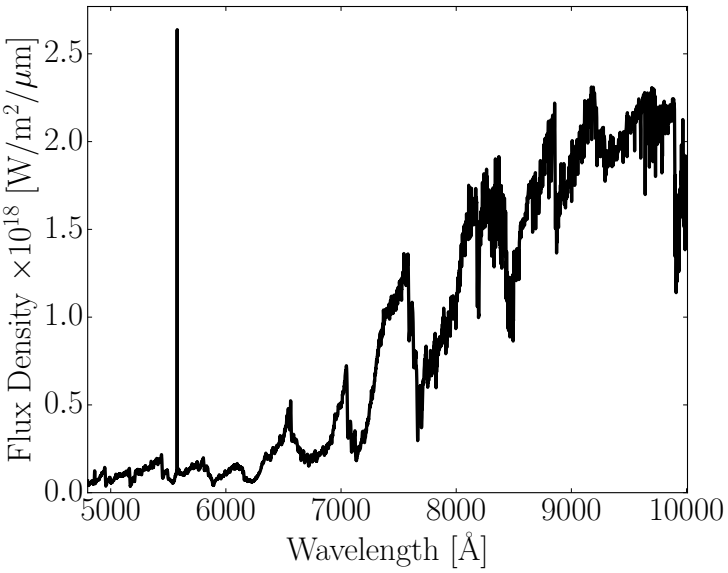

Figure 80: Simulated high-resolution visible spectrum of Proxima Cen b with a 0.1 TW OI auroral emission at 5577 Å. A grey geometric albedo of 0.3 is assumed for the planet. The spectrum is calculated at quadrature phase and scaled to the observing distance (1.302 pc).

9.3.2 Contrast Ratios & Telescope Integration Times

To unambiguously detect a narrow emission feature, such as the example shown in Fig. 80, the telescope resolving power ($R \equiv \lambda/\Delta\lambda$) needs to be taken into account. The typical resolving power being considered for future space-based coronagraph mission concepts is $R \approx 100$, which is appropriate for the detection of molecular absorption bands in the optical and NIR given the relatively low planet-star contrast ratio (Robinson et al., 2016). However, at 5577Å an $R = 100$ spectrograph has a spectral element width of $\Delta\lambda \approx 56$Å, over $\sim 10^3$ times broader than the OI green line width. Future space-based high-contrast exoplanet imaging missions would need to fly with higher resolution spectrographs to detect the OI 5577Å line.

Fig. 81 shows planet-star contrast ratios in a spectral element centered on the 5577Å OI auroral line as a function of spectrograph resolving power and auroral power, assuming a FWHM of 0.014Å (i.e., negligible orbital broadening). The FWHM of the auroral line and equivalent width (W_λ) as a function of auroral power are represented as "resolving powers," where $R_{FWHM} = \lambda_{\rm OI}/FWHM$ and $R_W = \lambda_{\rm OI}/W_\lambda$, respectively. In Fig. 81, the dashed-white line gives the resolving power such that the spectral element width matches the equivalent width of the line at a given auroral power. The dashed-orange line gives the resolving power such that the spectral element width

matches the FWHM of the line. That is, a fixed FWHM=0.014Å yields $R_{FWHM} = 4 \times 10^5$. Optimal observations occur when the planet-star contrast (indicated by the contours) is highest. An increase in the contrast of the emission line is only achieved when the width of a spectral element is smaller than the equivalent width of the line. For resolving powers greater than R_{FWHM}, multiple spectral elements are needed to span the emission line, which may introduce additional unwanted read noise and dark current. Therefore, observations should be made in the wedge between the FWHM resolving power and the equivalent width resolving power. The predicted steady-state auroral emission (∼0.1 TW) requires that spectrographs achieve $R \gtrsim 10^5$.

Using the auroral power estimates from §9.2, I explore the feasibility of detecting the 5577Å OI auroral emission line with five different ground-based telescope configurations: the 3.6m High Accuracy Radial velocity Planet Searcher (HARPS), the 8.2m Very Large Telescope (VLT) with and without a coronagraph, and a Thirty Meter Telescope (TMT) concept with and without a coronagraph (Skidmore et al., 2015; Udry et al., 2014; Johns et al., 2012). I also model the detection using two future space-based coronagraph concepts: the 16m Large UV/Optical/IR Surveyor (LUVOIR; Kouveliotou et al., 2014; Dalcanton et al., 2015) and the 6.5m Habitable Exoplanet Imaging Mission (HabEx; Mennesson et al., 2016).

High spectral resolution coronagraphy with the VLT will require an update to the SPHERE high-contrast imager and a coupling with the ESPRESSO spectrograph, as described in Lovis et al. (2016). Note that given the VLT's 8.2m diameter, the SPHERE coronagraph must achieve an inner working angle no more than $\theta_{\mathrm{IWA}} = 2.7\lambda/D$ to observe at wavelengths as long as 5577Å given the maximal planet-star angular separation of 37 mas for Proxima Cen b. My HARPS, TMT, LUVOIR and HabEx telescope models use $R = 115,000$, while for VLT I use $R = 120,000$. All models assume a total telescope and instrument throughput of 5% and a quantum efficiency of 90%. Coronagraph noise estimates use the model presented in Robinson et al. (2016) with updated parameters from Meadows et al. (2016), and consider noise due to speckles, dark counts, read noise, telescope thermal emission, and zodi and exozodi light. Ground-based coronagraphy assumes a conservative design contrast of 10^{-5} (Dou et al., 2010; Guyon et al., 2012), while space-based assumes 10^{-10} (Meadows et al., 2016) unless stated otherwise. Typically, telescope detectors have a maximum exposure time to mitigate the damaging effect of cosmic ray strikes (see Robinson et al., 2016). Therefore, integration times longer than one hour require multiple readouts,

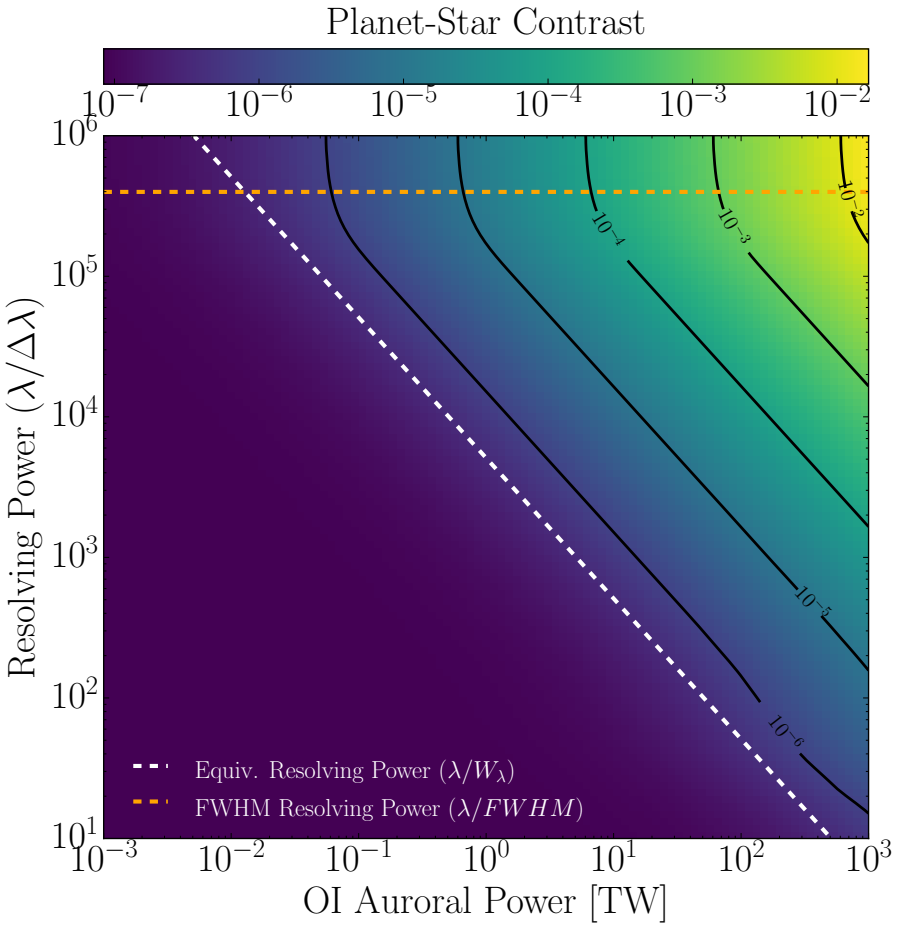

Figure 81: Planet-star contrast ratio contours as a function of telescope resolving power and OI auroral power. The full width at half maximum (FWHM; dashed-orange) of the line and equivalent width (W_λ) as a function of auroral power (dashed-white) are represented as "resolving powers", where $R_{FWHM} = \lambda_{\rm OI}/FWHM$ and $R_W = \lambda_{\rm OI}/W_\lambda$, respectively. The black contour lines show curves of constant planet-star contrast.

Table 7: Planet-Star contrast ratios and telescope integration times necessary to detect the 5577Å OI auroral line

		Telescope Integration Time [hours]							
Power [TW]	Contrast	HARPS	VLT	VLT + C	TMT	TMT + C	HabEx	LUVOIR	TMT + C*
0.001	9×10^{-8}	2×10^{13}	4×10^{12}	4×10^{8}	3×10^{11}	1×10^{7}	6×10^{8}	2×10^{7}	1×10^{5}
0.01	2×10^{-7}	2×10^{11}	4×10^{10}	4×10^{6}	3×10^{9}	1×10^{5}	6×10^{6}	2×10^{5}	2×10^{3}
0.1	$\mathbf{8 \times 10^{-7}}$	$\mathbf{2 \times 10^{9}}$	$\mathbf{4 \times 10^{8}}$	$\mathbf{4 \times 10^{4}}$	$\mathbf{3 \times 10^{7}}$	$\mathbf{1 \times 10^{3}}$	$\mathbf{6 \times 10^{4}}$	$\mathbf{2 \times 10^{3}}$	**40**
1	7×10^{-6}	2×10^{7}	4×10^{6}	4×10^{2}	3×10^{5}	10	7×10^{2}	30	3
10	7×10^{-5}	2×10^{5}	4×10^{4}	7	3×10^{3}	4×10^{-1}	10	1	3×10^{-1}
100	7×10^{-4}	2×10^{3}	4×10^{2}	4×10^{-1}	30	3×10^{-2}	6×10^{-1}	9×10^{-2}	3×10^{-2}
1000	7×10^{-3}	20	4	3×10^{-2}	3×10^{-1}	3×10^{-3}	6×10^{-2}	9×10^{-3}	3×10^{-3}

Note. — Integration times refer to the time required to achieve a signal-to-noise of 6 on the auroral emission above the continuum assuming a telescope throughput of 5%, a spectrograph with resolution $\lambda/\Delta\lambda = 115,000$ for HARPS, TMT, HabEx and LUVOIR and $\lambda/\Delta\lambda = 120,000$ for VLT. "+ C" indicates the use of a coronagraph and associated noise sources discussed in Robinson et al. (2016). Auroral power of order 0.1 TW (boldface) corresponds to the predicted steady-state value in §9.2 while $\sim 1 - 100$ TW corresponds to the predicted auroral power arising from a combination of substorm event and CMEs. TMT + C* denotes a coronagraph-equipped TMT concept with a design contrast of $C = 10^{-7}$ and negligble instrumental noise.

introducing more detector noise. Non-coronagraph telescope calculations assume only stellar noise at the photon limit; their values are therefore lower limits, and may increase significantly due to stellar activity (see §9.4). To prevent significant orbital broadening, I assume that observations are made for one hour at a time at or close to quadrature; longer exposure times are achieved by stacking multiple observations. For exposure times much longer than an hour, stacking will appreciably increase the read noise and dark current for coronagraph observations, where the star is nulled, but not for the non-coronagraph observations, where the stellar photons dominate the noise budget.

My integration time calculations follow those described in Robinson et al. (2016). For the stellar spectrum I adopt the steady-state Proxima Centauri spectrum of Meadows et al. (2016) and neglect the impact of flares on the stellar continuum. I assume that the quoted auroral power emitted via the 5577Å OI line is constant throughout the entire observation.

For observations without a coronagraph, both the stellar flux and reflected stellar flux define the continuum from which we wish to resolve the auroral emission feature. Observations with a coronagraph need only resolve the auroral emission above the coronagraph noise and reflected stellar continuum. With these considerations in mind, I simulate the net planetary emission as a combination of reflected stellar continuum and auroral emission. I compute the flux from the reflected stellar continuum by assuming that the planet is a Lambertian scatterer at quadrature with a planetary geometric albedo of 0.3 and a planetary radius of $1.07 R_\oplus$ following Barnes et al. (2016). I then inject the expected flux from the auroral line at its wavelength. I integrate over all spectral elements that contain the auroral line flux, taking the auroral photon count rate as the signal and all other sources as noise as in Robinson et al. (2016). For the oxygen 5577Å line width of ∼0.014Å (§ 9.3.1) and my nominal resolving power, this corresponds to one spectral element.

Table 7 shows the integration times required to achieve a signal-to-noise of 6 on the 5577Å OI auroral emission line above the stellar and reflected planetary continuum as a function of auroral power. I simulated contrast ratios and integration times for emitted auroral powers at the OI 5577Å line ranging from $10^{-3} - 10^{3}$ TW to bracket all potential auroral fluxes. The 10^{-3} TW lower limit corresponds to a strong 5577Å emission from Earth (Earth total electromagnetic auroral power is of order 10^{-2} TW with 5577Å typically ∼2% of this value). The upper limit of 10^{3} TW is an extreme case that is an order of magnitude stronger than the largest value predicted in §9.2. Values

in between correspond to the different cases considered in Table 6, which depend on the planetary dipole moment, magnetospheric substorm activity, and whether CME conditions are present. For reference, the estimated steady-state Proxima Cen b value calculated in §9.2 is ∼0.1 TW.

The weak 10^{-3} TW aurora is indistinguishable from the purely reflecting planet-star contrast near the 5577Å OI auroral emission line (Turbet et al., 2016; Meadows et al., 2016) and effectively demonstrates why high resolution spectroscopy is not typically considered for high-contrast imaging. For an Earth-like planet, the auroral power estimates from §9.2 (∼0.1 TW) make detecting the OI emission line infeasible with current instruments, even though the contrast ratio at the line is relatively strong ($\sim 8 \times 10^{-7}$). Although unlikely, if the auroral power were much higher ($\sim 10^3$ TW) and sustained over the period of the observation, detection of OI emission could be possible with current instruments in tens of hours. Realistically, however, these estimates suggest that current instruments are likely not capable of detecting an OI aurora on Proxima Cen b.

For a SPHERE-ESPRESSO coupling (Lovis et al., 2016), the integration times required to detect an OI auroral line are slightly more favorable over a wide range of possible auroral powers, but still prohibitively long under most plausible circumstances. If Proxima Cen b has a Neptune-strength magnetic dipole moment, and observations were made during substorm conditions (when the power in the OI 5577Å line reaches ∼1 TW with a contrast ratio of 7×10^{-6}), a coronagraph-equipped VLT would have to integrate for ∼400 hours. However, if observations coincided with periods of more vigorous stellar activity such as during a concurrent CME and substorm, the auroral output could reach ∼100 TW and contrast ratios of 7×10^{-4}. An upgraded SPHERE (denoted by VLT+C in Table 7) may be able to detect this signal in under an hour. Since CMEs and fast stellar wind streams can have timescales ∼10 hours, and substorms up to several days (Gonzalez et al., 1994, 1999), the high level of transient activity may be observable. Under near constant CME activity, storm conditions could potentially last for weeks or longer, improving the chances of detecting auroral emission (Khodachenko et al., 2007).

Even future observations with TMT, HabEx, and LUVOIR outfitted with instruments optimized for high-resolution, high-contrast coronagraphy will be unable to detect a steady-state 0.1 TW OI aurora on Proxima Cen b. However, a coronagraph-equipped TMT could detect a substorm strength aurora of ∼10 TW in about 10 hours, while LUVOIR could make the predicted substorm auroral observation in about 30 hours. Only auroral powers $\gtrsim$10 TW would be detectable

with HabEx. Auroral powers of order 100 TW arising from a concurrent CME and substorm could be observed by the TMT, HabEx, and LUVOIR in well under an hour.

Finally, I consider how improvements in ground-based instrumentation might expand the ability to detect exo-aurorae. In the top panel of Fig. 82, I model a coronagraph-equipped TMT concept that achieves a design contrast of $C = 10^{-7}$ and has negligible instrumental noise (e.g., no dark current and read noise). Low-resolution observations with a resolving power smaller than the equivalent width resolving power, $R < \lambda/W_\lambda$, yield longer integration times at fixed auroral power as the auroral signal is diluted by additional stellar continuum photons from larger spectral elements, which increases the noise. For high resolutions that exceed the equivalent width resolving power, $R > \lambda/W_\lambda$, the auroral signal dominates the planetary continuum as the spectral element more tightly bounds the narrow emission feature, yielding little additional improvement in integration times.

In the bottom panel of Fig. 82, I vary coronagraph design contrasts for observations with and without instrumental noise. I find that a TMT with coronagraphic starlight suppression, negligible instrumental noise, a design contrast of $C = 10^{-7}$ and $R > 10^5$ allows for a detection of the predicted steady-state OI auroral emission (auroral power of $\sim$0.1 TW) over about 40 hours (see also Table 7). The discontinuities that occur at high resolving powers are due to the need for additional spectral elements to span the width of the OI auroral line.

Despite the likely increased strength of aurorae on Proxima Cen b compared to Earth, observing a *steady-state* 0.1 TW aurora requires sufficiently long integration times that it is not currently feasible, nor will it be feasible with the next generation of instruments, unless ideal instrumental performance were achieved. OI auroral detection may only be possible if observations coincide with magnetospheric substorms or periods of vigorous stellar activity, such as CMEs, which can induce much stronger aurorae ranging from $1 - 10$ TW (and up to $\sim$100 TW if Proxima Cen b has a stronger magnetic dipole than Earth). These transient events are frequent on Proxima Centauri (Davenport et al., 2016) and may persist on timescales comparable to the integration times needed to detect strong aurorae.

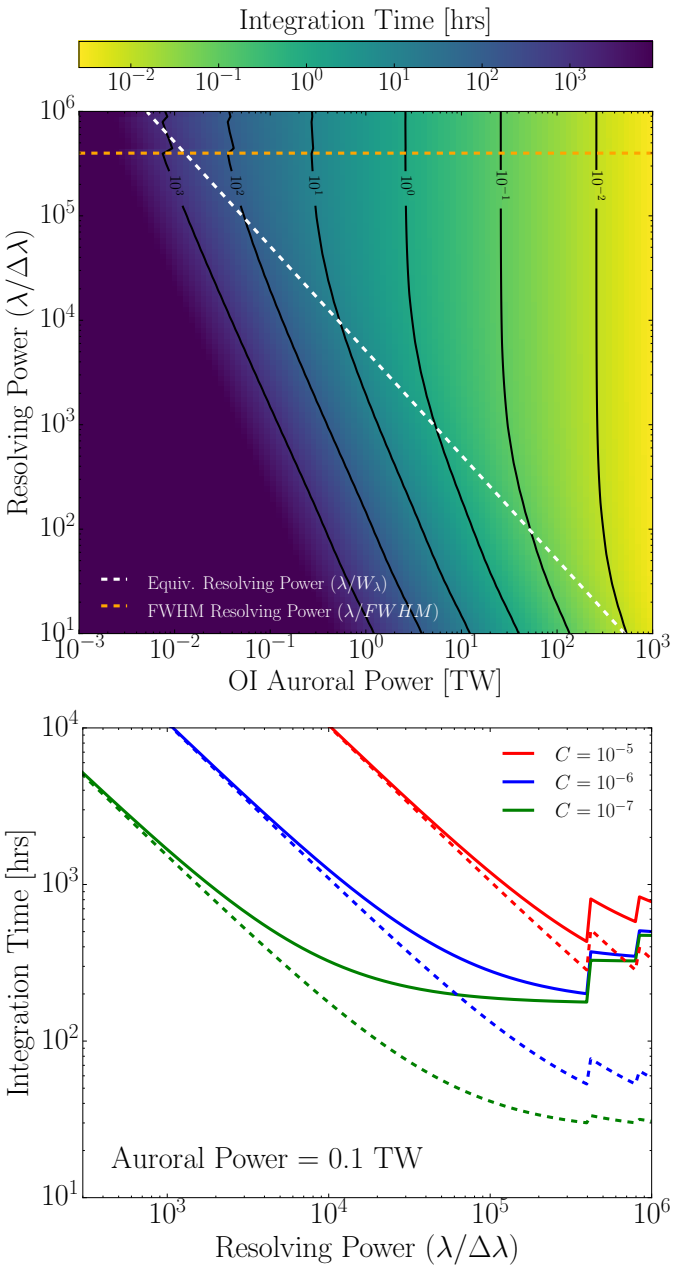

Figure 82: *Top*: Similar to Fig. 81, but displays telescope integration time contours as a function of telescope resolving power and OI auroral power for a coronagraph-equipped TMT concept with a design contrast of $C = 10^{-7}$. Dark current, read noise, and telescope thermal noise are set to zero here to simulate optimal detector performance that may be achieved by future instruments. *Bottom*: Telescope integration time as a function of resolving power for a coronagraph-equipped TMT concept for three different design contrasts. The solid curves denote integration times that include all modeled noise sources while the dashed curves assume negligible instrumental noise.

9.4 Search in the HARPS Data

The ESO Archive[20] hosts 319 HARPS spectra of Proxima Centauri taken between 2004 and 2016 and totaling about 70 hours of exposure time. The spectra were taken in the wavelength range $3782 - 6913$Å with a resolving power $R = 115,000$, yielding a wavelength resolution $\Delta\lambda \approx 0.05$Å at 5577Å. Each wavelength bin was oversampled by a factor of about 5. Given the estimates in Table 7, if Proxima Cen b's auroral power were on the order of 10^3 TW (however unlikely), the OI line could be detectable in this dataset. I therefore downloaded all spectra to conduct a search for the OI emission feature of Proxima Cen b. The method I outline below is similar to so-called "spectral deconvolution" techniques used to detect molecular absportion in exoplanet atmospheres (e.g., Sparks and Ford, 2002; Riaud and Schneider, 2007; Kawahara et al., 2014; Snellen et al., 2015).

I first shifted all spectra to the stellar rest frame by cross-correlating them against each other and calibrating the wavelength array to the stellar Na D I and II lines. Next, I removed stellar lines by performing weighted principal component analysis (WPCA; Delchambre, 2015) on a 250Å window centered at 5577Å. Each spectrum was then fit with a linear combination of the first 10 principal components, a number which I obtained by optimizing the recovery efficiency of injected planetary signals (see below); the fit was then subtracted, reducing the noise in the vicinity of 5577Å by a factor of $\sim$7. In order to obtain the principal components, I weighted each spectrum by the square root of its exposure time and assigned weights of zero to the individual telluric 5577Å airglow features, as these are among the strongest features in any individual spectrum and may incorrectly bias the principal components in the stellar frame. I remove Earth airglow separately below.

Next, I Doppler-shifted all spectra into the frame of Proxima Cen b. Since the orbital inclination i is unconstrained, I performed a grid search, varying i in one degree increments from 30° to 90°. I did not consider inclinations lower than 30° due to the difficulty of deconvolving stellar and planetary signals in near face-on orbits. I further varied the planet period P and planet mean longitude λ across a range spanning $\pm 3\sigma$ about the best fit values reported in Table 1 of Anglada-Escudé et al. (2016), in increments of 0.25σ; in total, I considered 37,440 different orbital configurations for the planet. For simplicity, the eccentricity was assumed to be zero, the

[20] `http://archive.eso.org/`

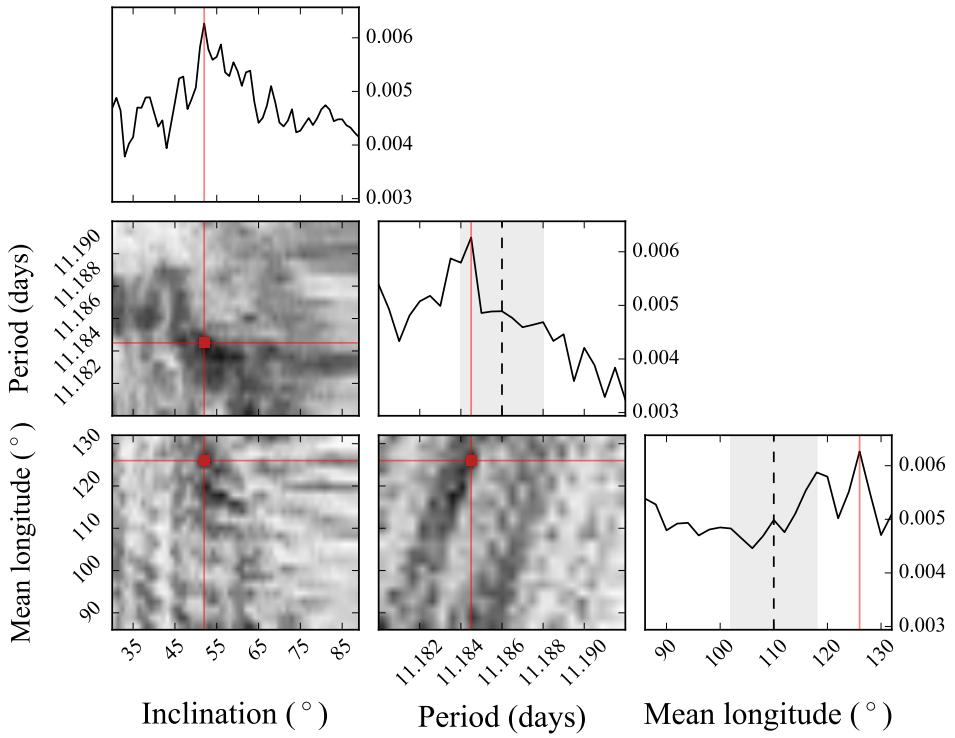

Figure 83: Results from the grid search over inclination (i), period (P), and mean longitude (λ) for the strongest 5577Å planetary signal. The inclination grid spans the range $30° - 90°$ in increments of $1°$. The period and mean longitude grids are centered on the best-fit values reported in Anglada-Escudé et al. (2016) and span the $\pm 3\sigma$ range in increments of 0.25σ. In total, 37,440 different orbital configurations for Proxima Cen b were considered. The curves along the main diagonal show the fractional amplitude of the bin centered on the OI line as a function of inclination (top left), period (center), and mean longitude (bottom right). In the period and mean longitude plots, the dashed line is the value reported in the discovery paper, with the 1σ bounds shaded in gray. The colormaps show the joint distributions of signal strengths for pairs of the three orbital parameters (black highest, white lowest). The peak signal is indicated by the red lines and occurs at $i = 52°$, $P = 11.1845$ days, and $\lambda = 126°$, with detection significance $\sim 0.7\sigma$. As I argue below, this signal has a very high false alarm probability (FAP ~ 0.2) and is entirely consistent with noise.

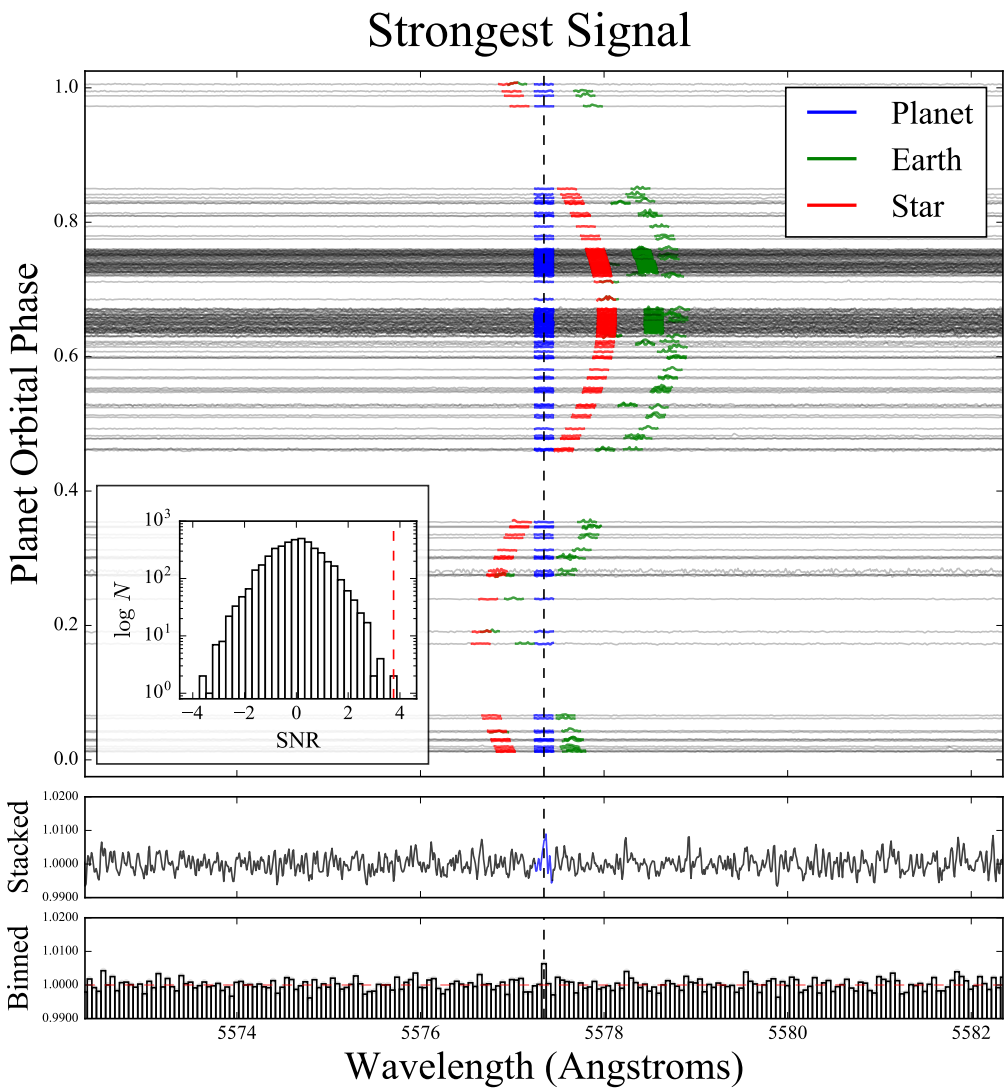

Figure 84: The HARPS spectra of Proxima Centauri. After removing stellar and telluric lines, the individual spectra are Doppler-shifted into the frame of Proxima Cen b according to the orbital parameters corresponding to the peak signal in Fig. 83. The spectra are then normalized and distributed vertically on the main subplot according to the planet's orbital phase. Blue regions indicate a small (0.2Å) window centered on the 5577Å oxygen feature in the planet frame. Red and green regions indicate the same window in the star and Earth frames, respectively; note the residual telluric airglow features in many of the spectra. The bottom subplots show the stacked spectrum in the planet frame and the stacked spectrum after downsampling to bins of size equal to the instrumental FWHM of the line (0.05Å). The peak recovered by the grid search is evident in both the stacked and the binned flux. The inset at the center left shows a histogram of the amplitude of deviations from the median in bins across a 250Å window centered on the OI line, indicating a signal-to-noise ratio (SNR) of about 4 in the 5577Å bin. Despite the apparent strength of this detection, further analysis yields a detection significance of only $\sim 0.7\sigma$, with false alarm probability $\sim$20% (see Fig. 86).

planet mass was set to $1.27 M_\oplus / \sin i$, and the stellar mass was held fixed at $0.12 M_\odot$. The latter parameter is considerably uncertain; however, changing the stellar mass changes the amplitude of the planetary RV signal, making the stellar mass degenerate with the inclination of Proxima Cen b's orbit. A grid search over the stellar mass is therefore redundant as long as we treat the inclination above as an "effective" inclination for $M_\star = 0.12 M_\odot$.

After Doppler-shifting each spectrum, I translated them back to the original wavelength grid by linear interpolation. Once in the planet frame, I identified and interpolated over $> 10\sigma$ outliers in each wavelength bin of the normalized spectra outside the 0.2Å window centered on the OI line. I found that this successfully removed telluric airglow and prevented outlier features in individual spectra from contributing to the stacked spectrum. I purposefully did not perform this outlier removal step in the vicinity of the (putative) planetary 5577Å line to prevent time-variable emission from being removed. Note that, in principle, this could result in a false detection of a planetary signal due to the presence of a large (non-planetary) outlier in a single spectrum. In the event that a signal were recovered, a detailed analysis of the spectrum/spectra it originated from would be necessary to rule out this possibility.

For each orbital configuration, I then co-added all spectra in the planet frame, omitting spectra in which the planetary 5577Å window overlapped with either the stellar or telluric 5577Å windows to avoid contamination from OI emission by those sources. For orbits close to edge-on, this reduced the total exposure time from 70 to about 50 hours, and less for lower inclination orbits. In order to remove correlated stellar noise, I then applied a high pass median filter of window size 1Å.

Finally, I binned the stacked spectra to 0.05Å-wide bins, with the central bin centered at 5577.345Å, the empirical wavelength of the OI green line (Cabannes and Dufay, 1955; Chamberlain, 1961). The bin size is the HARPS resolution at that wavelength, and closely matches the FWHM of the telluric OI lines in the dataset. As I argued in §9.3, a higher resolution spectrograph (with less instrumental broadening) would allow for smaller bin sizes and higher contrast in the OI line. I then measured the amplitude of the 5577.345Å bin relative to the spectrum mean.

The results of my grid search are shown in the triangle plot in Fig. 83. Along the main diagonal, I plot the maximum fractional strength of the 5577Å signal as a function of each of the orbital parameters (the inclination i, the period P, and the mean longitude λ). Below those plots, I show the two-parameter joint distributions of the maximum signal strength, where darker colors

correspond to higher values. A peak is visible at an (effective) inclination of 52°, a period of 11.1845 days, and a mean longitude of 126°. In Fig. 84 I show the spectra Doppler-shifted into the planet frame according to these orbital parameters. Each of the processed spectra are normalized and distributed vertically along the main subplot according to the planetary phase at the time the observation was made. The location of the expected OI planetary feature is indicated in blue; I show the same window in the frame of Proxima Centauri (red) and Earth (green), where residual telluric emission is clearly visible. As mentioned above, spectra in which Proxima Cen b's 5577Å window overlaps with either the stellar or telluric windows are omitted. When stacking the spectra below, I also I masked and interpolated over 0.2Å windows centered on the telluric features.

The stacked flux in the planetary frame is shown below the main plot, where the peak at 5577Å is visible. Below it, I show the stacked flux binned to 0.05Å bins; the feature also stands out here. The inset at the bottom left of the main subplot shows a histogram of the SNR of all the bins in a 250Å window centered on the OI line; the 5577Å feature (indicated by a dashed red line) is one of only two with SNR ∼4. However, this signal is consistent with correlated stellar noise and is in no way a detection of planetary 5577Å emission. To show this, I performed the same grid search used to generate Fig. 83 in each of 2,250 wavelength bins on either side of 5577Å (where we do not expect significant planetary emission) and computed the strongest recovered signal in each wavelength bin. The results are shown in Fig. 85, where the 5577Å bin is indicated by a red circle. The dashed red line indicates the fractional strength of the signal in that bin, ∼0.006, which roughly corresponds to a planet-star contrast of the same magnitude if the signal were real. However, it is clear from the figure that over 20% of the bins in the range $5465 - 5670$Å have a stronger peak signal than that in the 5577Å bin, from which I estimate a detection significance of $\sim 0.7\sigma$ and a false alarm probability (FAP) for my recovered peak of ∼0.2. These data are also shown as a histogram in Fig. 86, where the number of signals are plotted as a function of their fractional strength. It is evident from these figures that the peak shown in Fig. 84 is fully consistent with noise.

A striking feature of Fig. 85 is the correlated nature of the noise; for instance, no bins in the range $5500 - 5575$Å have particularly strong signals, while the 5577Å peak is one of many in its immediate vicinity. This is likely a sign of correlated stellar noise that was improperly removed by PCA, either due to high temporal stellar variability or nonlinear correlations across the spectrum

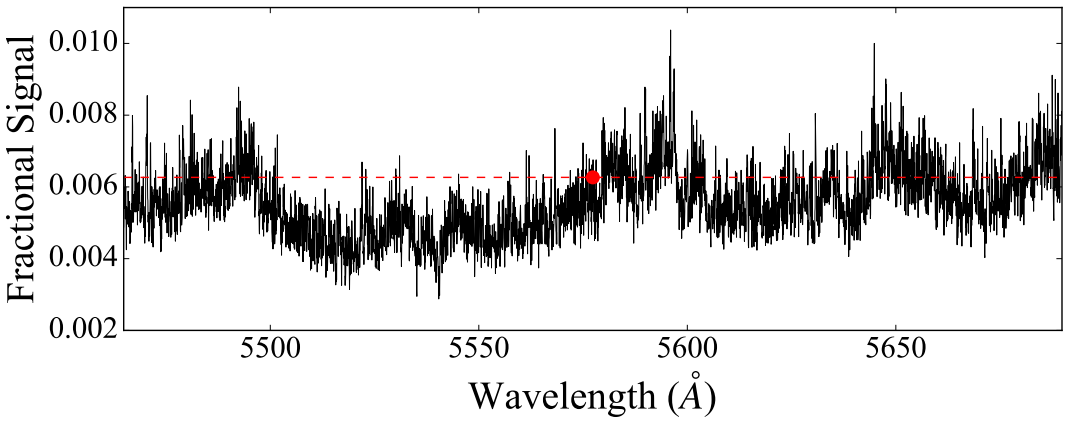

Figure 85: The peak signal in each wavelength bin in the vicinity of the 5577Å line. The fractional signal (y axis) is the flux in the bin divided by the continuum, and would roughly correspond to a planet-star contrast ratio if the signal were real. The peak signal at the 5577Å line (0.7σ) is indicated by the dashed red line. About 20% of the bins display stronger peak signals than the 5577Å bin, leading to a FAP for the 5577Å signal of ∼20%. Note also the strong correlated noise as a function of wavelength, likely due to improperly subtracted time-variable stellar features.

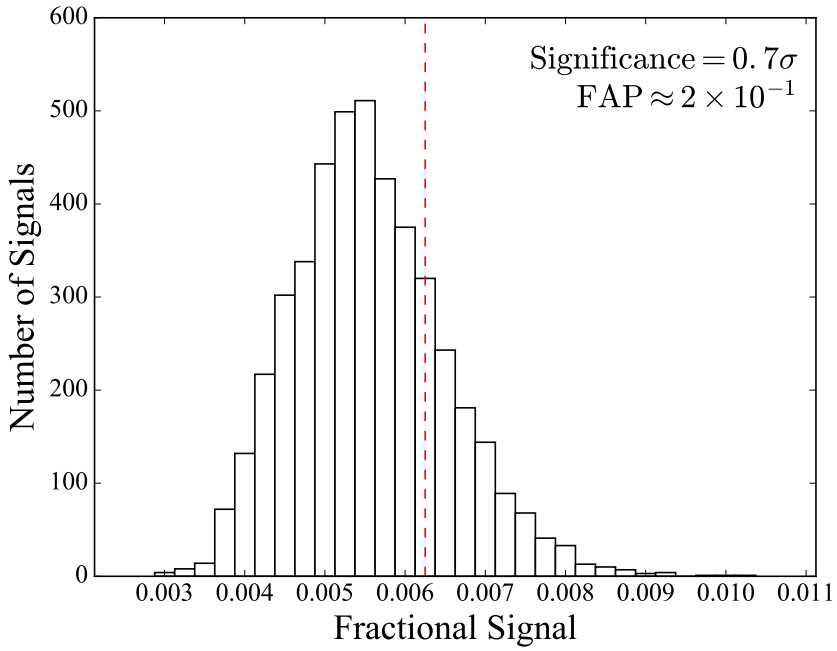

Figure 86: The distribution of the signal strength over the wavelength grid in Fig. 85. The bin corresponding to the peak 5577Å signal is indicated with a red dashed line; given the large FAP, the recovered signal is fully consistent with noise.

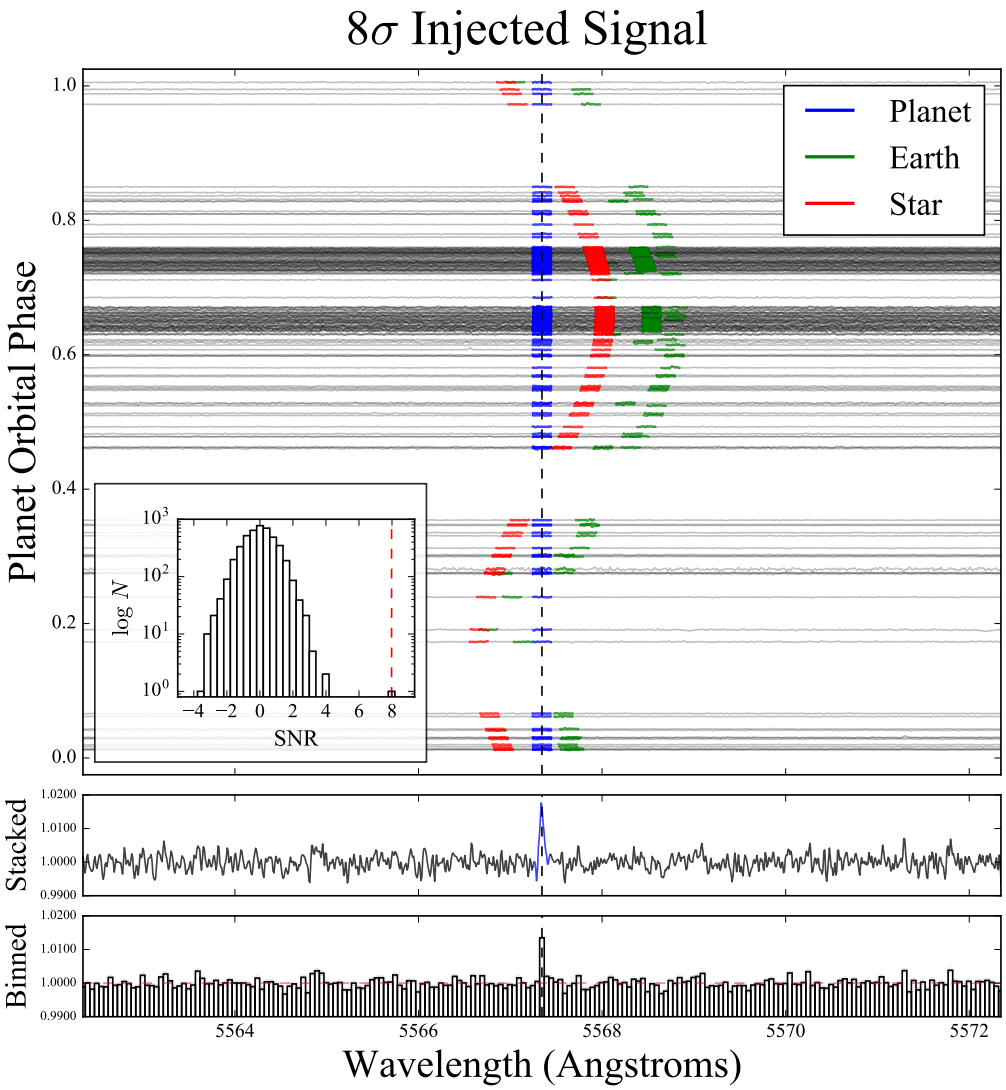

Figure 87: Similar to Fig. 84, but for an emission feature injected into the raw data at 5567.345Å (10Å blueward of the OI line, where no emission is expected) with contrast 1.8×10^{-2}, corresponding to a power of 2.6×10^3 TW. My method recovers the signal in the stacked, binned spectrum with SNR ∼8 and a detection significance of 8σ, my nominal detection threshold. The non-detection in Fig. 84 therefore constrains the auroral power on Proxima Cen b to be $\lesssim 3 \times 10^3$ TW, consistent with the calculations in §9.2.

that cannot be captured by PCA. At present, this noise is the limiting factor in my ability to recover auroral emission by Proxima Cen b. Future coronagraphic observations of Proxima Centauri (§9.3) should greatly reduce this noise by nulling most of the starlight.

In order to quantify the constraints my non-detection imposes on the properties of Proxima Cen b, I inject Gaussian OI emission signals with FWHM 0.05Å and of varying contrast into each of the raw spectra and attempt to recover them via the procedure described above. I find that an OI auroral signal with planet-star contrast 1.8×10^{-2} yields a detection with SNR $\sim$8 (Fig. 87). By performing the same wavelength search as before (Figs. 85 and 86), I estimate the significance of this detection to be 8σ, which I conservatively choose to be my nominal detection threshold for the HARPS dataset. From Table 7, this contrast is equivalent to an OI auroral power of $\sim 3 \times 10^3$ TW, my empirical upper limit on the strength of Proxima Cen b's 5577Å emission. Note that by scaling the estimates in that table, it should take only a few hours (versus the 50 hours used in the injection/recovery step) to detect such a signal with HARPS. This order-of-magnitude difference is likely related to the correlated stellar noise discussed above. Integration times for telescopes without coronagraphs may therefore be significantly longer than the estimates quoted in that table.

Finally, I also performed similar searches for the red oxygen lines (6300.308 and 6363.790Å) and the 3914.4Å UV nitrogen line, which are prominent in Earth's aurora, but find no significant peaks. Given the lower power in the red lines relative to the green line, and the low transmissivity of Earth's atmosphere and lower detector efficiencies in the UV, this non-detection is consistent with the non-detection of the 5577Å feature.

9.5 Discussion & Conclusions

My calculations above assume that Proxima Cen b is a terrestrial planet with an Earth-like atmosphere. Although its radius is not known — making an estimate of its density impossible — the planet is statistically likely to be rocky. The *a priori* probability distribution for the inclination of an exoplanet is $\mathrm{P}(i)\mathrm{d}i = \sin(i)\mathrm{d}i$. With 68% confidence, the inclination of Proxima Cen b is greater than 47°, and with 95% confidence it is greater than 18°. Given $m_p \sin i = 1.27\mathrm{M}_\oplus$, this corresponds to true planet masses smaller than $1.7\mathrm{M}_\oplus$ (68% confidence) and $4.1\mathrm{M}_\oplus$ (95% confidence). Recent exoplanet population studies suggest that the transition between rocky and gaseous exo-

planets occurs at a radius of about $\sim 1.6 R_\oplus$ (Rogers, 2015; Wolfgang and Lopez, 2015). However, a corresponding value for the transition *mass* is still uncertain, and since the radius of Proxima Cen b has not been measured, I cannot argue for its terrestrial nature based on its mass alone. Nevertheless, we can obtain predictions for its radius under certain assumptions. Assuming it is a rocky planet with Earth-like composition, we may use the scaling law from Fortney et al. (2007) to obtain radii of $1.16 R_\oplus$ ($m_p = 1.7 M_\oplus$) and $1.45 R_\oplus$ ($m_p = 4.1 M_\oplus$). Using the mass-radius grids of Lopez et al. (2012) and assuming instead that Proxima Cen b is a super-Earth/mini-Neptune with a thin H/He envelope with mass equal to 1% the planet mass, the radii jump to $1.7 R_\oplus$ ($m_p = 1.7 M_\oplus$) and $2.1 R_\oplus$ ($m_p = 4.1 M_\oplus$). Planet occurrence rate calculations for cool M dwarfs (Dressing and Charbonneau, 2013) suggest that there is a steep drop in the number of short-period planets per star with radii above $1.4 R_\oplus$ ($0.19^{+0.07}_{-0.05}$), compared to those with radii below $1.4 R_\oplus$, which are more than twice as common ($0.46^{+0.09}_{-0.06}$). This suggests that Proxima Cen b is more likely to be terrestrial than Neptune-like. A thin ($\lesssim 1\%$ by mass) H/He veneer is still possible, but given the extended pre-main sequence phase of the host star, past hydrodynamic escape is likely to have blown it off (Luger et al., 2015a; Barnes et al., 2016).

Nevertheless, I cannot definitively rule out the possibility that Proxima Cen b has an atmosphere dominated by H/He, in which case one would not expect OI auroral emission. A search for Lyman-Werner H_2 emission in the UV would be more appropriate in this case. Although broader than the lines I consider here, this emission is likely stronger, and is unlikely to be confused with stellar emission, given that it is molecular in origin. But perhaps more importantly, a robust *non*-detection of this and other H/He features could rule out a large gaseous envelope and confirm the terrestrial nature of the planet. That said, I am currently unable to efficiently probe near-face-on orbits due to the much smaller Doppler shift of the planetary lines. Observations made exclusively at quadrature, when the planet RV is highest, may help with this in the future.

Alternatively, Proxima Cen b could be terrestrial but be significantly larger than the Earth, with mass as high as $\sim 4 M_\oplus$ and radius $\sim 1.5 R_\oplus$. Since the scaling methods used in §9.2 implicitly assume Proxima Centauri b is similar to Earth in size, the auroral strength could be different than what I estimate. Assuming a global field, a larger planetary radius (and therefore core radius) could increase the magnetospheric cross-section to the stellar wind, leading to an increase in the emitted power. Furthermore, assuming an Earth-like atmospheric composition, the higher surface

gravity would decrease the ionospheric scale height, which could lead to larger magnetic field parallel potential drops in auroral acceleration regions. This would increase the upward flowing current and therefore the downward accelerated electron beams into the upper atmosphere, which could also enhance the auroral signal. Moreover, an increased atmospheric density at the depth where precipitating electrons deposit their energy could also change recombination rates and alter the energy level distribution of the O atoms, which would in turn affect the auroral strength in different lines. A quantitative estimate of these effects is beyond the scope of this study, as it would require both modeling the changes to the atmospheric structure and solving the Boltzmann kinetic transport equation.

Assuming Proxima Cen b is terrestrial, my HARPS search constrains its auroral power to be $< 3 \times 10^3$ TW. This is consistent with the calculations in §9.2, which suggest the OI auroral power on Proxima Cen b is likely $\sim$0.1 TW, or $\sim$100$\times$ that of the Earth during steady-state solar wind conditions. Those calculations, however, ignore transient increases in stellar magnetic activity, which can enhance the auroral signal and the diffuse airglow emission of the planet. As discussed in §9.2, transient magnetospheric activity could result in auroral power for the 5577Å line up to $10 - 100$ TW, lasting from $10 - 10^3$ minutes (§9.2.4). In addition, for a planet that is under near constant CME activity, it is possible that the storm conditions last for weeks or longer (Gonzalez et al., 1994, 1999). Spectra taken during periods of vigorous stellar activity could thus enhance the chances of detecting auroral emission.

Even if Proxima Cen b is terrestrial, an auroral signal is not guaranteed to be present. The existence of an atmosphere is still an open question, owing to vigorous past hydrodynamic escape (Luger and Barnes, 2015; Barnes et al., 2016), observed persistent stellar activity (Davenport et al., 2016) and an observationally unconstrained planetary magnetic field. If an atmosphere is in fact present, it may not be Earth-like; instead, it could be dominated by CO_2 (e.g. Meadows et al., 2016). Airglow and auroral 5577Å emission are still expected for such an atmosphere, since atomic oxygen is produced by photodissociation of CO_2; in fact, OI 5577Å emission has been observed at both Mars and Venus, both of which are CO_2-dominated (e.g. Bertaux et al., 2005; Slanger et al., 2001). In particular, Slanger et al. (2001) and Slanger et al. (2006) found that the Venusian airglow strength is comparable to that of Earth. Given that Venus receives about twice the solar flux Earth receives, airglow and/or auroral emission from a CO_2-rich Proxima Cen b could be a

factor of ∼2 weaker than the values I predict in this chapter, although detailed photochemical modeling is required to accurately model this scenario.

In the case that the atmosphere is oxygen-rich, nitrogen may need to be present to enhance the auroral signal. On Earth, the OI green line emission results primarily from O_2^+ dissociative recombination, as well as collisions with excited N_2 and direct electron impact (Strickland et al., 2000). It is unclear whether or not other molecular species could play a similar role if N_2 is not the bulk atmospheric constituent. However, the detection of the 3914Å N_2^+ band could be a good diagnostic in the UV, where the star is even fainter. If we assume that the power of the 3914Å nitrogen band is comparable to that of the OI line (which is typical for higher energy magnetospheric particle populations), then the planet-star contrast in the N_2^+ band would be an order of magnitude greater than at 5577Å. Since the strength of the N_2^+ band scales with magnetospheric parameters, stellar activity could cause strong transient features in the UV, which may be observable. Note, however, that limitations in UV detector efficiencies may complicate the detection of nitrogen and other UV aurorae.

My integration times for the predicted steady-state 5577Å OI auroral line render its detection infeasible for current facilities. However, if key design goals are met for future coronagraphs, steady-state aurorae may be more easily detected. As shown in Fig. 81, achieving the optimal star-planet contrast ratio at the emission line requires that the width of a spectral element (resolving power) is smaller (greater) than the line's equivalent width. For my predicted steady-state auroral emission (∼0.1 TW), this requires future spectrographs to achieve $R \gtrsim 10^5$. High-resolution spectroscopy is also needed to resolve the Doppler shift of the planetary auroral emission (∼1Å) and place strong constraints on the eccentricity and mass/inclination of the planet. Such constraints would lead to greater confidence in the terrestrial nature of Proxima Cen b. Furthermore, since read noise and dark current dominate the coronagraph instrumental noise budget, the development of low-noise detectors, e.g. MKIDS (Mazin et al., 2012, 2015), would significantly help the detection sensitivity and would allow such high-resolution spectroscopy to be downbinned to the lower resolution typically considered for direct exoplanet spectroscopy. For instance, if future detectors render read noise and dark current negligible, a LUVOIR concept telescope could observe a 0.1 TW OI 5577Å auroral emission feature in 100 hours as opposed to the 2×10^3 hours for the noised observations considered in §9.3. In addition, low-noise detectors would make stacking

short observations (required to mitigate broadening of the line due to the planet's orbital motion) more feasible. If TMT is built with a coronagraph that can achieve a design contrast of 10^{-7} and negligible instrumental noise, it could observe steady-state auroral emission ($\sim$0.1 TW) in a few nights. However, such an observation would require the development of an effective AO system in the optical.

Alternatively, observations made during periods of vigorous stellar activity may enhance the detectability of exo-aurorae on Proxima Cen b. Transient magnetospheric activity could increase auroral power to $1 - 100$ TW, depending on the planetary magenetic dipole strength, allowing future coronagraph-equipped TMT and LUVOIR telescopes to detect auroral emission in 1 hour or less. This is comparable to estimated CME timescales (Khodachenko et al., 2007) and much shorter than the timescales of long-lasting solar storm conditions (Gonzalez et al., 1994, 1999). Future observing missions similar to the MOST campaign (e.g. Davenport et al., 2016) could be used to characterize and monitor Proxima Centauri's activity levels to constrain the star's activity cycles. Such missions could aid in scheduling spectroscopic observations of Proxima Centauri. Observing the star following a CME-like event or long duration fast solar streams could enhance detectability of the planetary auroral signal.

The methods of exo-auroral detection discussed here are not limited to Proxima Cen b, but may be applicable to any exoplanet orbiting a nearby late-type star or brown dwarf. For example, the recently discovered TRAPPIST-1 system (Gillon et al., 2016) consists of three planets orbiting an active late M8 ultracool dwarf only 12 pc away; one planet in the system, TRAPPIST-1d, potentially lies in the habitable zone. Since TRAPPIST-1 is a later type star than Proxima Centauri, it is likely more active (e.g., West et al., 2008) and hence could generate larger particle fluxes and a stronger interplanetary magnetic field than Proxima Centauri, leading to more powerful aurorae on its planets. Additionally, the redder blackbody spectrum of TRAPPIST-1 makes it a factor of about 6 dimmer than Proxima Centauri at the OI 5577Å line, resulting in far more favorable contrast ratios. However, due to its distance, auroral emission from this system will be $\sim$100$\times$ dimmer than that from Proxima Centauri, likely making its detection infeasible. For coronagraphic observations, the distance to the TRAPPIST system would require an inner working angle smaller than the diffraction limit to extend as long as 5577Å for all known TRAPPIST-1 planets observed with a 10m class telescope, further complicating the observation.

Another planet to consider is GJ1132b, which orbits a M3.5 star 12 pc away (Berta-Thompson et al., 2015). Since it receives $\sim$19$\times$ the Earth's flux and may have an O_2 rich atmosphere (Schaefer et al., 2016), it could display strong auroral emission. However, as with the TRAPPIST-1 system, its distance makes auroral characterization difficult. Moreover, the earlier type host emits a larger fraction of its light in the optical, resulting in a poorer auroral contrast ratio.

Finally, exoplanets orbiting nearby brown dwarfs may be prime targets for exo-auroral searches. Early-type brown dwarfs display significant magnetic activity (West et al., 2008) and are significantly fainter than M dwarfs in the optical, potentially enhancing the detectability of the 5577Å signal from planets in orbit around them. Although no short-period exoplanets are currently known to orbit nearby brown dwarfs (He et al., 2017), the methods described in this chapter may be used as means of exoplanet detection, as suggested by Sparks and Ford (2002). Since the stack-and-search method described in §9.4 does not require previous RV observations of a system, a long baseline of spectroscopic observations of nearby M dwarfs and brown dwarfs could be used to search for Doppler-shifted 5577Å OI emission. My method is particularly sensitive to short-period terrestrial planets, whose auroral power (if an atmosphere is present) is large and whose large RV will Doppler-shift the signal by one or more Å. However, since the stack-and-search method is best suited to detect steady-state emission, exoauroral searches will likely have to wait for a future generation of space-based telescopes or noiseless ground-based ELTs capable of detecting sub-TW aurorae. These searches may someday reveal the presence of unknown nearby terrestrial exoplanets, including ones in the habitable zone.

All code used to generate the tables and figures in this chapter is open source and available at `https://github.com/rodluger/exoaurora`. A static version of the code is archived at `https://doi.org/10.5281/zenodo.192459`.

10 Planet-Planet Occultations

10.1 Introduction

Among the most remarkable accomplishments of the *Kepler* telescope was the discovery of a large population of tightly-packed, close-in, and highly coplanar multi-planet exoplanet systems. These include Kepler-32 (Swift et al., 2013), Kepler-444 (Campante et al., 2015), and Kepler-80 (MacDonald et al., 2016), all of which host five transiting planets, although dozens of other such systems with three or more planets are known (Borucki et al., 2011; Lissauer et al., 2011).

These Kepler systems were bested by the discovery of the TRAPPIST-1 system (Gillon et al., 2016, 2017), which hosts *seven* transiting terrestrial-size planets within less than 0.1 AU of their star, an ultracool dwarf only 12 pc away. The system is so packed that a transit occurs, on average, about 6% of the time, placing the TRAPPIST-1 light curves among the most information-rich transiting exoplanet datasets. Transit timing variation (TTV) analyses of these light curves have so far placed strong constraints on the masses and eccentricities of these planets (Gillon et al., 2017) and dynamical studies of resonances in the system led to the precise prediction of the period of TRAPPIST-1h, the farthest-out planet (Luger et al., 2017). A spectroscopic analysis of the transit light curves of TRAPPIST-1b and c with the Hubble Space Telescope (HST) has further ruled out cloud-free hydrogen-dominated atmospheres for those planets (de Wit et al., 2016). In the near future, the James Webb Space Telescope (JWST) is expected to detect secondary eclipses of the shortest-period TRAPPIST-1 planets and to reveal more detailed information about their atmospheres via transit transmission and secondary eclipse spectroscopy (Barstow and Irwin, 2016; Lustig-Yaeger et al., in preparation).

As new generations of telescopes continue to enable the detection of smaller signals, it is worthwhile to explore new methods to study exoplanets. In this work, I consider the detectability of planet-planet occultations (PPOs; Ragozzine and Holman, 2010) in the TRAPPIST-1 system and in other systems of multiple, close-in transiting planets. PPOs occur when one planet transits the disk of another planet in the same system as seen by a distant observer, producing a dip in the light curve due to the interception of light emitted from or reflected off of the occulted planet. While PPOs are in this sense analogous to transits, their signals are orders of magnitude weaker, given the large star/planet surface brightness ratio. PPOs are also in general quite rare, given the

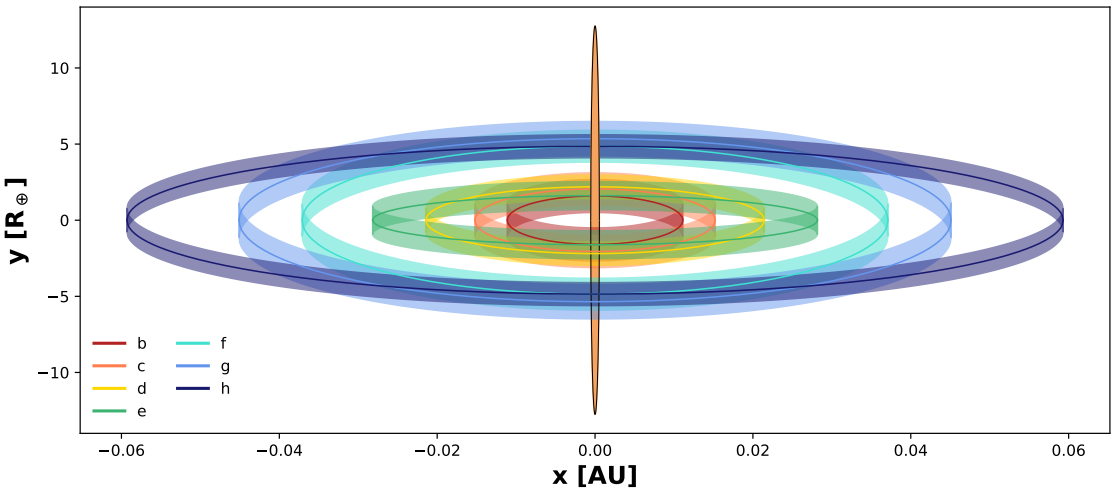

Figure 88: The orbits of the seven planets in the TRAPPIST-1 system as seen from Earth, assuming the mean orbital parameters from Gillon et al. (2017) and Luger et al. (2017). The thickness of each orbital track is the planet diameter. The aspect ratio of the plot is 100:1, but all horizontal and vertical distances are to scale. The star is shown in orange for reference. Because of the compactness of the system, its near edge-on orientation, and the dynamical coldness of the disk, the orbital paths of all planets overlap with those of their neighbors over a significant fraction of their orbits. For this particular configuration, planet-planet occultations occur between the set of planets $\{b, c, d, e\}$ and the set $\{f, g, h\}$.

small planet sizes (relative to the star) and the low probability of the overlap of their disks on the sky plane. For this reason they are also intermittent and typically short-lived.

However, the geometry of the orbits in TRAPPIST-1, and in several of the known multi-planet systems, is particularly favorable for planet-planet occultations. Because of the extreme coplanarity of the system, the orbital tracks of the seven TRAPPIST-1 planets overlap over considerable portions of their orbits as seen from Earth. In Figure 88 I plot these tracks for the mean orbital parameters reported in Gillon et al. (2017) and Luger et al. (2017). To make the orbital paths apparent, this figure is plotted with an aspect ratio of 100:1, but horizontal and vertical distances are each to scale. The star is shown in orange at the center and appears as a thin ellipse due to the vertical stretch. The widths of the orbital tracks are the planet diameters.

It is evident from the figure that the orbital paths of all seven TRAPPIST-1 planets overlap with those of their neighbors over large fractions of their orbits. Even when the observational uncertainties in the planet inclinations, radii, semi-major axes, and eccentricities are accounted for, the orbital tracks of all planets and those of their neighbors overlap for virtually any configuration

of the system allowed by the orbital constraints. The figure assumes the planets have the same longitude of ascending node Ω, which for an edge-on system is the polar angle of the orbit in the x-y plane. While large scatter in the values of Ω for the TRAPPIST-1 planets could disrupt the alignment of their orbits, I show that such configurations can be ruled out at high confidence with minimal assumptions (see §10.4.1.1).

The crossing orbital tracks suggest PPOs may be common among the TRAPPIST-1 planets and among planets in other compact, coplanar multi-planet systems. In particular, orbital crossings on the same side of the star can lead to long-lived PPOs, given the low sky-projected relative velocity of prograde, neighboring planets—in many cases, PPOs among these systems last significantly longer than transit or secondary eclipse. Furthermore, because the signal of a PPO event is the missing flux from a planet as it is occulted, complete planet-planet occultations can be as deep as secondary eclipses. For the TRAPPIST-1 system, PPOs are most detectable in emission at long wavelengths, where the planet/star contrast is most favorable due to the linear relationship between flux and temperature in the Rayleigh-Jeans limit.

The detection of one or more PPOs in a multi-planet system pins down the relative orbital positions of pairs of planets to extremely high precision, placing strong constraints on the mutual inclinations of the planets as well as their eccentricities and the relative orientation of their orbits on the sky. Moreover, because PPOs occur in general off the face of the star, the timing of a PPO event is affected by the planets' orbital eccentricities (§10.3.1.1), but also probes a planet's TTV curve at an orbital phase inaccessible to either transit or secondary eclipse, potentially breaking eccentricity-eccentricity and mass-eccentricity degeneracies inherent to traditional TTV measurements (Lithwick et al., 2012; Deck and Agol, 2015). Similarly to secondary eclipses, PPOs can provide constraints on the albedo and/or temperature of an occulted planet, with the added benefit that they allow one to sample both the dayside *and* nightside of the occulted planet when the occultation occurs far from the disk of the star. Finally, at high time resolution, the shape of a PPO light curve can constrain the two-dimensional surface reflectance or emission map, potentially at a range of wavelengths. Because PPOs are aperiodic, different occultations of a given planet will occur at different phases and with different impact parameters, in principle allowing one to construct crude multi-wavelength surface maps of the entire planet surface.

In this work I assess the frequency and dynamical properties of PPOs in the TRAPPIST-

1 system and their detectability with current and future instruments, including the James Webb Space Telescope (JWST) and the Origins Space Telescope (OST). I develop a framework to predict and model PPOs with the new, open-source, photodynamical model `planetplanet`. In §10.2 I review other works that have previously considered PPOs or similar events, and in §10.3 I describe my methodology for predicting, modeling, and extracting orbital information and crude surface maps from PPOs. I present my results for TRAPPIST-1 in §10.4 and discuss my findings and their applicability to other systems in §10.5.

10.2 Planet-planet occultations in the literature

Planet-planet occultations have previously been considered in other works as a method to detect or characterize planets. Ragozzine and Holman (2010) first introduced the term, citing previous work by Cabrera and Schneider (2007), Sato and Asada (2009), and Sato and Asada (2010) on the detectability of mutual transits of binary planets and planets with large moons. Ragozzine and Holman (2010) extended the idea to occultations among planets on independent astrocentric orbits. The authors pointed out that these so-called "planet-planet occultations" could place strong constraints on the three-dimensional architecture of exoplanet systems and even permit surface mapping with future telescopes, citing JWST as a potential facility to detect PPOs in emitted light. More recently, Brakensiek and Ragozzine (2016) developed code to calculate geometric probabilities of transits in exoplanet systems, with applications to PPOs and mutual transits. Veras and Breedt (2017) studied eclipses in exoplanet systems, with applications to TRAPPIST-1 and other compact systems. The coplanarity of TRAPPIST-1 enhances the likelihood that planets in that system will eclipse one another when seen from an observer situated on the surface of one of the planets. In principle, the shadow cast on the eclipsed planet could be detected from Earth, but the signal would be orders of magnitude weaker than a PPO observed in thermal light (Ragozzine and Holman, 2010).

Other relevant studies include those of Kipping (2011) and Pál (2012), who developed algorithms to model light curves of planet–moon and mutual planet–planet transits, respectively. The latter case, in which two planets occult each other as they transit their host star, has been studied by several other authors (e.g., Hirano et al., 2012; Masuda et al., 2013; Masuda, 2014). Hirano et al. (2012) presented the first claimed detection of such an event, seen as a brightening in the

light curve of Kepler-89 during a simultaneous transit of Kepler-89b and Kepler-89d. During a mutual transit the planet-planet occultation occurs on the face of the star, resulting in a signal that is typically orders of magnitude stronger than that of a typical PPO; these events are therefore detectable in white light with photometers like *Kepler*.

Closer to home, occultations have been studied, predicted, and observed among solar system bodies. The first ever recorded occultation of two planets in the solar system was that of Jupiter by Mars in the year 1170, which was observed by the monk Gervase of Canterbury and by Chinese astronomers (Stubbs, 1879; Hilton et al., 1988). Albers (1979) computed the ephemerides of past and future planet-planet occultations in the solar system, reporting two in the 19th century and five in the 21st century (but none in the 20th century!), all of which involve one of Mercury and Venus and one of the superior planets. Occultations among moons in the solar system have also been studied. Recently, de Kleer et al. (2017) used the Large Binocular Telescope (LBT) to observe an occultation of Io by Europa, deriving high resolution maps of the volcano Loki Patera from interferometric imaging of the moon's thermal emission during the occultation.

However, despite the considerable number of studies on planet-planet occultations, a detailed assessment of the detectability of PPOs in exoplanet systems and a framework to predict, model, and extract information from these events is still absent from the literature. In the next sections I discuss my approach to modeling and detecting PPOs in TRAPPIST-1 and in other exoplanet systems.

10.3 Methods

In order to model planet-planet occultation light curves, I developed the open source[21] software package `planetplanet`, a photodynamical code written in `C` and wrapped in `Python`. Given initial orbital, stellar, and planetary parameters, `planetplanet` integrates the system forward in time using an N-body code and computes light curves for all planet-planet occultations, as well as all transits, secondary eclipses, and planet phase curves. In §10.3.1 I discuss the dynamical theory of PPOs and its implementation in `planetplanet`, and in §10.3.2 I describe my approach to the photometric modelling of planet-planet occultations.

[21] `https://github.com/rodluger/planetplanet`

10.3.1 PPO Dynamics

Unlike transits or secondary eclipses, planet-planet occultations are episodic. Unless the ratio of the orbital periods of two planets is *exactly* the ratio of two integers, PPOs among the two planets will occur at different phases, with different impact parameters, and with varying durations. While all adjacent pairs of planets in TRAPPIST-1 are close to two-body resonances, departures from exact commensurability cause PPOs in this system to be aperiodic. Nonetheless, precise knowledge of the orbital parameters of pairs of planets can allow one to deterministically predict the times of future PPOs. By the same token, if the orbital parameters are not well known, the detection of PPOs can be used to impose strict constraints on their values. Below I describe techniques to constrain the eccentricity vectors and other three-dimensional orbital information via the detection of PPOs. In general, this is best done with an N-body code (§10.3.1.4), but it is instructive to first consider a few analytic methods.

10.3.1.1 Constraining the eccentricities

In the limit of two edge-on planets on plane-parallel orbits, and weak planet-planet perturbations, the timing of each PPO may be used to constrain a combination of the free eccentricity vectors of both planets involved. For a pair of planets, labelled i, j, with ephemerides $(t_{0,i}, P_i)$ and $(t_{0,j}, P_j)$ with $t_{0,i}$ the time of transit and P_i the orbital period, sorted by $P_i < P_j$, one may compute the longitudes at which the times of planet-planet occultations would occur if they were on *circular*, coplanar, and edge-on orbits:

$$\alpha \cos(\lambda_{k,i}) = \cos(\lambda_{k,j}) \tag{169}$$

$$\lambda_{k,i} = \frac{2\pi(t_k^0 - t_{0,i})}{P_i} - \frac{\pi}{2}, \tag{170}$$

where $\alpha = (P_i/P_j)^{2/3}$ and t_k^0 is the midpoint of the kth PPO between planets i and j, assuming zero eccentricity, and $\lambda_{k,j}$ is the mean longitude at time t_k^0 for planet j, measured from the sky plane. This is a transcendental equation that may be solved numerically for the roots t_k^0. The superscript zero indicates that this assumes zero eccentricity and edge-on orbits for each planet. These times *only* depend upon the ephemerides, which are measured precisely from the transit times (although they may have small uncertainties associated with each planet's TTVs). Note

that I have also neglected the masses of the planets in this equation.

One can compare this to the actual time of the kth PPO, t_k, for the more general case of non-zero eccentricity, albeit still ignoring inclination and TTVs. The difference between the eccentric and circular PPO times, $\delta t_k = t_k - t_k^0$, may be expanded to first order in eccentricity; assuming the eccentricities are small, as they are for packed planetary systems due to tidal damping and stability considerations, truncating second order eccentricity terms is well justified. Figure 89 illustrates a method to calculate the timing offset analytically as a function of the eccentricity vectors of the planets. Using the epicyclic approximation, one can compute the offset of the position of the planet relative to the guiding center, as well as the offset due to the fact that the transit time is shifted by the epicyclic motion of the planet. The offset between the x-position of the planets in the eccentric and circular case, Δx, divided by the relative x-velocity, Δv, gives the timing offset, $\delta t_k = \Delta x / \Delta v$.

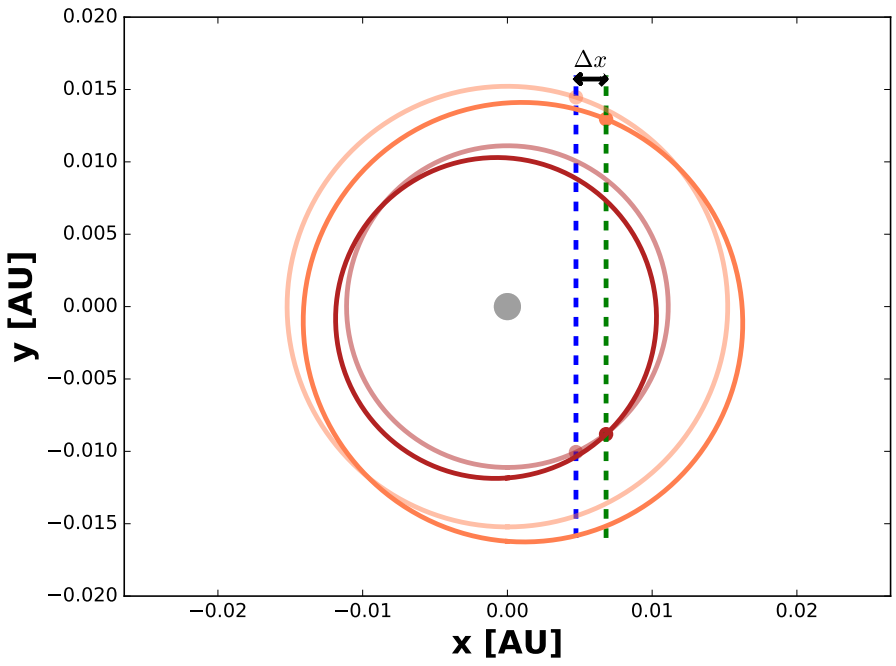

Figure 89: Pole-on diagram of coplanar orbits for planets b and c. The light colored circles (red and orange) are the orbits assuming zero eccentricity, while the darker ellipses are for $e_b = e_c = 0.1$. The position of the planets at the PPO in the circular case (blue dashed line) is offset from the eccentric case (green dashed line) by Δx. The observer is located at $(0, -\infty)$.

To first order in eccentricity,

$$\delta t_k = \tau_k \left(\mathbf{v}_j \cdot \mathbf{e}_j - \alpha \mathbf{v}_i \cdot \mathbf{e}_i \right), \tag{171}$$

where

$$\tau_k = \frac{P_j}{4\pi} \left(\sin \lambda_j - \alpha^{-1/2} \sin \lambda_i \right)^{-1}, \tag{172}$$

$$\mathbf{v}_i = \left[\cos\left(2\lambda_i\right) - 4 \sin \lambda_i - 3, \sin\left(2\lambda_i\right) \right], \tag{173}$$

$$\mathbf{e}_i = \left[e_i \cos \omega_i, e_i \sin \omega_i \right]. \tag{174}$$

This equation breaks down when $\sin \lambda_j \approx \alpha^{-1/2} \sin \lambda_i$ as the denominator approaches zero; in this limit the sky-projected acceleration can be included to obtain an accurate expression for δt_k, but the solution for the eccentricity vector in terms of the time offset is no longer linear. The mean longitude, λ_j, assumes zero eccentricity, and is zero when the planet crosses the sky plane in the direction away from the observer. Hence, transits occur when $\lambda_j \approx -\pi/2$ and secondary eclipses when $\lambda_j \approx \pi/2$. The cosine component of these vectors can be very similar in behavior due to the relation that causes PPOs to occur, Equation (169); however, these components can be constrained by the timing of secondary eclipses, which should be as detectable as (or more than) the PPO events. Note that since the eccentric orbits have a different breadth than the circular orbits, there may be occasional events that occur in the eccentric case that do not occur in the circular case (or vice versa); these will place a strict limit on the eccentricity vectors, but I ignore these for this analysis.

Equation (171) depends on the four components of the eccentricity vectors of both planets, and so at least four planet-planet occultations are needed to obtain a unique solution for the eccentricities of the planets. Given a set of measured PPO times, $\delta t_{k,obs}$, with uncertainties σ_k, then the eccentricities may be fit for directly by linear regression. Let $\mathbf{x} = [\mathbf{e}_i, \mathbf{e}_j]$ be the set of four eccentricity-vector parameters for the two planets. Letting $\delta t_k = \mathbf{y}_k \cdot \mathbf{x}$, where $\mathbf{y}_k = \tau_k \left[-\alpha \mathbf{v}_i, \mathbf{v}_j \right]$, one may write the goodness of fit metric as

$$\chi^2 = \sum_{k=1}^{K} \frac{\left(\delta t_{k,obs} - \mathbf{y}_k \cdot \mathbf{x} \right)^2}{\sigma_k^2} \tag{175}$$

for the observation of K planet-planet events between planets i and j. Since this equation is quadratic in $\mathbf{x}$, χ^2 has a unique minimum given by setting $d\chi^2/d\mathbf{x} = 0$. Taking the derivative with respect to x_m for $m = 1$ to 4, and setting these to zero, yields four equations for the four unknowns:

$$\mathbf{M}\mathbf{x} = \mathbf{b}, \tag{176}$$

where

$$M_{mn} = \sum_{k=1}^{K} \frac{y_{kn} y_{km}}{\sigma_k^2} \tag{177}$$

and

$$b_m = \sum_{k=1}^{K} \frac{y_{km} \delta t_{k,obs}}{\sigma_k^2}. \tag{178}$$

This may be inverted to solve for

$$\hat{\mathbf{x}} = \mathbf{M}^{-1}\mathbf{b}, \tag{179}$$

where $\mathbf{\Sigma} = \mathbf{M}^{-1}$ is the covariance matrix with Σ_{mm} the uncertainty on each eccentricity vector component.

An example of this technique is shown in Figure 90. Together with co-I Eric Agol, I simulated the timing offsets versus the time of planet-planet occultation events between TRAPPIST-1b and c for 30 days assuming edge-on Keplerian orbits, with 5-minute Gaussian random timing noise added to each PPO time. We then used these timing offsets to recover the eccentricity vectors of the planets with Equation (179). The input and recovered eccentricity vectors are listed in Table 8; our solutions agree within the uncertainties. The uncertainties on the eccentricity vectors will depend upon the period ratios of the planets, the timing precision achieved, and the sampling of the times; I provide code for carrying out these simulations to assist with experimental design in the `github` repository.

Note that in this simulation the timing offsets are larger than the amplitude of TTVs for TRAPPIST-1b and c, which are of order 1 minute, so that TTVs can be approximately neglected for these planets. However, if the eccentricities were much smaller for these planets, as they may

Table 8: Simulated recovery of the eccentricities of TRAPPIST-1b and c

	Input	Recovered
$e_b \cos \omega_b$	-0.0128	-0.0102 ± 0.0021
$e_b \sin \omega_b$	-0.0225	-0.0289 ± 0.0142
$e_c \cos \omega_c$	0.0067	0.0078 ± 0.0012
$e_c \sin \omega_c$	0.0053	0.0107 ± 0.0064

likely be due to tidal damping (Luger et al., 2017), then TTVs must be accounted for, which can be accomplished with a photodynamical code (see §10.3.1.3).

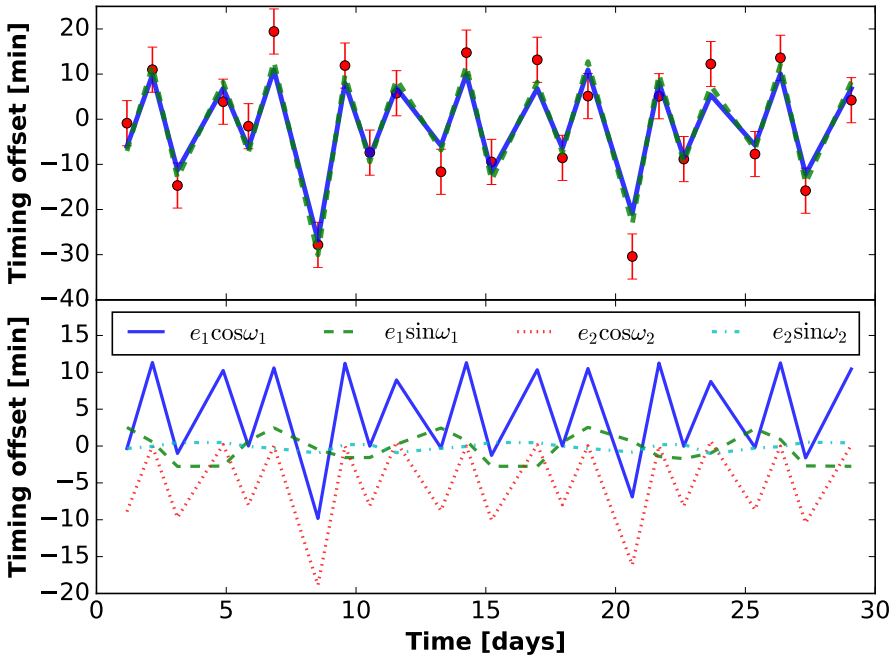

Figure 90: (Top) Simulated timing offsets for TRAPPIST-1b,c with an arbitrary eccentricity (red dots) and 5-minute Gaussian noise; the blue curve shows the noise-free computation. The optimized linear fit captures the timing offsets well (dashed green), and recovers the correct eccentricity vectors (Table 8). (Bottom) The amplitude due to each eccentricity vector component is shown; each of these has a different functional form, which is what allows the eccentricity vectors to be recovered.

Thanks to the linear dependence, the eccentricity uncertainties are proportional to the timing uncertainties. The timing uncertainties may be affected by the radiative intensity asymmetry of the planets, if accounted for inaccurately in the photometric analysis. As shown below, an airless body can mimic an occultation timing offset of several minutes; this gives an upper limit on the

expected timing error since the recirculation can be modeled, even if only approximately. The impact of intensity asymmetry on timing uncertainty will require further simulations to quantify more precisely.

Measuring the eccentricities can help to break the eccentricity-eccentricity degeneracy (Lithwick et al., 2012), and may also help break the mass-eccentricity degeneracy (Lithwick et al., 2012; Deck and Agol, 2015) that affects TTV measurements. Breaking these degeneracies could lead to more precise measurements of the mass ratios of the planets to the star, as well as better dynamical constraints and estimates of the tidal effects on the planets, such as constraints on the planetary Love numbers (Mardling, 2007; Batygin et al., 2009). The timing offset of secondary eclipses (Charbonneau et al., 2005) can impose a constraint on $e_i \cos \omega_i$, and thus fewer PPOs may be needed to obtain a unique solution for the planets' eccentricities; this could be incorporated into the linear analysis. I note, finally, that if PPOs occur between different pairs of planets, then the above solution may be extended to all planets involved to obtain a simultaneous solution for the eccentricity vectors of all N planets. However, this subsection ignores the mutual inclinations between planets and between system and observer; this will not be satisfied in general, and can provide additional constraints upon the orbital geometry, described next.

10.3.1.2 Constraining longitudes of nodes

The foregoing analysis assumed edge-on, plane-parallel orbits for the planets; in real planet systems, these assumptions will always be broken to some extent. The inclination of the planets' orbits (I) can be obtained from measurement of the transit/eclipse impact parameter(s) (Winn, 2010), while the longitudes of nodes (Ω) are more difficult to constrain, even with the measurement of TTVs (Nesvorný and Vokrouhlický, 2014), although it can be constrained in some cases in part due to transit duration variations (TDVs; Carter et al., 2012).

In contrast, planet-planet occultations can yield strong constraints on Ω. Assuming circular orbits and identical inclinations, the impact parameter will vary with the longitude of the planet-planet event as

$$\begin{aligned} b &= |a_i \cos \lambda_i \sin \Delta\Omega_{i,j}|/(r_i + r_j) \\ &= |a_j \cos \lambda_j \sin \Delta\Omega_{i,j}|/(r_i + r_j), \end{aligned} \tag{180}$$

where a_i and r_i are the semi-major axis and radius of the i^{th} planet, respectively. This will yield the absolute value of $\Delta\Omega_{i,j} = \Omega_i - \Omega_j$, the difference between the longitude of nodes of the two planets. The same calculation applies to planet-planet occultations that occur during transit (Pál, 2012), but in that case the lever arm is *much* smaller, and thus the constraint on $|\Delta\Omega|$ tends to be poor.

The inclination, I, combined with $|\Delta\Omega_{i,j}|$ will yield a constraint on the full geometry of the system, but with the discrete degeneracy between each pair of planets that undergo multiple occultations due to the absolute value. However, the measurement of $|\Delta\Omega|$ for more than two planets will yield a unique solution in most cases, thus constraining the full three-dimensional geometry of the planet system, up to an unknown sky angle.

10.3.1.3 De-aliasing transit timing

Planet-planet perturbations induce variations in the orbital elements at frequencies that are differences of integer multiples of the orbital frequencies of the planets, $P_{j,k} = (j/P_1 - k/P_2)^{-1}$, where j and k are integers, and P_1 and P_2 are the periods of a pair of adjacent planets (Deck and Agol, 2015). Transits are observed every orbital period of each planet, causing aliasing which induces a degeneracy between different k (j) values for the inner (outer) planet due to sampling on the period P_1 (P_2). This aliasing is the origin of the mass-eccentricity degeneracy (Lithwick et al., 2012; Deck and Agol, 2015). In contrast, planet-planet occultations will sample the mutual planetary perturbations at a *different* orbital phase, which may allow for the measurement of the free orbital eccentricity, as discussed above, but also may allow for the de-aliasing of different components of transit-timing variations. If sufficient precision can be obtained to measure the variation of the times of PPOs due to forced eccentricity and period variations caused by dynamical interactions, these timing measurements can help break the TTV mass-eccentricity degeneracy. In principle, the full orbital properties of the planets should be modeled to create a photodynamical model (Carter et al., 2012); I describe this approach next.

10.3.1.4 N-body code

In order to model the general case of eccentric, non-coplanar, massive planets subject to transit timing variations, I use the `REBOUND` N-body code (Rein and Liu, 2012) to model PPOs in

`planetplanet`. Starting from an initial state informed by the observational constraints of a multi-planet system, I integrate the orbital positions of all bodies forward in time using `REBOUND`. I implement both the high order integrator `IAS15` (Rein and Spiegel, 2015) and the symplectic integrator `WHFAST` (Rein and Tamayo, 2015). I track the relative sky-projected positions of all bodies, taking Keplerian steps on a finer subgrid to resolve all planet-planet events. At each subgrid timestep, I evaluate the impact parameters between all pairs $\{i, j\}$ of planets, given by

$$b_{ij} = \frac{\sqrt{(x_i - x_j)^2 + (y_i - y_j)^2}}{r_i + r_j}, \tag{181}$$

where x_i and y_i are the sky-projected Cartesian coordinates of the i^{th} planet and r_i is its radius. Occultations between two planets occur when $b_{ij} < 1$. I discuss the implementation of this procedure for the TRAPPIST-1 system in detail in §10.4.1.

10.3.2 PPO photometry

Given the full orbital solution for the system obtained with `REBOUND`, `planetplanet` computes light curves for all transits, secondary eclipses, and planet-planet occultations, as well as phase curves for all planets in the system. Since planet-planet occultations are observable primarily in the mid- and far-infrared, in what follows I model only the thermal emission from planets and neglect reflected light from the star, although the latter may easily be computed with `planetplanet` given the symmetry of the problem.

I developed a novel scheme for fast computation of occultation light curves optimized for the case of a body whose thermal emission map is radially symmetric about an *arbitrary* point on its surface. This is the case for a limb-darkened star or limb-darkened/ limb-brightened planet with a thick atmosphere and homogeneous cloud cover, whose surface brightness is to a very good approximation symmetric about the center of its disk. It is also the case for a planet (or moon) with a thin (or nonexistent) atmosphere, whose surface brightness is symmetric about the substellar point but in general lacks radial symmetry in the observer's frame. In the sections below, I describe the geometry of these two limiting cases for a planet's thermal emission map and discuss my integration scheme. At present, I only model planets in these two limits. The presence of spatially variable surface features and/or inhomogeneous clouds break these symmetries, and will be left for future investigation.

10.3.2.1 Planets with thick atmospheres

In the limit that a planet's atmosphere is thick, thermal emission from the planet is spatially decoupled from the illumination pattern of the star and the planet appears as a radially symmetric, limb-darkened (or limb-brightened) disk. Barring inhomogeneities due to clouds, this is approximately the case for Venus (Murray et al., 1963), and if one neglects star spots, it is also a good approximation for stars. The case of a spherical body transiting a limb-darkened disk is well studied (Mandel and Agol, 2002) and the occultation light curve is analytic under the linear, quadratic, and nonlinear limb darkening laws (Claret, 2000). However, the calculation involves multiple evaluations of elliptic integrals or hypergeometric functions, which are expensive to compute. Often it is more expedient to exploit the radial symmetry of the problem to reduce a two-dimensional integral to a one-dimensional integral and solve the latter numerically; this is the approach taken in the nonlinear limb-darkening law case in code provided by Mandel and Agol (2002). Recently, Kreidberg (2015) developed a similar method for the `batman` transit calculation package. The `batman` code discretizes the stellar surface as a series of spherical segments, each with a constant intensity given by an arbitrary radial limb darkening profile. Seen in projection, these are concentric rings. The flux occulted by a transiting planet is then simply the weighted sum over the areas of overlap between the planet and each of the rings, which are expressed in terms of arccosine functions (Mandel and Agol, 2002). These numerical methods allow for the fast computation of light curves under arbitrary limb darkening laws.

In order to treat occultations of limb-darkened/brightened planets (and stars), I adopt a similar integration scheme, discretizing the body's surface with N concentric circles that bound $N + 1$ spherical segments equally spaced in ϕ, the angle between the line of sight and the vector normal to the sphere. Each circular boundary is centered at the origin and has a radius

$$a = r_{\mathcal{P}} \sin \phi \tag{182}$$

where $r_{\mathcal{P}}$ is the radius of the body. This results in an adaptive grid in the radial coordinate, a, in which the radial spacing between rings decreases toward the limb, where the change in intensity

is fastest. If $\Delta\phi$ is the grid spacing in ϕ, the spacing in a is

$$\Delta a \approx \Delta\phi\sqrt{1-a^2}, \tag{183}$$

which in the limit $a \to 1$ is identical to that used in `batman`.

I adopt the following generic limb darkening law:

$$B_\lambda(\phi) = B_\lambda^0 \left[1 - \sum_{i=1}^{n} u_i(\lambda)(1-\cos\phi)^i\right], \tag{184}$$

where B_λ is a spectral radiance (measured in units of power per unit area per unit solid angle per unit wavelength), B_λ^0 is the spectral radiance at the center of the disk (see below), λ is the wavelength, and u_i is the i^{th} limb darkening coefficient, which I allow to be an arbitrary function of wavelength. The variation in the intensity of the disk towards the limb can be a strong function of wavelength, even for a perfect blackbody, and it is essential to allow for this when modeling occultations in the mid-infrared. Note that for $u_i(\lambda) =$ constant and $n = 1$ or $n = 2$, this is equivalent to the standard linear or quadratic limb darkening laws (Claret, 2000), respectively. Also note that this definition is equivalent to the limb-darkening formulation of Giménez (2006), although with differently defined coefficients.

One may find the normalization constant B_λ^0 by requiring that the blackbody intensity integrated over the planetary (or stellar) disk be equal to that of a blackbody at the effective temperature of the body, T_{eff}:

$$\int_0^{2\pi}\int_0^{r_\mathcal{P}} B_\lambda\left(\phi(r)\right) r dr = \pi r_\mathcal{P}^2 B_{\lambda,T_{\text{eff}}}, \tag{185}$$

where $r_\mathcal{P}$ is the radius of the planet and $B_{\lambda,T_{\text{eff}}}$ is the blackbody intensity at a wavelength λ and an effective temperature T_{eff} (Equation 192 with $T = T_{\text{eff}}$). Defining

$$\mu \equiv \cos\phi = \sqrt{1 - \left(\frac{r}{r_\mathcal{P}}\right)^2}, \tag{186}$$

one may write Equation (185) as

$$2\pi \int_0^1 B_\lambda(\mu)\mu d\mu = \pi B_{\lambda,T_{\rm eff}}. \tag{187}$$

Plugging in Equation (184) and solving for B_λ^0,

$$B_\lambda^0 = \frac{B_{\lambda,T_{\rm eff}}}{2\int_0^1 \left[1 - \sum_{i=1}^n u_i(\lambda)(1-\mu)^i\right]\mu d\mu}, \tag{188}$$

which may be simplified to yield

$$B_\lambda^0 = \frac{B_{\lambda,T_{\rm eff}}}{1 - 2\sum_{i=1}^n \frac{u_i(\lambda)}{(i+1)(i+2)}}. \tag{189}$$

In principle, my method is flexible enough to allow for any limb darkening law, including those that cannot be expressed as simple polynomial functions of ϕ. As both the temperature and the abundance of various species in a planetary atmosphere can change drastically with altitude, light emerging from the limb (where the effective emission layer is higher in the atmosphere) may have a drastically different spectrum than light emerging from the center of the planet disk. Because low-order polynomials are not flexible enough to accurately model this, my model also accepts as input a grid of emission spectra over the planet disk, provided the emission profile is radially symmetric:

$$B_\lambda(\phi) = B_{i,j}(\lambda, \phi), \tag{190}$$

where i and j are the corresponding indices in the λ and ϕ grids, respectively, such that $\lambda_{i-1} < \lambda \leq \lambda_i$ and $\phi_{j-1} < \phi \leq \phi_j$.

10.3.2.2 Eyeball planets

In the limit that a planet's atmosphere is thin enough as to have negligible thermal inertia and negligible heat transport, the temperature of a given region on the planet's surface is dictated entirely by radiative equilibrium with the incident stellar flux $F_\star$. If ϕ is the angle of the star measured from the zenith as seen from an observer at a point on the planet surface, the temperature

at that point is given by (e.g., Maurin et al., 2012)

$$T(\phi) = \max\left(T_{\rm night}, \left(\frac{F_\star(1-A)\cos\phi}{\sigma}\right)^{\frac{1}{4}}\right), \tag{191}$$

where A is the planetary albedo, σ is the Stefan-Boltzmann constant, and $T_{\rm night}$ is the temperature of the unilluminated night side, which may be nonzero due to, e.g., a geothermal heat flux. Note that for a planet at full phase, ϕ is identical to the angular measure used in §10.3.2.1. From here on out, I refer to ϕ as the "zenith angle," defined in the range $[0, \pi]$.

The corresponding spectral radiance, B_λ, is given by Planck's law,

$$B_\lambda(\lambda, \phi) = \frac{2hc^2}{\lambda^5}\frac{1}{e^{\frac{hc}{\lambda k_B T(\phi)}} - 1}, \tag{192}$$

where h is Planck's constant, c is the speed of light, k_B is Boltzmann's constant, and λ is the wavelength. While the finite angular size of the star as seen from its planets can result in illumination past the day/night terminator ($\phi = \frac{\pi}{2}$), the change in the radiance of the planet is negligible and I do not consider it here.

As before, let us discretize the radiance gradient. Regions of constant radiance are once again spherical segments. When the planet is seen at full phase, the curves bounding these regions are concentric circles centered on the sub-stellar point, as in the limb-darkened case. At half phase, these become lines perpendicular to the vector connecting the planet and the star. At intermediate viewing angles, these curves are segments of ellipses.

Figure 91 shows an example of a region of constant surface brightness extending from the sub-stellar point to a zenith angle ϕ. This region is a spherical cap, which in the observer's frame is bounded by an ellipse of semi-major and semi-minor axes

$$\begin{aligned} a &= r_\mathcal{P} \sin\phi \\ b &= r_\mathcal{P} \sin\phi |\sin\theta|, \end{aligned} \tag{193}$$

respectively, where $r_\mathcal{P}$ is the radius of the body and θ is the phase angle, the angle between the sub-stellar point and the axis perpendicular to the line of sight and parallel to the orbital plane (see Figure 91). For an orbit at arbitrary eccentricity viewed edge-on and aligned with the xz

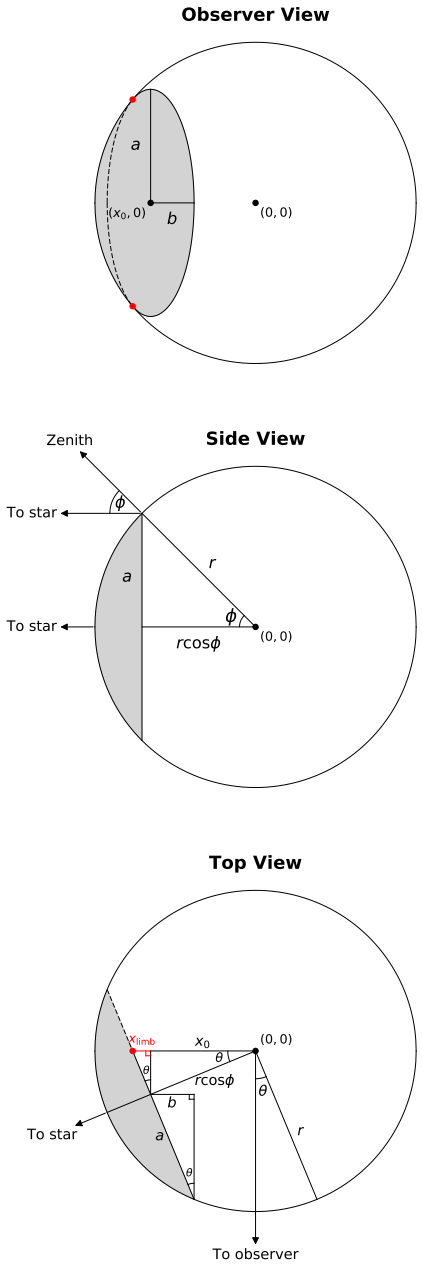

Figure 91: Geometry of a region of constant surface brightness on the planet surface, seen from three different vantage points. In general, such a region is a thin spherical segment of radius $a = r \sin\phi$, where r is the radius of the body and ϕ is the zenith angle. The boundaries of each region are circles, which when projected onto the sky plane become ellipses with semi-major axis a and semi-minor axis $b = a \sin\theta$, where θ is the phase angle. In this example, I show a spherical cap extending to $\phi = \frac{\pi}{4}$ for a planet at $\theta = \frac{\pi}{8}$. Note that a portion of the elliptical boundary is behind the limb of the planet (dashed lines). See text for details.

plane, with the x axis pointing to the right on the sky and the z axis pointing into the sky, this angle is simply the mean longitude of the orbit and is given by

$$\theta = \arctan2\left(z, x\right), \tag{194}$$

where z and x are the coordinates of the planet in this frame and arctan2 is the two-argument arctangent function. An airless planet in an orbit with arbitrary inclination will still appear as an "eyeball" (§10.3.2.2) in projection, but at a different phase angle than that given in Equation (194). The planet disk will also appear rotated on the sky plane through some angle γ. If the planet is at the point (x, y, z) in a left-handed Cartesian coordinate system with the x axis pointing to the right on the sky, the y axis pointing up, and the z axis pointing into the sky, it is straightforward to show that the sub-stellar point is located at

$$\begin{aligned} x_\star &= x\left(1 - \frac{r_\mathcal{P}}{r}\right) \\ y_\star &= y\left(1 - \frac{r_\mathcal{P}}{r}\right) \\ z_\star &= z\left(1 - \frac{r_\mathcal{P}}{r}\right), \end{aligned} \tag{195}$$

where

$$r = \sqrt{x^2 + y^2 + z^2} \tag{196}$$

is the magnitude of the orbital radius vector. The angle by which the hotspot is rotated away from the horizontal on the sky, measured counter-clockwise from the $-x$ axis, is then

$$\gamma = \pi + \arctan2\left(y_\star - y, x_\star - x\right). \tag{197}$$

The distance from the center of the planet disk to the hotspot is simply

$$d = \sqrt{(x_\star - x)^2 + (y_\star - y)^2}. \tag{198}$$

Setting $\phi = 0$ in Equation (200), the effective phase angle for this planet is then

$$\theta = \begin{cases} \arccos\left(\frac{d}{r_{\mathcal{P}}}\right) & z^\star - z \leq 0 \\ -\arccos\left(\frac{d}{r_{\mathcal{P}}}\right) & z^\star - z > 0. \end{cases} \tag{199}$$

Note that in the edge-on limit ($y = 0$), this reduces to Equation (194).

As expected, when $\theta = \pm\frac{\pi}{2}$ (full/new phase), $a = b$ and the bounding region is a circle; when $\theta = 0$ (half phase), $b = 0$ and the bounding region is a line. In general, assuming the planet is centered at the origin, it is straightforward to show that the ellipse is centered at

$$\begin{aligned} x_0 &= -r_{\mathcal{P}} \cos\phi \cos\theta \\ y_0 &= 0. \end{aligned} \tag{200}$$

Note that one must also account for the fact that the ellipse may not be fully visible to the observer. In the example shown in the figure, the points where the ellipse crosses beyond the limb of the planet are indicated as red dots. It can be shown that the x coordinate of these points is

$$\begin{aligned} x_{\text{limb}} &= x_0 - r_{\mathcal{P}} \cos\phi \sin\theta \tan\theta \\ &= x_0 \sec^2\theta. \end{aligned} \tag{201}$$

As I discuss below, this formalism allows us to approximate the occulted planet flux as a sum over the integrals of elliptical segments, which are analytic when one of the axes of the ellipse is parallel to the x axis.

10.3.2.3 Integration scheme

Given the discretized radiance gradients discussed above, I wish to compute the total flux occulted by a body that passes in between the emitting body and the observer. This may be computed as the integral of the occulted planet's radiance evaluated over the region of overlap of two circles. In both the limb-darkened and airless body cases, all curves of constant radiance are ellipses (or circles), so the problem is reduced to a sum over one dimensional integrals of ellipses, which are analytic. Below I show how this integration may be fully automated by identifying all relevant

points of intersection between the ellipses, sorting them, and evaluating the integrals in the region between all pairs of adjacent points.

Circle-circle intersection. I begin by deriving the points of intersection between the disks of the occulted planet (radius $r_\mathcal{P}$) and the occultor (radius $r_\mathcal{O}$). If the occulted planet is centered at the origin and the occultor at the point $(x_\mathcal{O}, y_\mathcal{O})$, the x coordinates of the points of intersection of their disks are

$$x = a \pm b \tag{202}$$

where

$$\begin{aligned} a &= \frac{x_\mathcal{O}}{2d^2}\left(d^2 - r_\mathcal{P}^2 + r_\mathcal{O}^2\right), \\ b &= \frac{y_\mathcal{O}}{2d^2}\sqrt{4r_\mathcal{P}^2 r_\mathcal{O}^2 - (r_\mathcal{P}^2 + r_\mathcal{O}^2 - d^2)^2} \end{aligned} \tag{203}$$

and

$$d = \sqrt{x_\mathcal{O}^2 + y_\mathcal{O}^2} \tag{204}$$

is the distance between their centers.

Circle-ellipse intersection. Next, I must find the points of intersection between the occultor and each of the curves of constant radiance on the occulted body, which as I showed above are ellipses. The points of intersection of a circle centered at the origin with radius r and an ellipse centered at (x_0, y_0) with semi-major axis a parallel to the y axis and semi-minor axis b parallel to the x axis are given by the roots of the following equation:

$$y_0 \pm \frac{a}{b}\sqrt{b^2 - (x - x_0)^2} \pm \sqrt{r^2 - x^2} = 0. \tag{205}$$

This is equivalent to finding the roots of the following quartic polynomial:

$$c_4 x^4 + c_3 x^3 + c_2 x^2 + c_1 x + c_0 = 0 \tag{206}$$

where

$$
\begin{aligned}
c_0 &= C^2 - D(b^2 - x_0^2) \\
c_1 &= 2BC - 2Dx_0 \\
c_2 &= 2AC + B^2 + D \\
c_3 &= 2AB \\
c_4 &= A^2
\end{aligned}
\tag{207}
$$

and

$$
\begin{aligned}
A &= \frac{a^2}{b^2} - 1 \\
B &= -2x_0 \frac{a^2}{b^2} \\
C &= r^2 - y_0^2 - a^2 + x_0^2 \frac{a^2}{b^2} \\
D &= 4y_0^2 \frac{a^2}{b^2}.
\end{aligned}
\tag{208}
$$

In general, this polynomial has four roots, 0, 2 or 4 of which are real and correspond to the actual points of intersection. These roots may be found analytically (e.g., Hughes and Chraibi, 2011), but in practice the large number of arithmetic operations can result in significant errors due to limited machine precision. It is therefore preferable to solve Equation (206) with a numerical root-finding algorithm (e.g., Press et al., 1992).

Since in my integration scheme (see below) I place the occulted body at the origin, I have $x_0 = -r_\mathcal{P} \cos\phi \cos\theta - x_\mathcal{O}$ and $y_0 = -y_\mathcal{O}$ (see Equation 200). In the frame of the occulted planet, the points of intersection are $x_n + x_\mathcal{O}$, where x_n are the roots of the quartic polynomial above.

Integral of an ellipse. The function describing an ellipse centered at (x_0, y_0) and with semi-major axis a aligned with the y axis and semi-minor axis b aligned with the x axis is

$$
y(x) = y_0 \pm \frac{a}{b} \sqrt{b^2 - (x - x_0)^2}. \tag{209}
$$

The indefinite integral of Equation (209) is

$$\int y(x)dx = y_0 x \pm \frac{a}{2b}\left[z(x-x_0) + b^2 \arctan\left(\frac{x-x_0}{z}\right)\right] + C \tag{210}$$

where

$$z = \sqrt{(b+x-x_0)(b-x+x_0)}$$

and C is an arbitrary constant.

Zenith angle of a point on the occulted planet. The final derivation before I introduce my integration scheme concerns the zenith angle of a point (x, y, z) on the surface of a sphere, which I will use to determine the radiance of a given region on the planet disk. Given a planet of radius $r_{\mathcal{P}}$ in an edge-on orbit at phase angle θ and centered at the origin, the zenith angle may be computed as follows. If $\theta = \pm\frac{\pi}{2}$, curves of constant zenith angle are circles, and the zenith angle is simply

$$\phi = \begin{cases} \arcsin\left(\frac{\sqrt{x^2+y^2}}{r_{\mathcal{P}}}\right) & \theta = +\frac{\pi}{2} \\ \pi - \arcsin\left(\frac{\sqrt{x^2+y^2}}{r_{\mathcal{P}}}\right) & \theta = -\frac{\pi}{2}. \end{cases} \tag{211}$$

At intermediate angles, the curves of constant zenith angle are ellipses. Given the equations in §10.3.2.2, I may solve for the zenith angle of the ellipse that threads the point (x, y). This yields a quadratic equation in the quantity $z = \sin^2\phi$, which may be solved to yield

$$\phi = \begin{cases} \arcsin\sqrt{z} & \left(|\theta| \le \frac{\pi}{2} \text{ and } x \le x_{\text{term}}\right) \text{ or } \left(|\theta| > \frac{\pi}{2} \text{ and } x \ge -x_{\text{term}}\right) \\ \pi - \arcsin\sqrt{z} & \text{otherwise} \end{cases} \tag{212}$$

where

$$z = \frac{1}{2}\left[1 + \frac{y^2}{r_{\mathcal{P}}^2} + \left(1 - \frac{2x^2}{r_{\mathcal{P}}^2} - \frac{y^2}{r_{\mathcal{P}}^2}\right)\cos(2\theta) + \frac{2x}{r_{\mathcal{P}}}\sqrt{1 - \frac{x^2}{r_{\mathcal{P}}^2} - \frac{y^2}{r_{\mathcal{P}}^2}}\sin(2\theta)\right] \tag{213}$$

and

$$x_{\text{term}} = r_{\mathcal{P}} \sin\theta \sqrt{1 - \frac{y^2}{r_{\mathcal{P}}^2}} \tag{214}$$

is the x coordinate of the day/night terminator.

Integration scheme. Finally, in this section I describe my integration scheme to compute transit, secondary eclipse, and planet-planet occultation light curves, as well as phase curves for limb-darkened (§10.3.2.1) and eyeball (§10.3.2.2) planets. My method is completely general and works for an arbitrary number of occulting bodies. For simplicity, below I describe the single-occultor case, but extending the method to multiple occultors is trivial. The reader is referred to Figure 92 for an illustration of the method.

Let us label the occulted body by $\mathcal{P}$ and the occultor by $\mathcal{O}$. I place $\mathcal{P}$ at the origin and $\mathcal{O}$ at the point $(x_{\mathcal{O}}, y_{\mathcal{O}})$. $\mathcal{P}$ has radius $r_{\mathcal{P}}$; $\mathcal{O}$ has radius $r_{\mathcal{O}}$. I discretize the radiance gradient of $\mathcal{P}$ (Equation 184 or 192) with $N+1$ spherical segments of constant radiance bounded by N ellipses; these are fully described by Equations (182) and (193)–(201). I orient my coordinate system so that the sub-stellar point is on the x axis.

I wish to compute the total flux of $\mathcal{P}$ occulted by $\mathcal{O}$. To this end, I

1. Identify all functions $\mathbf{f} = \{f_0, f_1, ..., f_{n_f}\}$ in the problem. These are the functions describing $\mathcal{P}$, $\mathcal{O}$, and each of the ellipses in $\mathcal{P}$, which are given by Equation (209). For $\mathcal{P}$, $x_0 = y_0 = 0$ and $a = b = r_{\mathcal{P}}$. For $\mathcal{O}$, $x_0 = x_{\mathcal{O}}$, $y_0 = y_{\mathcal{O}}$, and $a = b = r_{\mathcal{O}}$. For the ellipses, the values of a, b, x_0 and y_0 are given in §10.3.2.2. In the case that the ellipses are circles (for a limb-darkened body or for an airless body at full or new phase), $a = b = r\sin\phi$ and $x_0 = 0$.

2. Compute the antiderivatives of each of the functions in $\mathbf{f}$ to form the set $\mathbf{F} = \{\int f_0(x)dx$, $\int f_1(x)dx$, ..., $\int f_{n_f}(x)dx\}$. These are given by Equation (210).

3. Identify the points of intersection between all curves. These are the intersections of $\mathcal{P}$ and $\mathcal{O}$ (Equation 202) and the intersections of $\mathcal{O}$ and each of the ellipses (solutions to Equation 206). Reject points that lie beyond the limb of $\mathcal{P}$ (Equation 201).

4. Identify the extrema of all functions that lie within or on the edge of both $\mathcal{P}$ and $\mathcal{O}$. The extrema of the ellipses are given by either $x = x_0 \pm b$ or $x = x_0 \pm x_{\text{limb}}$, depending on

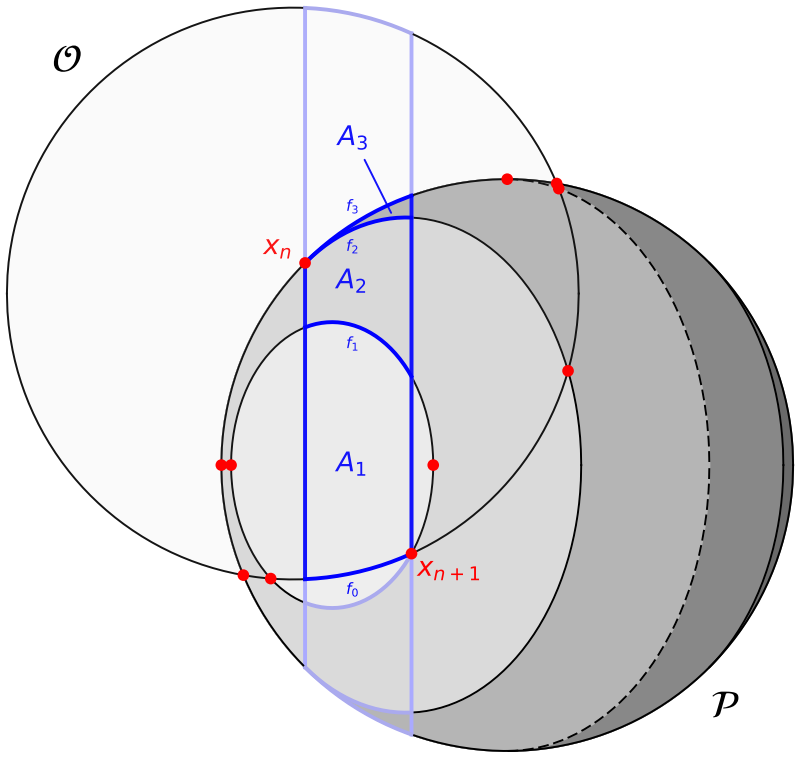

Figure 92: An example of the integration scheme for a planet-planet occultation. The occultor $\mathcal{O}$ is at the top left and the occulted planet $\mathcal{P}$ is at the bottom right. The latter is an airless body at a phase angle $\theta = \frac{\pi}{4}$ with a radiance given by Equation (192). This gradient is discretized into regions of constant radiance, shaded accordingly in the figure. The day/night terminator is indicated by the dashed curve, and the sub-stellar point faces the observer. The flux of $\mathcal{P}$ that is occulted by $\mathcal{O}$ is computed by identifying all intersection points x_n between curves (red points) and summing the integrals over each of the regions in both $\mathcal{P}$ and $\mathcal{O}$ that are bounded by adjacent pairs of these points. As the boundary functions f_j of these regions (dark blue curves) are either circles or ellipses, all integrals are analytic. The total occulted flux is then the product of the area A_j of each region and its radiance.

whether or not the entire ellipse is visible. The extrema of $\mathcal{P}$ and $\mathcal{O}$ are simply $x = \pm r_{\mathcal{P}}$ and $x = x_{\mathcal{O}} \pm r_{\mathcal{O}}$, respectively.

5. Sort the points identified in steps 3 and 4 in increasing order of their x coordinate to obtain the set of integration limits $\mathbf{x} = \{x_0, x_1, ..., x_{n_x}\}$.

6. For each adjacent pair of limits $\{x_n, x_{n+1}\}$ in $\mathbf{x}$, identify the the members of $\mathbf{f}$ that are defined over the interval (x_n, x_{n+1}). Evaluate each function at the midpoint, $x_{n+\frac{1}{2}} = (x_n + x_{n+1})/2$, and discard those that lie outside of either $\mathcal{P}$ or $\mathcal{O}$. Sort the remaining functions in increasing order of their value at $x_{n+\frac{1}{2}}$ to obtain the set $\mathbf{f^n}$ and the set of corresponding antiderivatives $\mathbf{F^n}$.

7. Each adjacent pair of functions $\{F_i^n, F_{i+1}^n\}$ in $\mathbf{F^n}$ bounds an area A_j of constant radiance on $\mathcal{P}$ that is occulted by $\mathcal{O}$. For each such pair, compute the area by evaluating

$$A_j = I_{i+1} - I_i, \tag{215}$$

where

$$I_{i+1} = F_{i+1}^n(x_{n+1}) - F_{i+1}^n(x_n) \tag{216}$$

and

$$I_i = F_i^n(x_{n+1}) - F_i^n(x_n) \tag{217}$$

are the integrals of the upper and lower functions bounding the region, respectively.

8. Compute the zenith angle of this region from the equations derived earlier in this section; the radiance $B_{\lambda,j}$ of this region is that of the discretized radiance grid at this zenith angle.

9. The occulted flux per unit wavelength from the j^{th} region is then simply

$$\Delta\mathcal{F}_{\lambda,j} = \frac{A_j B_{\lambda,j}}{d^2}. \tag{218}$$

where d is the distance to the system.

The total occulted flux per unit wavelength $\Delta\mathcal{F}_\lambda$ is the sum of all $\Delta\mathcal{F}_{\lambda,j}$ computed by iterating over all pairs of adjacent limits (step 6) and all pairs of adjacent boundary functions (step 7). Note that $\Delta\mathcal{F}_\lambda$ is technically a spectral flux density (units of power per unit area per unit wavelength), which may be integrated over a given wavelength range to give a flux.

Once $\Delta\mathcal{F}_\lambda$ is known, the flux from the body received at Earth is

$$\mathcal{F}_\lambda = \mathcal{F}_\lambda^0 - \Delta\mathcal{F}_\lambda, \tag{219}$$

where $\mathcal{F}_\lambda^0$ is the flux received at Earth if no occultation were occuring. This may be computed by following the steps above with a fictitious occultor $\mathcal{O}$ covering the entire disk of the planet. In the limb-darkened planet limit, this term is constant and need only be computed once; for eyeball planets, $\mathcal{F}_\lambda^0(t)$ is the planet phase curve, and must be computed at every step. However, in the limit that the orbital parameters are constant in time, the phase curve need only be computed over a single orbit of the planet, which greatly reduces computation time.

10.3.2.4 Phase curves

For an airless planet, the spatially variable radiance map results in a periodic phase curve signal over the course of its orbit. `planetplanet` computes phase curves using the same integration scheme described above. Figure 93 shows a sample phase curve for an inclined, eccentric "eyeball" planet.

10.3.2.5 Model validation

I validate my light curve model against `pysyzygy`[22] and `batman` (Kreidberg, 2015) (for transit light curves) and against a two-dimensional integration of the radiance map (for planet-planet occultations). `pysyzygy` is a standard implementation of the Mandel and Agol (2002) transit model for limb-darkened light curves, while `batman` is an adaptive algorithm that, similarly to `planetplanet`, discretizes the surface brightness gradient into concentric rings and computes light curves semi-analytically with high precision. I find that the transit light curves generated by `planetplanet` agree with those of `pysyzygy` and `batman` at the < 1 ppm level over a range of orbital parameters and limb darkening coefficients.

[22] `https://github.com/rodluger/pysyzygy/`

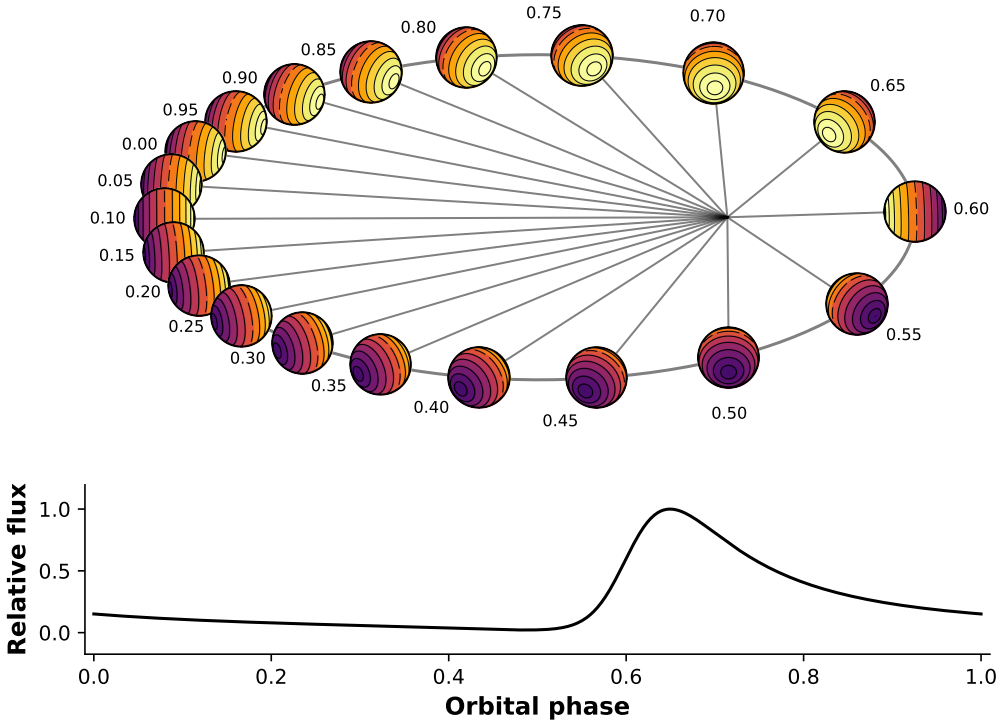

Figure 93: Orbital geometry (top) and phase curve (bottom) of an "eyeball" planet in an eccentric, inclined orbit computed with `planetplanet` using the relations derived in §10.3.2.3. The eccentricity vector is $(e \sin\omega, e\cos\omega) = (0, 0.5)$, with pericenter to the right of the plot. The orbital plane is parallel to the x axis ($\Omega = 0°$) and inclined into the sky plane with $I = 60°$. Orbital phases are labeled; these are defined such that transit would occur at a phase of 0.5. The lower panel shows the phase curve of the planet in arbitrary units.

I validate my occultation light curves by direct integration of the planet emission profile in the area of intersection between the planet disk and the occultor disk. I discretize the planet disk with a Cartesian grid and compute the radiance at each point from Equation (192), with a temperature given by Equation (191) and zenith angle given by Equations (212)–(214). I find that my code matches the occultation light curves generated in this fashion to within the error level of the direct integration procedure.

10.4 Application to TRAPPIST-1

In this section I present my results for the TRAPPIST-1 system, but my methodology is general and can be used to model planet-planet occultations in any nearby multi-planetary system.

10.4.1 Dynamics

The frequency of occurrence and the dynamical properties of planet-planet occultations (including, for instance, their distribution in orbital phase, their durations, and their impact parameters) are extremely sensitive to the full three-dimensional architecture of a planetary system. Transit photometry and TTV analyses have thus far constrained many of the orbital parameters of the seven planets in TRAPPIST-1. In my analyses below, I compute the statistics of planet-planet occultations by sampling the posterior distributions reported in Gillon et al. (2017) and Luger et al. (2017). At present, the eccentricities of the planets are constrained with only upper limits ($\lesssim 0.01$). For these, I use estimates based on a migration and tidal evolution model (Luger et al., 2017), drawing the longitude of pericenter from uniform distributions in the range $(0, 2\pi]$. Table 9 shows the distributions assumed for the orbital parameters of each of the planets.

Currently, the largest source of uncertainty relevant to PPOs are the longitudes of ascending nodes (Ω) of the planets, which are completely unconstrained. For an edge-on system like TRAPPIST-1, Ω is the angle of rotation of the orbital plane on the sky relative to some reference direction. In general, transit light curves are completely insensitive to the value of Ω; however, the relative value $\Delta\Omega$ for each pair of planets controls their sky-projected separation, such that planets with large $\Delta\Omega$ may never occult each other away from the disk of the star. In the following section, I describe a Monte Carlo technique co-I Eric Agol and I developed to place constraints on $\Delta\Omega$ for the TRAPPIST-1 planets.

10.4.1.1 Coplanarity of TRAPPIST-1

For a perfectly coplanar, circular planetary system, the planets should have $T' = T/\sqrt{(1+p)^2 - b^2} \propto P^{1/3}$, where b is the impact parameter, $p = r_{\mathcal{P}}/R_\star$ is the planet/star radius ratio, a is the semi-major axis, P is the period, and T is the transit duration from first contact to last. Figure 94 shows the value of T' for each of the TRAPPIST-1 planets alongside this relation. The transit durations satisfy the above relation surprisingly well, suggesting a very coplanar system.

The coplanarity of the system can be used to place strong prior bounds on the values of $\Delta\Omega$ for each of the planets. To quantify this coplanarity, we generated random planetary systems with planets in different orbital planes and simulated observations of each system, assigning a probability to each based on how well the observed transit durations matched the simulated durations.

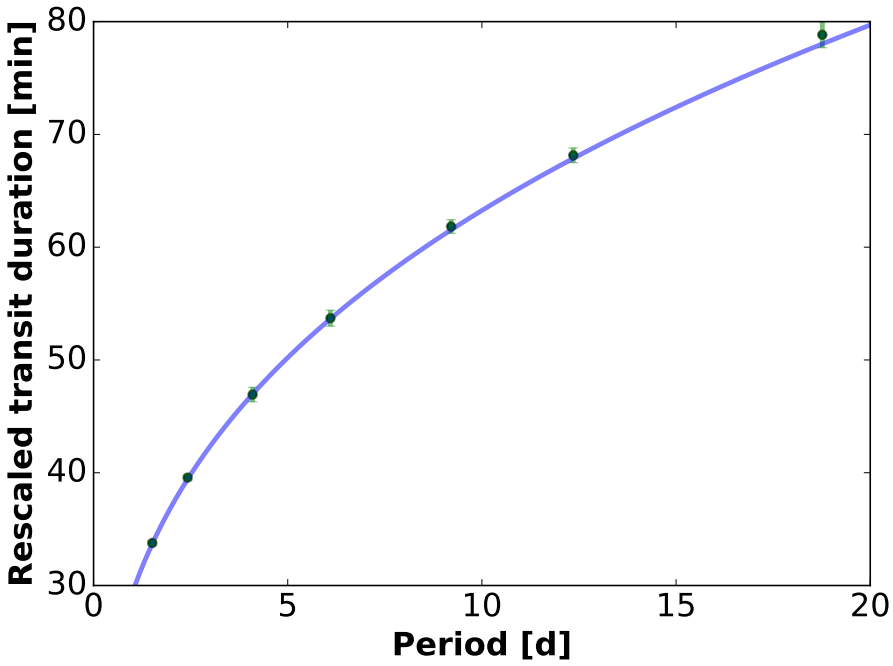

Figure 94: Rescaled transit duration, T', versus orbital period. The median impact parameter, b, and radius ratio, $p = r_p/R_\star$, is used to compute $T' = T/((1+p)^2 - b^2)^{1/2}$. A curve with $T' \propto P^{1/3}$ is overplotted assuming a stellar density of $\rho_\star = 51\rho_\odot$.

We randomized orbits with periods drawn from the observed values of the seven planets and zero eccentricity. We justify the assumption of zero eccentricity since for the small ($\lesssim 0.01$) values estimated from the tidal/migration simulations (Luger et al., 2017), the transit duration should only be affected by <1%, which is comparable to the uncertainty on the transit durations of the outer five planets. The inner two planets have smaller duration uncertainty, but are expected to have even smaller eccentricity due to tidal circularization. We drew the angular momentum vectors of the planets in each simulated system from a Gaussian in polar angle, ϑ, with an assumed width, σ_ϑ, relative to the system angular momentum axis. We drew the azimuthal angle of the angular momentum vector uniformly, and finally drew an observer from a uniform location within 2° of edge-on (since outside this region one or more of the planets are not seen to transit). We made a grid of (σ_ϑ, $\rho_\star$) values, where $\rho_\star$ is the stellar density, and computed the expected transit duration for each planet, drawing the planet radius ratios from the values observed by Gillon et al. (2017) and Luger et al. (2017). Note that we did not fit the observed impact parameters, as these are less well constrained from the observed transits and are thus strongly correlated with one another

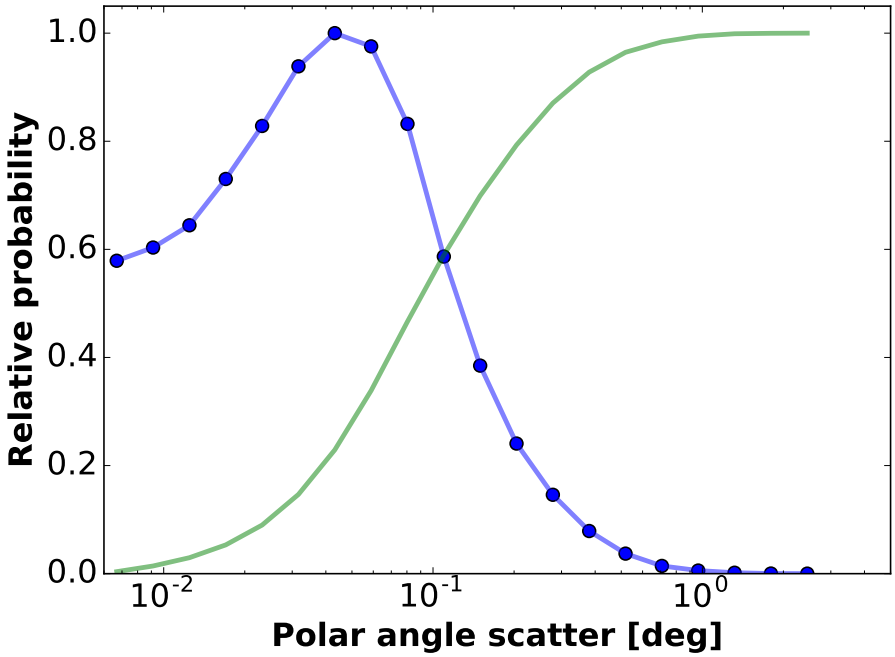

Figure 95: Probability of polar angle scatter, σ_ϑ, marginalized over $\rho_\star$ and assuming circular orbits for the planets (blue dots/line). Cumulative probability distribution shown in green.

and with the transit durations. The transit durations, on the other hand, are not correlated amongst the planets as they are well constrained by the data. For each grid point of $(\sigma_\vartheta, \rho_\star)$, we ran 10^6 simulations of planet system plus observer, computing the probability of the transit durations compared to the durations measured by Gillon et al. (2017); if one or more simulated planets does not transit, we set the probability to zero. We summed up the probabilities over all of the simulations for each grid point, assigning the summed probability to each grid point. This procedure yields a joint probability distribution estimate of the density of the star and the coplanarity of the planets.

Figure 95 shows the inferred value of σ_ϑ, which indicates that this system is *extremely* coplanar, with $\sigma_\vartheta < 0.3°$ at 90% confidence. Figure 96 shows the inferred stellar density of $\rho_\star/\rho_\odot = 51.14 \pm 0.67$, which has a narrower distribution than that computed in Gillon et al. (2017). The more precise value of the stellar density results from the assumption that the planets' angular momentum vectors are drawn from a single distribution rather than allowing the inclinations of each planet to vary independently. We have tried relaxing the circularity constraint, and found a consistent result, as expected.

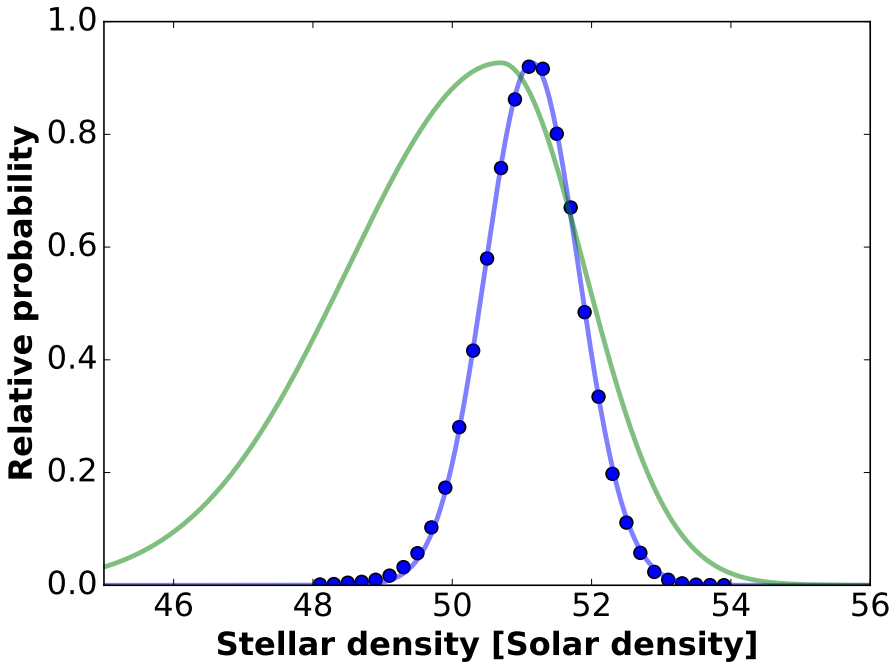

Figure 96: Density of the star, $\rho_\star$, in units of the Solar density, $\rho_\odot$, marginalized over σ_ϑ. Best-fit Gaussian (blue curve), and double-sided Gaussian (green) with uncertainties from Gillon et al. (2017).

10.4.1.2 Sample integration of TRAPPIST-1

At present, the uncertainties on the orbital parameters of the TRAPPIST-1 planets are too large to permit a deterministic prediction of the time of planet-planet occultations. Nevertheless, I may compute the frequency of occultations and their distribution in phase space by drawing planet properties from their respective distributions (Table 9). I assume a stellar mass $M_\star = 0.0802 \pm 0.0073\ \mathrm{M}_\odot$ (Gillon et al., 2017) and radius $R_\star = 0.121 \pm 0.003\ \mathrm{R}_\odot$ (Burgasser and Mamajek, 2017). Planet radii are self-consistently computed from the stellar radius and the transit depths in Table 9. I draw the longitudes of ascending node of each planet from a Gaussian with standard deviation σ_Ω, which for simplicity I take to be the same as σ_ϑ, whose distribution was derived in §10.4.1.1. This is technically incorrect, as the polar angle, ϑ, of the angular momentum vector of a planet has both an inclination component and a longitude of ascending node component, requiring $\sigma_\Omega \leq \sigma_\vartheta$. The σ_Ω distribution does not have an analytic expression resulting from my dynamical simulations, but my choice is conservative in the sense that the mutual inclinations of the planets are slightly *smaller* than I assume here. In practice, however, since the system is so coplanar, this

choice does not significantly affect my conclusions.

For a given set of orbital parameters, I use `planetplanet` to integrate the orbits forward in time with the `REBOUND` N-body code in order to capture dynamical interactions among the planets. I use the 11^{th} order `WHFAST` symplectic integrator with a timestep of one hour, which I find leads to a maximum error in the sky-projected separation of two planets of less than one percent the planet radii over one (Earth) year. I oversample each timestep on a finer 10 second grid, over which I take Keplerian steps. I use Equation (181) to determine when occultations occur.

Figure 97 shows the results of a sample 3-year integration of the TRAPPIST-1 system. System parameters were drawn from the distributions in Table 9. The figure shows a top-down view of the system, with the observer located towards the bottom and planets orbiting counter-clockwise. Each planet-planet occultation event is indicated as a circle, whose color, size, and transparency indicate the occultor, the duration of the event, and the minimum impact parameter of the occultation, respectively; refer to the legends for details. The "X"s indicate mutual transits, which are planet-planet occultations on the disk of the star (see §10.5.6).

For this particular set of system parameters, 1,036 planet-planet occultations occur among the TRAPPIST-1 planets over the course of 3 years, averaging to about one per (Earth) day. I find that this is a typical rate when marginalizing over the uncertainties on the orbital parameters. Nevertheless, it is evident from the figure that the vast majority of occultations are very short-lived ($\ll 10$ minutes), which poses challenges to their detectability. Longer, non-grazing occultations are less frequent, but still occur several tens of times per year for certain planet pairs (see §10.4.1.3). Occultations of b by c (orange circles along the innermost orbital track) are the most common, and occur at all orbital phases of b. Their duration is a strong function of phase: occultations on the near side of the star (mean longitude $\lambda_b \approx -90°$) last upwards of 30 minutes, while those on the far side of the star ($\lambda_b \approx 90°$) last on the order of tens of seconds. This is because the former case happens when the sky-projected velocity vectors of b and c are aligned, resulting in the smallest relative velocity among the two planets, while the latter case happens when c is on the near side of the star and b is on the far side, such that the two planets are moving in opposite directions on the sky and thus have large relative velocities. Two occultations of b stand out, as they are much longer than one would expect given the relative velocities of the planets. These are double (or triple) occultations, in which the relative velocity of the two planets changes sign one

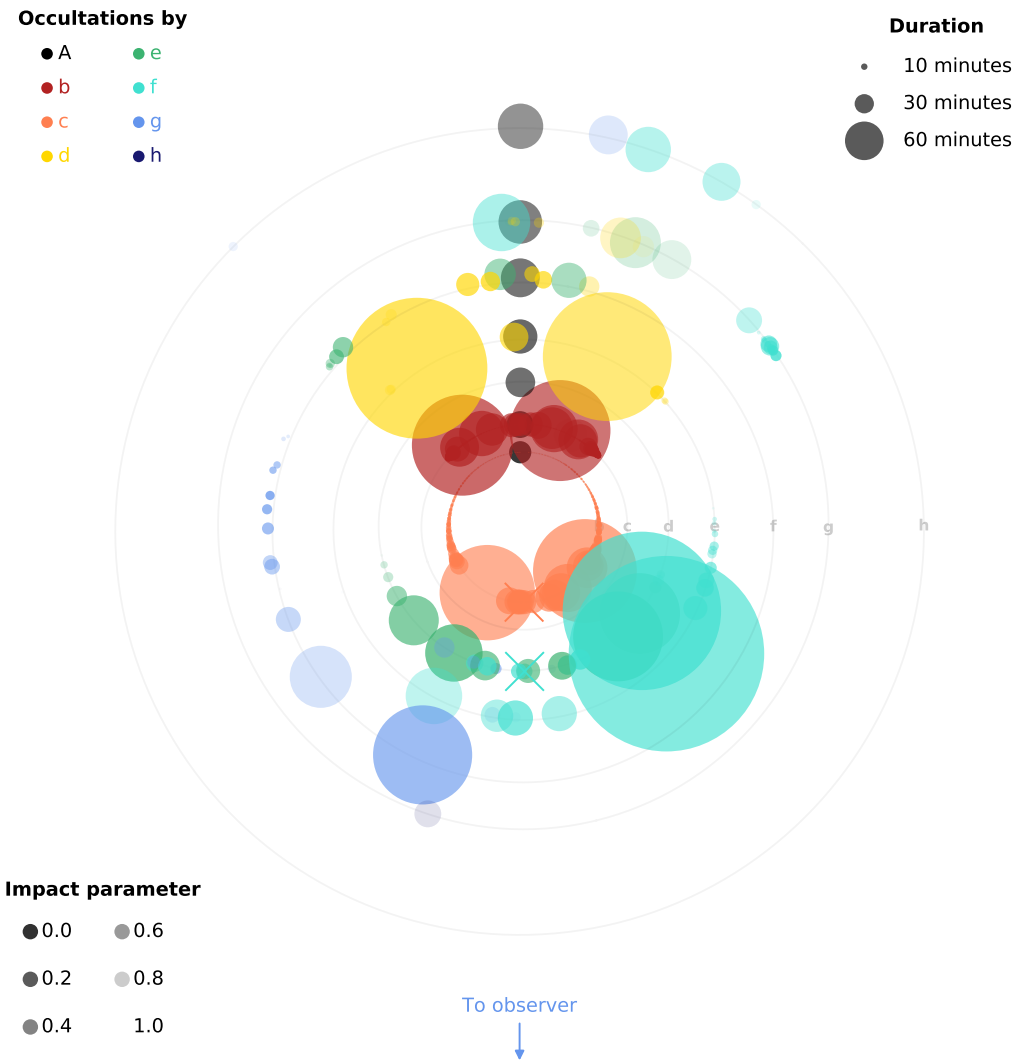

Figure 97: Three years of planet-planet occultations in TRAPPIST-1, for a single random draw from the orbital parameter distributions in Table 9. The system is seen from above, with the observer located towards the bottom of the plot and planets orbiting counter-clockwise. The initial orbital outlines of each of the seven planets are shown in grey, and each occultation is indicated as a colored circle placed at the location of the occulted planet at the time of the event. Circle colors correspond to different occultors: black is an occultation by the star (i.e., secondary eclipse), red is an occultation by b, and so forth (see legend at top left). Circle sizes are proportional to the event duration (legend at top right), and the transparency is proportional to the impact parameter (legend at lower right). The "X"s indicate mutual transits, or planet-planet occultations occurring on the face of the star.

or more times over the course of the event and the occultation occurs more than once. The large orange circle left of the center of the plot is an occultation in which b has just passed quadrature ($\lambda_b \approx 180°$), moving primarily towards the observer, while c is closer to transit, moving faster on the sky plane. The occultation is prolonged due to the fact that as b passes quadrature it speeds up on the sky and catches up to c. Eventually, as b approaches transit, its sky velocity exceeds that of c and it emerges from the occultation.

The second most common—and likely the most observable—set of occultations are those of c by b (red circles). Given semi-major axes a_b and a_c of b and c, respectively, these necessarily occur when c is on the far side of the star with mean longitude

$$\frac{\pi}{2} - \sin^{-1}\left(\frac{a_b}{a_c}\right) \leq \lambda_c \leq \frac{\pi}{2} + \sin^{-1}\left(\frac{a_b}{a_c}\right)$$
$$43° \lesssim \lambda_c \lesssim 137°. \tag{220}$$

These are therefore occultations of primarily the dayside of c, which could have the strongest signal (see §10.4.2). When b is also on the far side of the star, these occultations last between 30 and 60 minutes. Moreover, since b is slightly larger than c, these are often full or near-full occultations, although the impact parameter is sensitive to the longitudes of ascending nodes of the two planets. Two double occultations are again visible as the largest red circles.

Occultations among the other planets are less frequent and tend to be more intermittent because of the longer orbital periods. Interestingly, while occultations of b and c tend to be symmetric about the line of sight to the star, occultations of the other planets may preferentially occur on one side of the star. This is likely due to the near-resonant state of the planetary system. For this particular set of orbital parameters, the asymmetry is best seen in occultations of f by g, which occur solely to the left of the star.

One final consideration is that occultations of the farther-out planets can be much longer because of their smaller relative velocities. In the figure, one occultation of e by f exceeds 5 hours, and two occultations of e by d exceed 3 hours. Provided stellar variability is properly removed, longer occultations are in principle more detectable (see §10.4.2).

10.4.1.3 PPO statistics for TRAPPIST-1

The results presented in the previous section were for a single draw from the distributions of orbital parameters allowed by observations of the system up to the present time. Planet-planet occultations are aperiodic, and therefore different realizations of the system lead to very distinct occultation maps like the one shown in Figure 97. It is therefore instructive to consider the statistics of PPOs in TRAPPIST-1 when marginalizing over the uncertainties on the orbital parameters, so that I may quantify the expected frequency of PPOs given all the information known at present. I therefore ran an ensemble of 22,000 1-year integrations of TRAPPIST-1, each time drawing orbital parameters from their respective distributions and integrating the system forward in time as before.

Figure 98 shows the posterior distributions for the mean longitude λ, impact parameter b, and duration Δt of all planet-planet occultations of TRAPPIST-1b. These include occultations by all planets, though occultations by c are the most common. PPOs occurring behind the star (i.e., during secondary eclipse) are not included in the figure, as they are not observable. The plots at the top of each column are the fully marginalized distributions. These show that occultations of b occur primarily near quadrature ($\lambda_b \approx 0^\circ$ or 180°), and somewhat less frequently near transit ($\lambda_b \approx -90^\circ$). Occultations occur over the full range of impact parameters, with grazing ones being somewhat more common. The duration peaks around 5 minutes, with a long tail toward long (> 1 hour) occultations. Below the marginalized distributions, I plot the joint posterior distributions for each combination of the three parameters. The strongest correlation is that between phase and duration. The shortest events are those where b is on the far side of the star near secondary eclipse, when the occultor is necessarily on the near side of the star moving in the opposite direction. Conversely, for $\lambda_b \approx -150^\circ$ or -30°, the long occultations of b by c discussed in §10.4.1.2 are possible, resulting in events that can exceed two hours.

Figure 99 shows the same distributions, but for occultations of TRAPPIST-1c. The distributions of impact parameter and duration are quite similar to those of b, but the phase posterior now peaks at $\lambda_c \approx 43^\circ$ or 137°, as expected from Equation (220). These are primarily occultations by b when its sky-projected velocity is close to zero, resulting in a much higher occultation probability. Even though my calculations marginalize over the uncertainties on all orbital parameters, the two peaks in the histogram are *extremely* narrow, with FWHM $\sim 2^\circ$, as this width is due primarily to

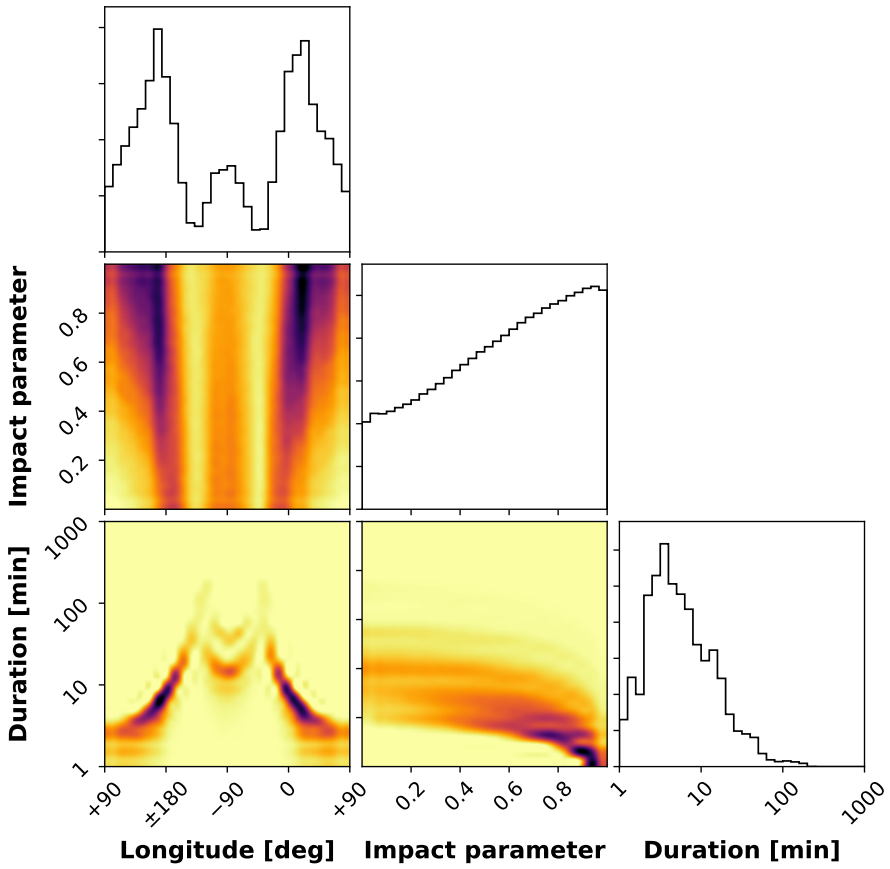

Figure 98: Posterior distributions for the mean longitude, impact parameter, and duration of all occultations of TRAPPIST-1b, plotted using the `corner` package (Foreman-Mackey, 2016). Panels at the top of each column are the marginalized distributions; other panels are the joint posteriors, where the darkness is linearly proportional to the probability density. These distributions are marginalized over the uncertainties on the orbital parameters of all bodies in the system. Occultations of b are piled up in phase near quadrature and transit. Grazing occultations are most common, but near-full occultations still occur frequently. Occultations lasting about 5 minutes are typical, with a long tail toward durations exceeding one hour. The longest events are occultations by c occurring between transit and quadrature.

the uncertainty on the semi-major axes of the two planets. As I argue in §10.5, this is the ideal mean longitude at which to observe TRAPPIST-1 in search for occultations of c.

Plots for occultations of the other five planets are qualitatively similar and are available on `github`. In general, occultations of all planets are most common in between quadrature and secondary eclipse, when most of the planet's dayside and part of the nightside is facing the observer. This makes planet-planet occultations a powerful method to probe the day-night temperature contrast of the TRAPPIST-1 planets, which I discuss further in §10.4.2.

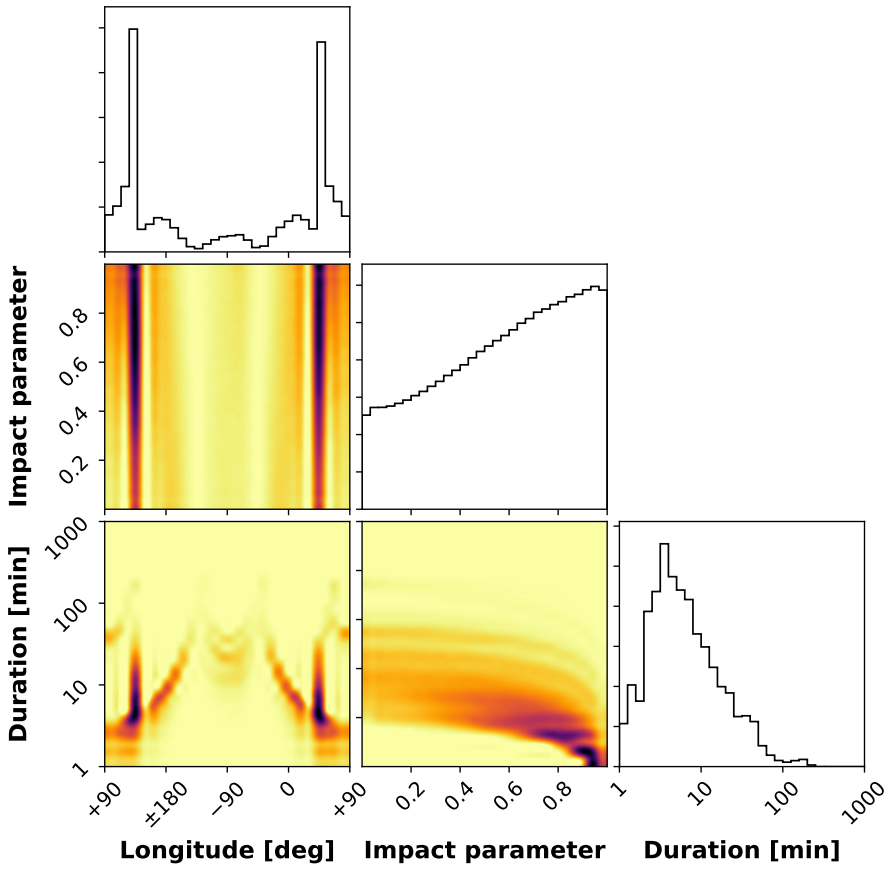

Figure 99: Same as Figure 98, but for occultations of TRAPPIST-1c. The distributions are qualitatively similar to those of b, although now the vast majority of occultations are by b, occurring halfway between quadrature and secondary eclipse.

Finally, in Figure 100, I marginalize over all system parameters and plot histograms of the number of occultations of each of the planets expected over the course of one (Earth) year. I count only "potentially observable" occultations, defined as those with impact parameter less than 0.5 and durations longer than 10 minutes. The mean and standard deviation from a Gaussian fit to these histograms is shown in the legend. Given our current knowledge of the orbital parameters of TRAPPIST-1, I expect there to be 21.1 ± 9.7 occultations of TRAPPIST-1b in any given year. Occultations of c are just as common, occuring at a rate of 21.3 ± 8.8 yr^{-1}. The other planets are occulted less frequently, but are still expected to undergo several occultations per year.

Note that there is a peak at 0 yr^{-1} in the distribution for TRAPPIST-1b, indicating that for about 4% of system realizations there are no occultations of planet b. These are likely realizations

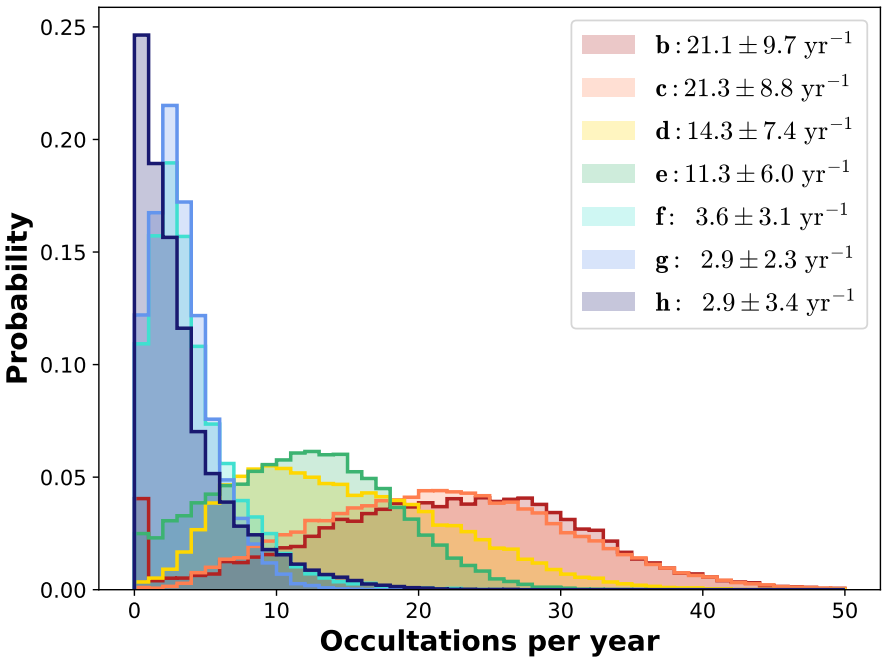

Figure 100: Frequency of planet-planet occultations in TRAPPIST-1 for each of the planets, marginalized over all system parameters. Only occultations with impact parameter $b < 0.5$ and duration $\Delta t > 10$ minutes are included. The legend shows the expectation value and the standard deviation for the occultation frequency of each planet.

in which b and c are mutually inclined by a significant amount, precluding occultations on the near side of the star. Interestingly, there is no such peak for TRAPPIST-1c: fewer than 0.1% of runs result in no significant occultation of c. In other words, potentially observable occultations of TRAPPIST-1c occur at least once per Earth year with greater than 99.9% confidence, with an expectation value of 21 yr^{-1}.

10.4.2 Photometry

I used `planetplanet` to produce synthetic light curves of TRAPPIST-1 over a range of wavelengths, modeling all transits, secondary eclipses, planet-planet occultations, and phase curves. Stellar variability is not modeled. In what follows, I assume an albedo of 0.3 for all planets for definiteness, unless otherwise stated. Airless bodies are assumed to have a fixed nightside temperature $T_{\rm night} = 40$ K, and bodies with thick atmospheres are assumed to have a uniform surface radiance with no limb darkening, though these assumptions can be changed or relaxed in `planetplanet`. The stellar luminosity is taken to be $L_\star = 0.000524 \pm 0.000034$ L$_\odot$ (Gillon et al., 2017) and the effective

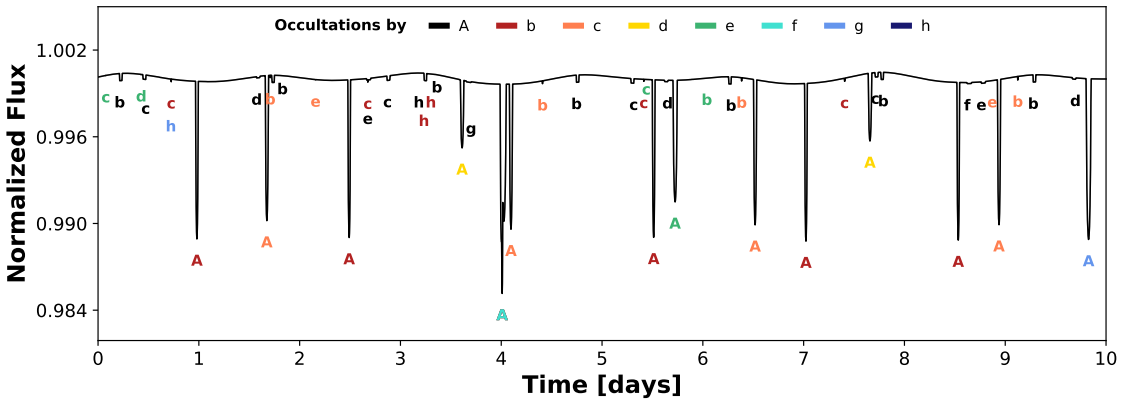

Figure 101: A light curve of TRAPPIST-1 over ten days at 15μm, generated from a random draw from the orbital parameter distributions (Table 9). All transits, secondary eclipses, and planet-planet occultations are labeled with the occulted body name and colored according to the occultor (see legend at the top). The planets are assumed to have thin atmospheres, giving rise to a bright dayside and a dark nightside. This results in distinct phase curves, which are visible in the figure; the overall signal is dominated by the phase curves of b and c. Stellar variability is not modeled. Note that a simultaneous (but not mutual) transit of b and f occurs at $t = 4$ days.

temperature is computed from this value via the radius (§10.4.1.2).

10.4.2.1 Sample light curve

Figure 101 shows the full light curve of the TRAPPIST-1 system over the course of ten days for a random realization. The light curve is computed at a wavelength of 15 μm and is normalized to a mean of unity. Planets are assumed to be airless, giving rise to the prominent phase curve modulation. All events are labeled with the name of the occulted body and colored according to the occultor. Transits by nearly all planets (labeled "A") are evident, with depths on the order of 1%. Secondary eclipses (black labels) are also visible for most planets, with depths ranging from 0.05% (500 ppm) for b to 0.0015% (15 ppm) for h. These are also the magnitudes of the phase curve amplitudes; the overall phase curve is dominated by that of b and c. All other labeled events are planet-planet occultations (4 of b, 6 of c, 1 of d, 2 of e, and 3 of h). The occultation of c by b at $t = 7.4$ days has a depth of 200 ppm, and that of b by c at $t = 9.1$ days has a depth of 250 ppm, which is within a factor of 2 of the depth of their secondary eclipses.

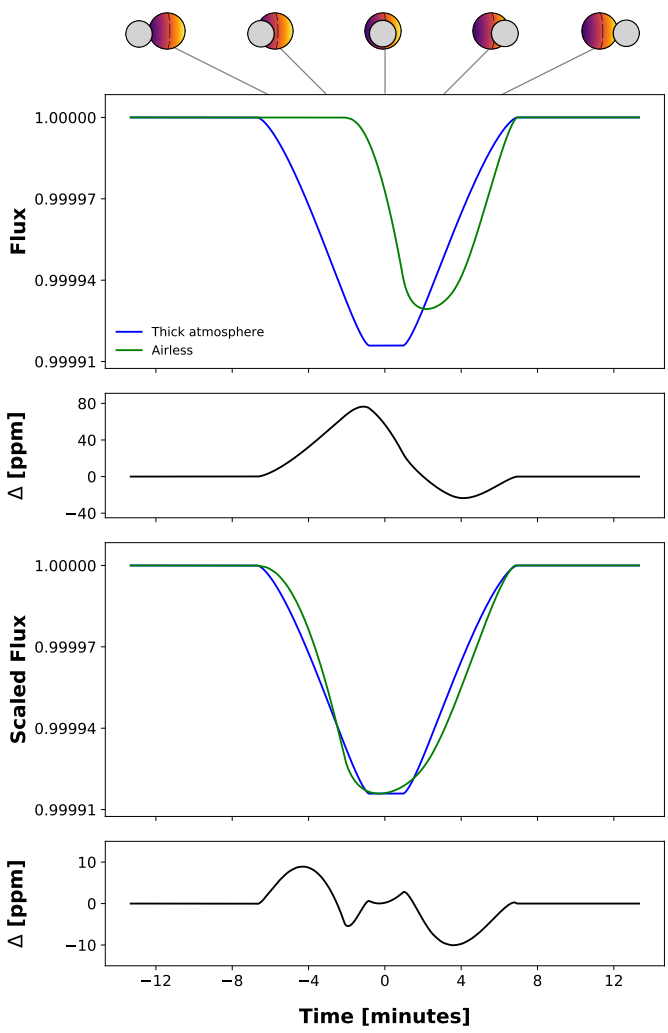

Figure 102: An occultation of TRAPPIST-1c by TRAPPIST-1d at 15μm as c approaches quadrature, for two different atmospheric regimes: the thick atmosphere limit (blue curves) and the airless body limit (green curves). In the former case, the planet disk is radially symmetric, and the light curve is symmetric about the midpoint of the occultation ($t = 0$). In the latter case, the stark day/night temperature contrast leads to an asymmetry in the light curve and a shift in the time of flux minimum. The light curves at 15 μm are shown in the top panel. Below it, I plot the difference of the two curves, showing that the day/night temperature contrast corresponds to a $\sim$80 ppm signal. However, if the time of occultation, the albedo of the planet, and the duration of the event are not known *a priori*, discrimination between an airless planet and one with a thick atmosphere must be made based on the shape alone. To this end, in the third panel, I shift and scale the green curve so that the timing, duration, and depth coincide with those of the blue curve. The residuals are shown in the bottom panel: the curves are different at the $\sim$10 ppm level.

10.4.2.2 Sample occultation

Figure 102 shows a sample occultation of TRAPPIST-1c by TRAPPIST-1d. I chose this occultation in particular because it highlights the ability of PPOs to probe a planet's day/night temperature contrast and even generate crude surface maps in the mid-infrared. This occultation occurs shortly after c has passed quadrature, such that about half its dayside and half its nightside are visible. TRAPPIST-1d is on the near side of the star, close to halfway between quadrature and transit, and the occultation occurs as it overtakes c on the sky. At occultation center, d is fully within the disk of c, but because d is smaller, only about half of the disk of c is occulted. The illustration at the top of the figure shows the progression of the occultation: d is the grey disk, and c is the colored disk. Blue indicates the nightside and red indicates the dayside; the dashed grey line is the terminator. The occultation proceeds such that the nightside is occulted first, followed by the dayside. The event lasts a total of 15 minutes.

The top panel of the figure shows the occultation light curve normalized to a stellar continuum of unity. The light curve is plotted for two different regimes of TRAPPIST-1c: a thick atmosphere (blue; see §10.3.2.1) and a thin/negligible atmosphere (green; see §10.3.2.2). The former case leads to a light curve that is symmetric about the occultation midpoint, reflecting the assumed radial symmetry of the planet disk. The latter case leads to a distinctly asymmetric light curve, whose time of minimum flux is shifted by ∼2.5 minutes. In the airless regime, occultation of the cold ($T = 40$ K) nightside leads to a negligible change in flux; ingress only occurs when TRAPPIST-1d begins to occult the bright dayside of c and flux minimum occurs when the region in the vicinity of the substellar point is occulted. The occultation depth in the airless case is thus smaller by about 15% compared to the thick atmosphere case.

The next panel shows the difference between the two light curves in parts per million (ppm) of the total signal. If the timing of the occultation is precisely known *a priori*, these residuals suggest that one can discriminate between the airless case and the thick atmosphere case at the ∼80 ppm level, assuming the observation cadence is short enough and sources of correlated noise such as stellar variability can be properly removed.

However, in the case that the time of occultation is not precisely known—due to uncertainty in the orbital parameters of the two planets—there will be a degeneracy between the day/night temperature contrast of the planet and the orbital parameters of the occulted/occultor pair, as

a change in the latter can easily cause a similar shift in the occultation time. Moreover, if the albedo of the planet is unknown *a priori*, the depth of the occultation is similarly not sufficient to distinguish between an airless body and one with a thick atmosphere. Finally, the duration of the event is in general degenerate with the impact parameter: a shorter event can be explained by either a large day/night contrast or a grazing occultation. Therefore, absent constraints on the orbital parameters and the albedo of the planet, distinguishing between the two atmospheric limits must be done based on the shape of the occultation alone. To this end, in the third panel I plot the same light curves, but having shifted and scaled the light curve corresponding to the airless case so that its duration, time of minimum, and depth match those of the thick atmosphere case. I again plot the difference between the two below, which has a maximum of $\sim$10 ppm during both ingress and egress. This signal is unlikely to be detectable with JWST (see §10.4.3), but constraints on the orbital parameters (from TTVs and secondary eclipses) and on the albedo (from secondary eclipses) can break several of the degeneracies mentioned above, allowing one to constrain the day/night temperature contrast based on the timing, depth, and/or duration of the occultations. I return to this point in §10.5.3.

10.4.3 Observability with JWST

The James Webb Space Telescope (JWST), scheduled for launch in late 2018, will offer an unprecedented view of exoplanetary systems by accessing wavelengths in the range $\sim 0.6 - 30$ μm (Gardner et al., 2006). The mid-IR capability of JWST is uniquely suited for secondary eclipse and PPO observations due to the rise in signal contrast towards longer wavelengths. Since the signals of interest are photons emitted from the planet, stellar photons only contribute noise. It is therefore advantageous to observe such occultation events at longer wavelengths where the stellar flux declines along the Rayleigh-Jeans tail and the thermal flux from temperate planets peaks.

I assess the detectability of PPOs with JWST by simulating time-series filter photometry with the Mid-IR Instrument (MIRI). I consider both shot noise and radiative background noise. My background noise estimates follow the 6-component gray-body emission spectra from Glasse et al. (2015), which account for scattered and emitted zodiacal dust and observatory straylight. MIRI has 9 photometric filters spanning $\sim 5 - 30$ μm (Bouchet et al., 2015). I acquired publicly available MIRI filter response curves online[23]. Although I considered all 9 filters, I choose to focus here on

[23] `http://ircamera.as.arizona.edu/MIRI/pces.htm`

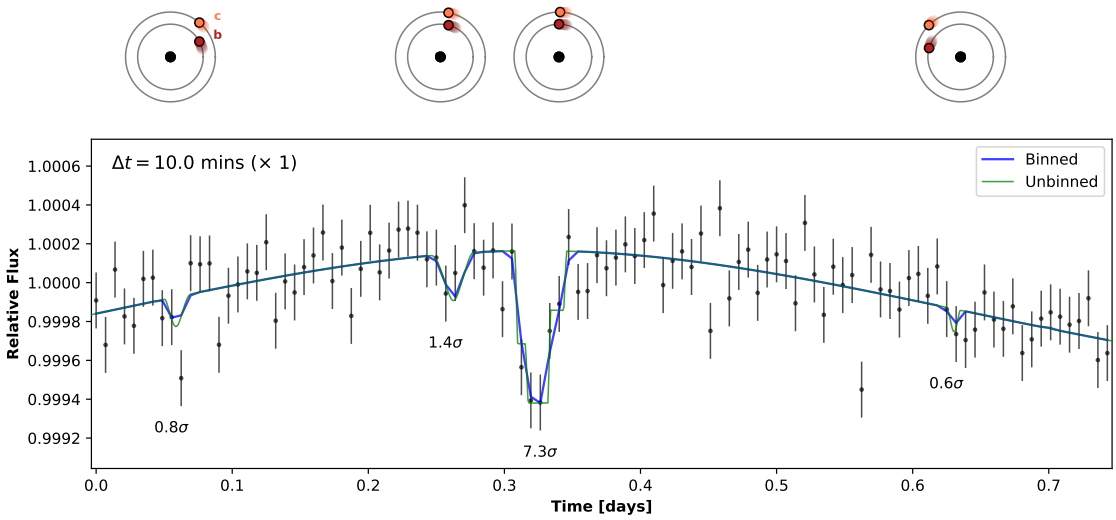

Figure 103: A simulated triple occultation of TRAPPIST-1c by TRAPPIST1-b seen with JWST/MIRI at 15μm with 10 minute exposures over 17 hours. The top panel shows the orbital positions of b and c during each of the events, seen from above the orbital plane; the observer is towards the bottom. The lower panel shows the full light curve (green), the light curve binned to the exposure time (blue), and the simulated observations with 1σ error bars (black). From left to right, c overtakes b and is occulted; b overtakes c and occults it again; b and c are successively occulted by the star; c overtakes b a final time and is occulted. While the simultaneous eclipse of b and c is detectable above the noise (SNR$\sim$7.3) in a single observation, the occultations of c by b are not. The deepest one, occuring at $t = 0.27$ days (when c is close to full phase), has SNR$\sim$1.4. Several of these must be stacked in order to permit detections of PPOs in TRAPPIST-1 with JWST.

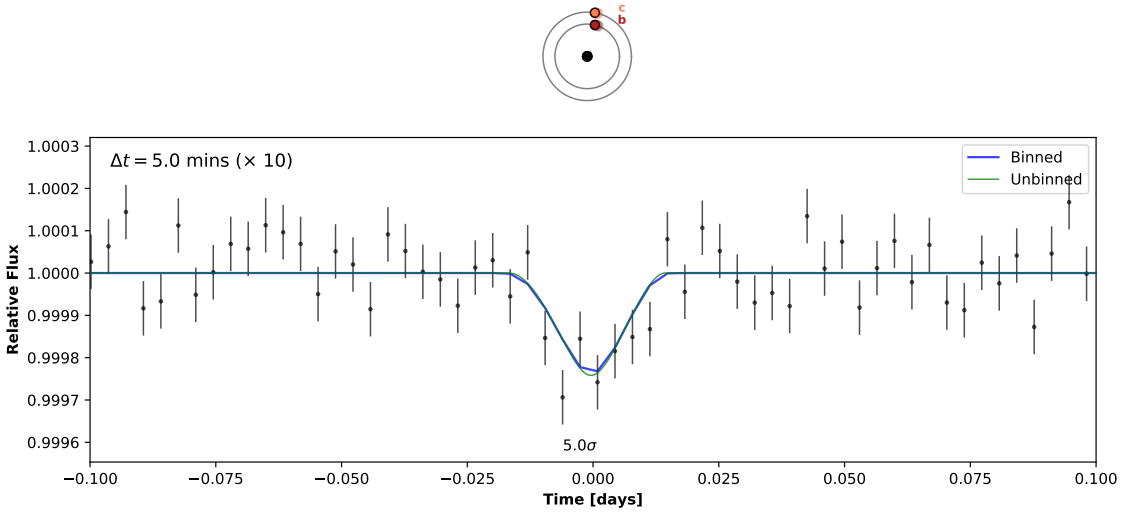

Figure 104: Similar to Figure 103, but showing ten stacked observations of an occultation of TRAPPIST-1c by TRAPPIST-1b seen with JWST/MIRI at 15μm with 5 minute exposures. The phase curves have been removed. The occultation is detectable above the noise with SNR$\sim$5.0.

the 15 μm filter F1500W because I find it to be optimal when considering the rise of both the signal contrast and the background noise towards longer wavelengths[24].

To construct light curves, I assume consecutive images are captured with the same exposure time and neglect readout and reset time. I interpolate the filter throughput curve T_λ to the high resolution spectrum grid F_λ output by the photodynamical model. I then calculate the number of photons detected from the system,

$$N_{\rm sys} = \frac{nAt}{hc} \sum_\lambda F_\lambda \lambda \Delta\lambda T_\lambda, \tag{221}$$

where A is the telescope aperture (25 m^2 for JWST), t is the exposure time, n is the number of observations (in the case of stacking), λ is the wavelength bin, $\Delta\lambda$ is the width of the wavelength bin, and h and c are Planck's constant and the speed of light, respectively. The calculation for the background photons $N_{\rm back}$ is analogous.

I also calculate the signal-to-noise ratio (SNR) for the detection of each occultation event. I borrow from formalism used to calculate the SNR on the detection of molecular features in spectra (Robinson et al., 2016, see their Equation 7), and generalize it to simulated time-series data. Accordingly, the SNR on an occultation is

$$\mathrm{SNR}_{\rm occ} = \frac{\sum_j |N_{{\rm cont},j} - N_{{\rm sys},j}|}{\sqrt{\sum_j (N_{{\rm sys},j} + N_{{\rm back},j})}}, \tag{222}$$

where $N_{\rm cont}$ is the number of photons one would expect to detect if no occultation event occurred, and the sum is taken over all observations j during the time of the occultation. I compute $N_{\rm cont}$ by fitting a 3$^{\rm rd}$ order polynomial to the continuum with all events masked.

Figure 103 shows a simulated triple occultation of TRAPPIST-1c by TRAPPIST-1b observed with JWST/MIRI at 15μm with 10 minute exposure times. Such triple events occur because of retrograde motion: when b is at quadrature and c is approaching secondary eclipse, the sky-projected velocity of b is smaller than that of c. But as b approaches secondary eclipse, it speeds up on the sky and overtakes c. On the opposite side of the star, b slows down once more and c overtakes it. A double secondary eclipse — of both planets — occurs at $t = 0.32$ days. The low frequency modulation of the light curve is due to the combined phase curves of b and c, which are

[24]Calculations using all 9 MIRI filters can be performed using my code on `github`.

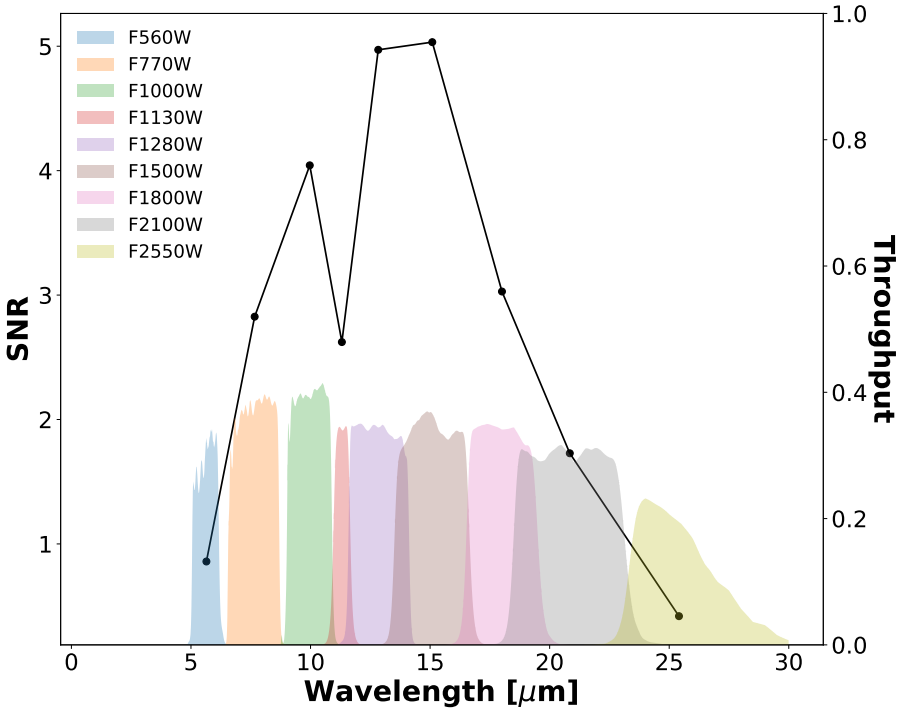

Figure 105: Expected signal-to-noise ratio (SNR; black curve) for ten stacked occultations of TRAPPIST-1c by TRAPPIST-1b observed in each of the nine JWST/MIRI photometric filter bands. Filter throughput curves are shown in color and are plotted on the right y-axis.

assumed to be airless bodies with an albedo of 0.3. For reference, diagrams of the orbital positions of the two bodies are shown above each of the occultations.

In green I show the high resolution noise-free light curve, and in blue I show the light curve binned to the exposure time; the latter is what I would observe in the limit of infinite SNR. Black dots are the simulated noised observations. At SNRs of ∼0.8, ∼1.4, and ∼0.6 (from left to right), none of the three occultations of planet c by b is statistically significant. The double secondary eclipse, on the other hand, is a detectable SNR∼7.3 event. Unless the albedo of TRAPPIST-1c is much lower than 0.3, or it has a much higher surface temperature, individual PPOs of TRAPPIST-1c may be just below the detectability threshold of JWST. As I discuss in §10.5.3, occultations of planet b by planet c should be marginally more detectable, but individual ones are also likely to be below the detectability threshold.

However, given the high frequency of occultations between TRAPPIST-1b and c (Figure 100), the observation of multiple PPOs with JWST could lead to a statistically significant detection. In Figure 104 I plot the result of stacking ten observations of the occultation at $t = 0.27$ days in

Figure 103, this time at a 5 minute cadence. The phase curves and the secondary eclipses have been subtracted out. In this case, stacking ten observations leads to a detectable, SNR$\sim$5.0 PPO.

Figure 105 shows the SNR on ten stacked occultations of planet c by planet b in all nine MIRI filters (5 μm $\lesssim \lambda \lesssim$ 26 μm). The F1280W (12.8 μm) and F1500W (15 μm) filters are approximately equal in their ability to optimally detect this particular PPO event given my assumption here that TRAPPIST-1c is an airless body. However, the presence of an atmosphere will affect the wavelength-dependent thermal emission from the planet and may change the optimal filter band. In particular, if TRAPPIST-1c has an atmosphere with CO_2 then I might expect relatively strong CO_2 absorption at 15 μm, as it is present in the spectra of Venus, Earth, and Mars. CO_2 absorption will make the planet darker at these wavelengths and therefore harder to detect in the F1500W filter. It is worth noting that secondary eclipse observations in both F1280W and F1500W may offer a simple test for a CO_2-bearing atmosphere by measuring the slope between the two eclipse depths. A positive slope may be fit by a featureless blackbody while a flat line or negative slope could indicate CO_2 absorption.

A word of caution is in order here. As PPOs are not periodic or even strictly repeatable events, it is not in general possible to simply stack multiple ones to improve the SNR of the detection, as I did above. Even for PPOs occurring at the same phase between the same two planets, slight differences in the impact parameter and the relative velocities of the two planets are likely to lead to deconstructive interference in the stacked signal. This is particularly problematic when the orbital parameters of the two planets are not well constrained (as is currently the case), since the timing of the occultation—and whether or not one occurs at all—is at best uncertain.

Instead, given a series of n observations of the system, one must jointly model all of them with a full photodynamical model to derive posterior probabilities of the quantities of interest, such as the effective temperature of the occulted planet or its day/night temperature contrast. This allows one to marginalize over the unknown orbital parameters—and over the uncertainty of whether or not an occultation occurs at all—instead of risking the contamination of real occultations with noise when stacking. Figure 104 is therefore merely illustrative of the fact that if the orbital parameters are well constrained, on the order of ten occultations of TRAPPIST-1c are necessary for a robust detection. I return to this point in §10.5.5.

10.4.4 Observability with OST

The Origins Space Telescope (OST; Cooray and Origins Space Telescope Study Team, 2017) is a NASA mission concept currently under study for the 2020 Astronomy and Astrophysics Decadal Survey. OST would be an actively-cooled, mid- to far-IR, large aperture (8–16 m) space telescope, and successor to JWST. OST's access to the far-IR may make it an ideal observatory for thermal emission studies of nearby exoplanets, including PPO observations.

Similar to my estimates for JWST, I assess the observability of PPOs in the TRAPPIST-1 system with OST. Given the active cooling of OST, I neglect background noise and assume photon-limited observations. I simulate time-series photometry for a 12 meter OST primary mirror and investigate filter configurations.

In Figure 106 I plot the signal, noise, and SNR for an occultation of TRAPPIST-1c by TRAPPIST-1b (the same one shown in Figure 104). For my photometry estimates I assume a tophat filter with 30% throughput in a 5 μm-wide filter centered at a range of wavelengths between 5 and 80 μm. Both the signal and noise terms (numerator and denominator in Equation 222, respectively) are monotonically increasing functions of wavelength, but with second derivatives of opposite sign. Interestingly, the SNR curve reveals that the optimal wavelength for the detection of this particular PPO event is $\sim$15 μm, similar to the conclusion I reached with JWST. However, because of the absence of background noise and the larger mirror size, PPOs are detectable at much higher SNR with OST.

In Figure 107 I plot the same triple occultation as in Figure 103, but this time observed at 5 min cadence with OST in a broad mid-IR filter between 10 and 30 μm with a 30% throughput. All three occultations of c by b are statistically significant events and are visible by eye in the data, particularly the SNR$\sim$8.5 event just prior to the double secondary eclipse.

10.4.5 Observability with Spitzer

Finally, I briefly explore the observability of PPOs in the TRAPPIST-1 system with the *Spitzer* Space Telescope. *Spitzer* has a small (85 cm diameter) primary mirror and is limited to photometry at 3.6 μm and 4.5 μm due to the loss of coolant, both of which are detrimental to PPO detectability. I find that a statistically significant (SNR $>$ 3) detection of the occultation of TRAPPIST-1c by TRAPPIST-1b considered in the previous sections would require the observation of on the order

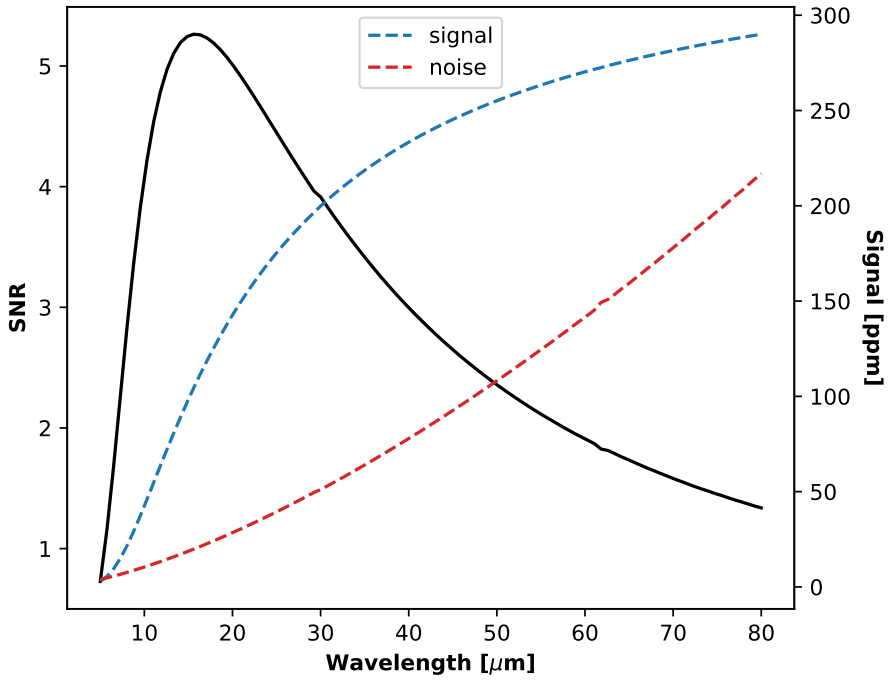

Figure 106: Expected signal-to-noise (SNR; black curve) for an occultation of TRAPPIST-1c by TRAPPIST-1b if observed with the Origins Space Telescope (OST). Dashed blue and red curves are plotted on the right y-axis and show PPO signal and noise terms (see Equation 222), respectively, in units of parts-per-million (ppm) of total photons observed from the system during the occultation. OST calculations assume 30% throughput in a 5 μm wide filter centered at each wavelength.

of 10^4 PPOs at 4.5 μm. However, if the brightness temperature of TRAPPIST-1c at 4.5 μm were ~700 K (~350 K hotter than the equilibrium value at zero albedo), a detection could be made with fewer than 100 stacked observations. This could in principle be possible if TRAPPIST-1c had a thick, greenhouse-heated atmosphere with low opacity at 4.5 μm; the same is true for TRAPPIST-1b. I therefore conclude that PPO observations with *Spitzer* are generally infeasible, except under vigorous greenhouse heating and the presence of transparent atmospheric windows coincident with the *Spitzer* filters.

10.5 Discussion

Here I discuss qualitatively some aspects of PPO detection, including planet mapping (§10.5.1), comparison with phase-curve measurements (§10.5.2), optimistic prospects for detection (§10.5.3), and potential pitfalls, including degeneracies and systematics (§10.5.4). I argue that photodynam-

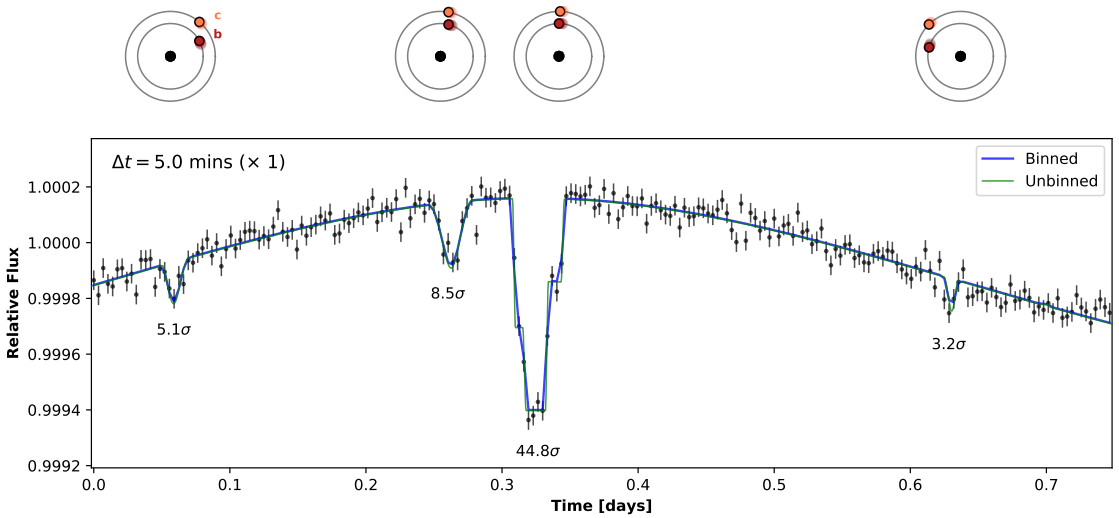

Figure 107: Similar to Figure 103, but for a single pointing of the Origins Space Telescope, assuming a 12m diameter mirror and an observation in a filter between 10-30 μm with a throughput of 30%. The exposure time is 5 minutes. All three planet-planet occultations of TRAPPIST-1c are detectable above the noise. The deepest one has SNR$\sim$8.5.

ical analysis with my code offers the best potential for detection of multiple PPO events (§10.5.5), which can account for mutual events (§10.5.6), as well as offset hot spots (§10.5.7) and tidal heating (§10.5.8). I end with a discussion of prior related work (§10.5.9) and the potential for application to other systems (§10.5.10), including non-transiting systems or white dwarfs.

10.5.1 PPO Mapping

Planet-planet occultations offer an opportunity to map the disk of the occulted planet (Ragozzine and Holman, 2010). High SNR, high cadence mid-IR observations of PPO light curves will, in principle, reveal the brightness temperature map, including any heterogeneities, along circular arcs having the radius of the occultor. While PPO mapping is similar to secondary eclipse mapping (e.g. Williams et al., 2006; Rauscher et al., 2007; Agol et al., 2010; Majeau et al., 2012; de Wit et al., 2012), the events will differ in orbital phase, duration, depth, and impact parameter. In principle, different events allow one to probe the surface of the planet along different circular arcs, permitting the deconvolution of the signal into separate latitudinal and longitudinal maps. Moreover, occultations at different phases could allow one to map the entire surface of the planet—not just the dayside, as with secondary eclipse mapping.

The procedure outlined above may likely not be possible with JWST, given the low SNR of individual PPO events using the nominal noise model (§10.4.3); future generation telescopes such as OST (§10.4.4) may be required. However, there are several diagnostics in PPO events arising from the spatial distribution of thermal emission that could be possible with JWST:

1. A larger day/night contrast will cause deeper events at full phase and shallower events at new phase.

2. A larger day/night contrast will cause a shift in the timing of events (see Figure 102).

3. A phase offset of the peak emission due to recirculation may also cause a shift in timing (see §10.5.7).

4. A latitudinal offset of the peak emission may cause a depth variation (see §10.5.7).

In principle, these effects could be degenerate with the orbital dynamics of the system. For example, if the eccentricity is such that the occultations at new phase have a smaller impact parameter, this will cause shallower depths. This may be degenerate with a larger day/night temperature contrast, but the shorter duration and the timing shift could perhaps break this potential degeneracy. Moreover, the detection of multiple PPOs will help constrain the orbital parameters and eventually allow precise determination of the impact parameter and timing of individual events, which will break these degeneracies.

In §10.4.2.2 I argued how the signature of a strong day/night temperature contrast corresponds to a $\sim$ 80 ppm signal if the timing of the PPO is known *a priori*. Incidentally, this is slightly above the noise level in the simulated stacked observation with JWST shown in Figure 104, suggesting that $\sim$10 PPOs of c are needed to detect the day/night asymmetry with SNR $>$ 1. A statistically significant detection would require on the order of 100 PPOs with JWST, which could in principle be achievable over the mission's lifetime. However, targeted observations of occultations of c at new phase would likely be a more efficient way of constraining the day/night asymmetry, as the (non)detection of PPOs could place strong (upper) limits on its nightside emission. Alternatively, just a few PPOs seen with OST could lead to a statistically significant detection.

10.5.2 PPOs versus phase curves

Thermal phase curves are another method used to map short-period planets by measuring the day/night temperature contrast of an exoplanet from its flux as a function of that planet's orbital phase about the star (Knutson et al., 2007; Selsis et al., 2011; Maurin et al., 2012). However, these observations will be challenging to acquire with JWST for several reasons, both practical and fundamental. The practical limits are that phase curves require a significant observing time investment, since the phase curve must be mapped over a significant fraction of the orbit. The high demand for JWST will likely make it difficult to obtain observing time for such long observations. Additionally, over longer timescales, instrumental systematic variations can be significant. For example, in the case of 55 Cancri e, the systematic variations were several hundred times larger than the claimed phase amplitude precision measured with *Spitzer* (Demory et al., 2016a).

The fundamental limits are that multiple phase curves can produce complex photometric behavior due to the different amplitudes and frequencies of each planet (Kane and Gelino, 2013), which when added to stellar variability of a periodic or random nature, may make it difficult to uniquely recover longitudinal maps of any of the planets. In contrast, PPO events are much shorter in duration, and can be scheduled in advance (within the dynamical uncertainties), using a much smaller fraction of time in which the star has less chance to vary, and systematic variations will likely be less severe.

Finally, phase curves of transiting planets are only sensitive to the longitudinal surface brightness of the planet (Cowan and Agol, 2008; Cowan et al., 2009), and act as a low-pass filter, while PPO events can probe latitudinal brightness, like secondary eclipse mapping. However, unlike secondary eclipse mapping, which can only occur at full phase and probe the dayside, PPOs can occur at multiple phases allowing the disk to be mapped for different illumination fractions.

10.5.3 Best case scenarios

Since detecting single PPOs with JWST will be challenging (§10.4.3), I consider in this section the best case scenarios for their detection and recommend various observing strategies.

In my detectability discussion above, I focused on occultations of TRAPPIST-1c by TRAPPIST-1b. In principle, occultations of *b by c* should be deeper events, given that b is exposed to $\sim$2 times the irradiation of c. However, dayside occultations of b by c are only possible when c is

on the opposite side of the star, closest to the observer, in which case the relative velocity of the two planets is at a maximum. These occultations are thus extremely short-lived and will not be detectable. In contrast, long occultations of the nightside of b by c *are* possible, when both planets are on the near side of the star, and could be detectable if b has a bright nightside due to strong atmospheric recirculation. By the same token, if b has a cold night side, the nondetection of these occultations could place strong constraints on its day/night temperature contrast. Given the higher irradiation of b relative to c, fewer occultations would need to be observed to obtain statistically significant results.

If the TRAPPIST-1 planets are airless bodies, occultations of c by b like the ones shown in Figure 103 will have the strongest signal. In my detectability calculations I assumed an albedo of 0.3 for all planets, but this could be an overestimate for airless bodies: both the Moon and Mercury have albedos of about 0.1. A lower albedo could slightly improve the detectability of these events.

If, on the other hand, TRAPPIST-1b and c have thick atmospheres, they may be analogous to Venus, whose extreme greenhouse forcing results in a surface that is much hotter than the planet's equilibrium temperature. In general, the surface of such a planet will be radiatively decoupled from the effective emitting layer, which is at a much lower temperature, and PPOs may be difficult to observe. This is an issue particularly for a CO_2-dominated atmosphere, which has strong absorption bands near 15 μm. Nevertheless, if the orbital parameters are well constrained, non-detections of PPOs at 15 μm could in principle be used to infer the presence of CO_2 or other strongly absorbing atmospheric species.

Alternatively, observations in atmospheric windows—such as the 2.4 μm K band window in Venus' atmosphere (Arney et al., 2014)—could allow one to probe to a much deeper (and hotter) layer, significantly enhancing the detectability of PPOs. I leave this investigation to future work.

Whatever the atmospheric regime of TRAPPIST-1c, there are two particular strategies for choosing when to observe occultations by TRAPPIST-1b. First, the triple event shown in Figure 103 is perhaps the single best observational scenario, as three occultations of TRAPPIST-1c plus a secondary eclipse of both TRAPPIST-1c and b occur over the span of about 15 hours, yielding a large amount of information for a single pointing of JWST. These events are relatively common, and happen on both sides of the star; on the near side, it is TRAPPIST-1b which is occulted, but the orbital geometry is otherwise identical by symmetry. Once the orbital parameters

of b and c are better constrained, these events can be predicted using my photodynamical code.

Second, from Figure 99 it is clear that occultations of TRAPPIST-1c are most common when the planet is at a mean longitude $\lambda_c \approx 43°$ or 137°. As I discussed previously, the peak in the histogram has FWHM $\sim 2°$, corresponding to a time window of just under 20 minutes. Thus, even if the eccentricities and longitudes of ascending node of TRAPPSIT-1b and c are not well known, the likelihood of observing an occultation of c is maximum for an observation centered at this mean longitude lasting 20 (or more) minutes. In order for the occultation to occur, TRAPPIST-1b must be near quadrature, which can easily be predicted in advance. Alternatively, and by symmetry, when TRAPPIST-1c is at a mean longitude $\lambda_c \approx 43°$ or 137° and b is near quadrature, occultations of b are most common (note the peaks near 0° and 180° in Figure 98).

Finally, I note that occultations of the other planets are likely not observable with JWST, but may be detectable with OST. Nevertheless, occultations of b and c *by* the outer planets are potentially observable and can be long-lived, especially occultations of c by d, but are more rare.

10.5.4 Degeneracies and other issues

There exist several degeneracies in PPO light curves that merit special attention, particularly for low SNR observations. I already discussed how a timing offset due to a day/night temperature contrast could be degenerate with the orbital parameters (§10.5.1). The planetary albedo is also degenerate with the impact parameter of the occultation, since both can change the depth of the event; in principle, at high SNR the duration can be used to resolve this.

If the orbital parameters are *very* poorly constrained, it may not be possible to tell with certainty which pairs of planets are occulting. This could be an issue if simultaneous or near-simultaneous occultations occurred between two different sets of planets. This is likely not going to be the case for TRAPPIST-1, since only occultations of b and c are at present potentially detectable and the orbital parameters are fairly well constrained. In principle, however, this could lead to interesting degeneracies between the orbital and surface properties across multiple planets in the system.

It is important to remember, however, that transits occur about 6% of the time in TRAPPIST-1 (as do secondary eclipses). Moreover, the star flares at a rate of about 0.26 day^{-1} (Luger et al., 2017). It is therefore quite likely that some PPOs will occur during one of these events, which will

negatively impact their detectability.

Another degeneracy that may arise concerns the potential heterogeneity of the surface. Spatial albedo variations due to clouds or geography could lead to asymmetries in the occultation light curve that could in principle be mistaken for a day/night temperature contrast. I defer a proper treatment of heterogenous surfaces to a future paper.

Finally, I note that in my detectability calculations I have neglected stellar variability, even though brightness modulations due to spots are clearly visible in the *K2* light curve of the system (Luger et al., 2017). However, while this may complicate phase curve observations, it is unlikely to significantly impact the detection of PPOs, given the rotation period of the star (3.3 days; Luger et al., 2017) is two orders of magnitude greater than the typical duration of a PPO. Moreover, spot contrasts and granulation noise are both much weaker in the mid-IR. The same is true for flares, whose emission peaks at shorter wavelengths. Nevertheless, unmodelled stellar variability could in principle decrease the SNR of PPO events and increase the number of observations needed to achieve a detection.

10.5.5 Joint modeling

Since I expect individual PPOs to be close to the noise level in JWST observations, joint modeling of multiple events must be performed to obtain constraints on the orbital and atmospheric properties of the planets. As `planetplanet` is a full photodynamical model, it may be used in a Markov Chain Monte Carlo (MCMC) simulation to derive posterior probabilities of the quantities of interest. I described in §10.4.3 why this approach is preferable to "stacking": while multiple occultations of a given planet result in different light curves, all of these light curves are functions of the same set of model parameters Θ. Instead of stacking to infer the occultation depth (which will always be different), joint modeling can be performed to marginalize over all model parameters and constrain, for instance, the day/night temperature contrast (which is the same across all occultations, if the planet has weak temporal variability). Given a dataset of measured fluxes D, the probability of a given value of the day/night temperature contrast ΔT is

$$p(\Delta T|D) = \int p(\Delta T, \Theta|D) d\Theta \tag{223}$$

where Θ is the set of all other parameters in the model (eccentricities, albedos, etc.), over which I have marginalized.

MCMC provides an efficient and robust way of computing the expression above. More generally, it yields the posterior probability distributions of all parameters conditioned on the data and all prior information about the system, such as the current constraints on the orbital parameters (Table 9). I provide sample code for performing full photodynamical modeling of TRAPPIST-1 with MCMC on the project `github` page. For more details on MCMC and its implementation, the reader is referred to Mackay (2003) and Foreman-Mackey et al. (2013).

10.5.6 Mutual transits

The extreme coplanarity of the TRAPPIST-1 system suggests that mutual transits should also occur in the system. These are planet-planet occultations occurring on the face of the star, which have been studied for *Kepler* systems (e.g., Hirano et al., 2012; Masuda et al., 2013; Masuda, 2014) and for moon-planet occultations (e.g., Kipping, 2011; Pál, 2012). The advantage of mutual transits over PPOs is that since the star is being occulted, the signal of the mutual event is much stronger, often a significant fraction of the transit depth. Typical mutual transits result in a brightening close to transit center, due to the fact that the total area occulted on the star is less than the sum of the areas of the two planets (see, for example, Figure 3 in Pál 2012). However, for a coplanar system like TRAPPIST-1, mutual transits are much more infrequent than planet-planet occultations. This is because the region of overlap of the orbital paths of any two planets on the face of the star is much smaller than the total region of overlap (see Figure 88). Moreover, as I discussed in §10.3.1.2, PPOs are a much more powerful way to constrain the three-dimensional orbital architecture of the system because of their longer lever arm and potential to occur at any orbital phase. For this reason, I do not model mutual transits of the TRAPPIST-1 planets in this work. I note, however, that my generalized integration scheme (§10.3.2.3) allows one to easily compute occultation light curves for an arbitrary number of overlapping bodies. Figure 108 shows a hypothetical double mutual transit computed with `planetplanet`, in which three planets occult each other as they transit their star. The brightening due to the mutual transits is evident near transit center.

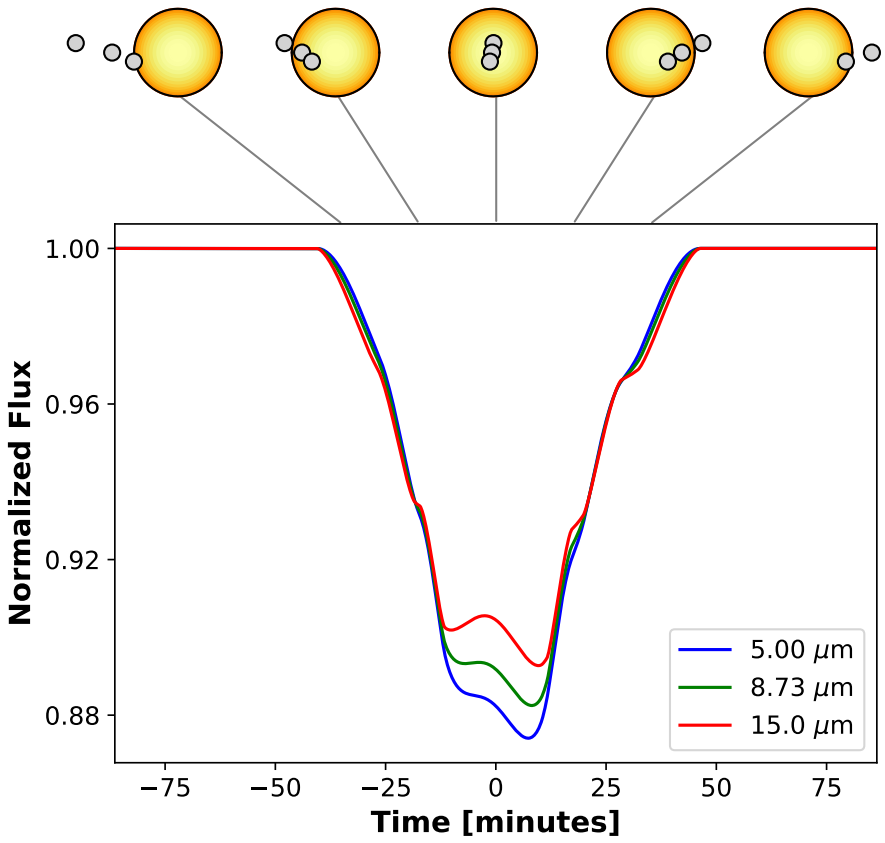

Figure 108: An example of a triple mutual transit across a star with an arbitrary limb darkening profile in a hypothetical system. The integration scheme in `planetplanet` is completely general and allows one to easily calculate occultation light curves for arbitrary numbers of overlapping bodies. In this example, I use a simple linear limb darkening law with a limb-to-center contrast that decreases linearly with wavelength, resulting in changes to the transit depth and shape at different wavelengths.

10.5.7 Hotspot offsets

Secondary eclipse measurements are capable of measuring latitudinal and longitudinal offsets in the position of peak emission of a planet due to atmospheric recirculation (Majeau et al., 2012; de Wit et al., 2012); planet-planet occultations are no different. The symmetry of the emission about the hotspot suggests I can use the same integration scheme discussed in §10.3.2.3 to model light curves for planets with shifted hotspots. `planetplanet` allows users to specify the latitudinal (Φ) and longitudinal (Λ) shifts when computing light curves. Below I derive the geometry of the general case of an "eyeball" planet with an arbitrary hotspot offset, and discuss how I adapt the

`planetplanet` integration scheme to handle these bodies in the same semi-analytic fashion. Note that, in general, the temperature distribution of an "eyeball" planet with an offset hotspot will not follow Equation (191), since that expression assumes the surface is in pure radiative equilibrium with the stellar flux. To this end, `planetplanet` allows users to input custom functions for $T(\phi)$.

Let us consider a frame in which the planet's orbit is in the xz plane and the planet is at full phase, so that its position vector is $(0, 0, r)$. In this frame, it is easy to show that the coordinates of the hotspot are

$$\begin{aligned} x' &= r_{\mathcal{P}} \cos\Phi \sin\Lambda \\ y' &= r_{\mathcal{P}} \sin\Phi \\ z' &= r_{\mathcal{P}} \cos\Phi \cos\Lambda + r. \end{aligned} \tag{224}$$

With some algebra, it can be shown that the hotspot coordinates in the sky plane are

$$\begin{aligned} x_\star &= \frac{zrx' - xyy' + xr_{xz}z'}{rr_{xz}} \\ y_\star &= \frac{r_{xz}y' + yz'}{r} \\ z_\star &= \frac{-xrx' - yzy' + zr_{xz}z'}{rr_{xz}}, \end{aligned} \tag{225}$$

where

$$r_{xz} = \sqrt{x^2 + z^2}. \tag{226}$$

Given the sky coordinates $(x^\star, y^\star, z^\star)$ of the hotspot, I may now use Equations (197) and (199) to obtain the phase angle θ and the rotation angle γ. As before, I rotate the problem by this angle to perform the integration in a frame in which the occulted body is symmetric about the x axis. Once again, `planetplanet` automatically handles all the transformations described above. An example of a planet with a hotspot offset is shown in Figure 109.

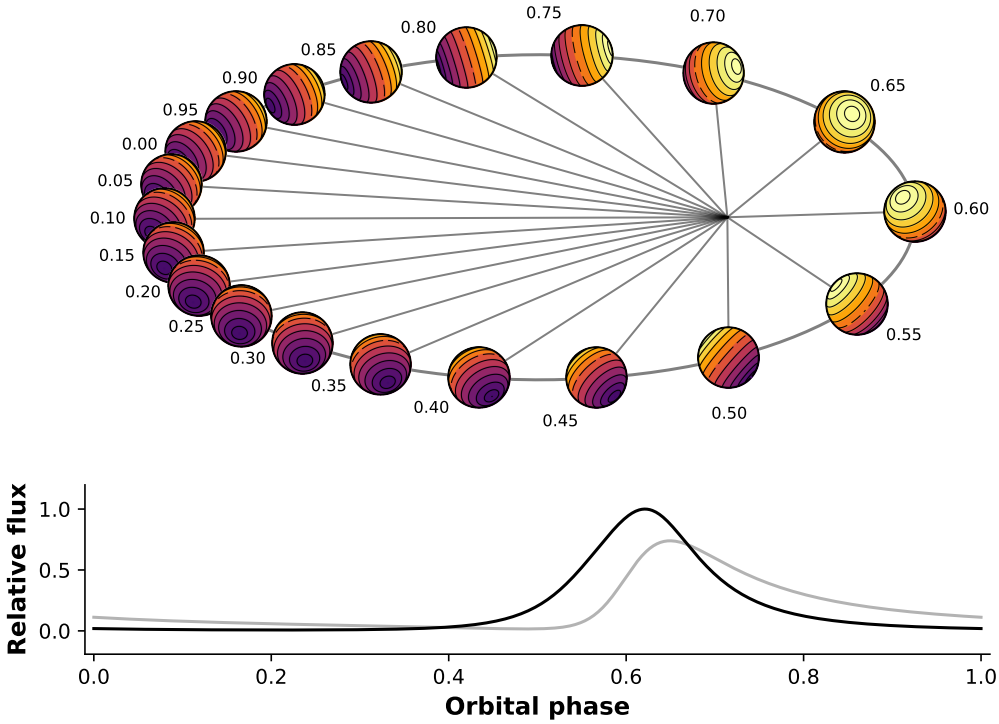

Figure 109: Similar to Figure 93, but for a planet with an offset hotspot. The orbital parameters are identical to those in the previous figure, but the hotspot has a latitudinal offset $\Phi = 30°$ (northward) and a longitudinal offset $\Lambda = 60°$ (eastward). The phase curve is shown as the black line in the lower panel; for reference, the phase curve for the default (no offset) case is shown in grey.

10.5.8 Tidal heating

My signal-to-noise considerations make the assumption that the planet temperatures are governed by the absorbed flux on each planet which is re-radiated in the thermal infrared. Another possibility is that planets may be tidally heated to the point that their thermal emission may be dominated by tidal heat flux (Selsis et al., 2013; Bolmont et al., 2013). However, on a short timescale this dissipation can cause the rotation of a planet to be synchronized (Jackson et al., 2008a) and the eccentricity to be damped, causing tidal forces to be diminished unless additional forcing maintains the planet in an asynchronous or eccentric state (Jackson et al., 2008b; Leconte et al., 2015; Arras and Socrates, 2010). If the eccentricity is maintained, tidal heating of short-period, eccentric exoplanets can compete with heating by the host star, even for modest eccentricities (Driscoll and Barnes, 2015).

Using Equation (2) from Driscoll and Barnes (2015), the ratio of the tidal heating flux to the

equilibrium flux is given by

$$\left(\frac{T_{tid}}{T_{eq}}\right)^4 = 1344\pi^3 \frac{|\mathrm{Im}(k_2)|}{1-A} \frac{M_\star}{L_\star} \frac{e^2 r_\mathcal{P}^3}{aP^3}, \tag{227}$$

where A is the albedo, T_{tid} is the effective temperature resulting from tidal flux, T_{eq} is the effective temperature resulting from absorbed flux, $\mathrm{Im}(k_2)$ is the imaginary component of the tidal Love number, $M_\star$ and $L_\star$ are the mass and luminosity of the star, and a, e, P, and $r_\mathcal{P}$ are the semi-major axis, eccentricity, orbital period, and radius of the planet, respectively. For TRAPPIST-1b, I find that

$$\left(\frac{T_{tid}}{T_{eq}}\right)^4 = 2.9 \times 10^6 |\mathrm{Im}(k_2)| \frac{e^2}{1-A}. \tag{228}$$

For $|\mathrm{Im}(k_2)| = 3 \times 10^{-3}$ and $A \ll 1$, these fluxes are equal for $e \approx 0.011$. However, this is about a factor of $\sim$10 larger than the eccentricity predicted by tidal evolution of the system (Luger et al., 2017). Nevertheless, if the eccentricity were this large, the tidal flux should give a higher signal-to-noise on PPO measurements. If strong enough, this could perhaps enable a constraint on $|\mathrm{Im}(k_2)|$ if A and e are known.

10.5.9 Comparison to other work

Previous studies have developed similar integration schemes to the one employed in `planetplanet` for computing various types of light curves. The `LUNA` code (Kipping, 2011) is a photodynamical algorithm for computing mutual transits of planet–moon systems across limb-darkened stars. `LUNA` computes mutual light curves by considering 34 different cases for the relative positions of the three disks (the star, the planet, and the moon) and deriving analytic expressions for the occulted stellar flux for each one. While extremely fast, `LUNA` accounts only for occultations of the stellar flux, and not for occultation of the flux from the planet and/or moon, and can therefore not be easily extended to the general case of planet-planet occultations.

`LUNA` is also limited to the three-body case and to the limb darkening laws that can be modeled analytically with the Mandel and Agol (2002) algorithm. Seeking to relax these assumptions and generalize the computation of mutual transits, Pál (2012) developed an algorithm to analytically compute light curves for arbitrary numbers of overlapping bodies transiting stars with various limb darkening profiles. This algorithm relies on a clever implementation of Green's theorem, which

allows one to analytically evaluate the surface integral of the stellar flux occulted by an arbitrary shape via computation of the line integral along the occultor boundary. In principle, the method is general enough that it can be extended to non-radially symmetric surface maps, as is the case for airless bodies or even bodies with arbitrary heterogeneous surface features. However, the algorithm requires finding (and integrating) a function whose exterior derivative is the surface brightness map, which can become arbitrarily complex for all but the simplest cases. Application of this algorithm to the airless body case would require finding the exterior derivative of Equation (192), with the temperature profile given by Equation (191) and the zenith angle given by Equations (212)–(214), which is intractable.

As I mentioned earlier, the `batman` code (Kreidberg, 2015) discretizes the radial brightness gradient of the star to compute fast transit light curves. `planetplanet` extends the `batman` algorithm to brightness gradients that are symmetric about *any* point on the surface of the body, not just the center of the disk. Because in the general case the curves of constant radiance are ellipses, `planetplanet` light curves are slower to evaluate (as they require solving a quartic equation to determine the points of intersection with the occultor; see §10.3.2.3). I therefore adopt the `batman` algorithm when modeling single-planet transits and occultations of planets whose emission is radially symmetric, making `planetplanet` fast for wavelength-dependent transit calculations.

10.5.10 Kepler-444 and other Systems

While I focused my discussion on TRAPPIST-1, planet-planet occultations likely occur in other compact, multi-planet systems and may be detectable in those. TRAPPIST-1 is the closest known edge-on, compact multi-planet system, which enhances its detectability, but because of the small sizes of its planets and their modest irradiation, there could exist other systems for which PPO detection is more favorable.

At 36 pc (three times more distant that TRAPPIST-1), Kepler-444 is a K dwarf hosting five edge-on, extremely coplanar sub-Earths with periods less than 10 days (Campante et al., 2015; Mills and Fabrycky, 2017). The Kepler-444 planets receive 50–100 times higher irradiation than the TRAPPIST-1 planets and, with higher equilibrium temperatures and stronger thermal emission, could in principle produce detectable PPOs. However, despite the fact that the Kepler-444 planets emit more flux than the TRAPPIST-1 planets, the larger luminosity of the star and greater distance

to the system strongly decrease the occultation signal and increase the noise on PPO observations. I find that the $\sim$7.7 μm JWST/MIRI filter (F770W) would be the optimal filter for secondary eclipse and PPO observations of Kepler-444b, but on the order of 100 events would be required to build up a SNR$\sim$3, which makes their detection infeasible. On the other hand, only $\sim$15 secondary eclipses/occultations of Kepler-444b would be required for a SNR$\sim$3 detection using my nominal OST setup from §10.4.4 with the same F770W filter.

Alternatively, surveys such as TESS and PLATO may soon discover nearby systems of multiply-transiting super-Earths and mini-Neptunes, which could have stronger PPO signals than TRAPPIST-1. In particular, PPO searches should focus on late-M or brown dwarf multi-planet systems with large, short-period planets that are equally distant or closer than TRAPPIST-1.

10.5.11 Other applications

In principle, PPOs could be used to search for non-transiting planets. This was pointed out by Ragozzine and Holman (2010), with the caveat that the rare and aperiodic nature of these events would make detections based on PPOs alone extremely challenging. However, there may exist systems of extremely coplanar planets that are inclined just enough so that they do not transit. Around main sequence stars, even small system inclinations will in general greatly reduce the frequency of PPOs, since the minimum projected separation of coplanar planets increases steeply with inclination. Compact planetary systems orbiting white dwarfs or other compact objects, on the other hand, can still be extremely close to edge-on and not transit. If planetary systems orbiting these stars are common, PPOs may be a way to detect and characterize them.

This last point merits further discussion. Not only may PPOs be ideally suited to detecting planets around white dwarfs, but they should also be *easier* to detect for these systems than for main-sequence stars (should they exist). In the Rayleigh-Jeans limit, the depth of a PPO scales as $T_p r_\mathcal{P}^2/(L_\star^{1/4} R_\star^{3/2})$, where T_p and $r_\mathcal{P}$ are the temperature and radius of the planet and $L_\star$ and $R_\star$ are the luminosity and radius of the star, respectively. For fixed planet properties and stellar luminosity, a planet orbiting a smaller star undergoes deeper occultations. In fact, if TRAPPIST-1 were a white dwarf with the same luminosity, the depth of PPO events in the system would be $\sim$30 times larger because of the relative faintness of the star in the infrared.

Finally, I only explored PPO detectability from space-based telescopes. It may be possible in

the future to consider the observation of PPOs with ground-based telescopes. This is challenging due to the small PPO depth, limited visibility from any given site, high sky-brightness in the mid-infrared, and limited sky transmittance. However, future extremely large telescopes (ELTs) will have a smaller PSF, decreasing the contribution from sky brightness to the noise, and a large collecting area, leading to smaller photon noise. In practice, though, variable sky brightness and variable seeing will likely limit the precision of ground-based ELTs, just as it limits the precision of transit and secondary eclipse measurements. The transmittance in the 20 μm Q band requires high and/or dry sites for observation, as well as limited cloud emissivity, all leading to a small probability of success. Nevertheless, the possibility of ground-based observations of PPOs may be worth examining in the future.

10.6 Conclusions and future directions

I have developed a formalism to predict, model, and extract orbital and atmospheric information from planet-planet occultations (PPOs) in extrasolar multi-planet systems. These events occur when a planet occults another planet in the same system, blocking thermal or reflected light from the occulted body and producing a photometric signal analagous to (and of similar depth to) a secondary eclipse. While PPOs are in general rare events, for extremely coplanar, edge-on, and compact systems such as TRAPPIST-1 they are frequent and occur preferentially at certain orbital phases, which may in some cases permit advanced scheduling of observations.

PPOs are a powerful tool to assess the day/night temperature contrast of exoplanets. Unlike secondary eclipse, which always probes a planet's dayside emission, PPOs can occur at any phase and thus can yield information about the entire surface of the planet, enabling crude two-dimensional surface maps that are potentially not degenerate in latitude. Moreover, PPOs yield strong constraints on the three-dimensional architecture of exoplanetary systems. Observations of multiple events can yield the full eccentricity vectors of pairs of planets, breaking both the eccentricity-eccentricity and eccentricity-mass degeneracy inherent in transit timing variation (TTV) analyses. Because they occur off the face of the star, PPOs also constrain the mutual inclinations of pairs of planets precisely.

At this time, the mutual inclinations of the TRAPPIST-1 planets are unconstrained. I developed a Monte Carlo framework to assess their values given current data on the system, finding

that the scatter in their mutual inclinations is less than 0.3° with 90% confidence, suggesting the TRAPPIST-1 system is *extremely* coplanar. I marginalize over the uncertainty on the mutual inclinations and all other orbital parameters to determine the frequency of PPOs in TRAPPIST-1. I find that occultations among the TRAPPIST-1 planets occur, on average, about once per (Earth) day. Most of these are grazing, short-lived occultations. Neglecting PPOs with impact parameter $b > 0.5$ and durations less than 10 minutes, I find that TRAPPIST-1b is occulted 21.1 ± 9.7 times per (Earth) year and TRAPPIST-1c is occulted 21.3 ± 8.8 times per (Earth) year. The outer planets are occulted less frequently, but even TRAPPIST-1h likely undergoes at least two occultations per year.

In order to model PPO light curves, I developed the open-source code `planetplanet`, a photodynamical software that uses an N-body solver to calculate the sky-projected separations of all pairs of planets and computes wavelength-dependent light curves for occultation events. I model planet thermal emission in two limits: the thick atmosphere limit, in which the planet is assumed to have symmetric emission about the center of its disk, and the thin atmosphere limit, in which its thermal emission is proportional to the stellar irradiation profile. I developed a novel integration scheme that takes advantage of the elliptical symmetry of the problem to compute fast PPO light curves and use it to also model transits, secondary eclipses, phase curves, and mutual transits.

I modelled the detectability of PPOs with future space-based telescopes. I find that observations with JWST/MIRI at 12.8 and/or 15 μm are best suited to study PPOs in TRAPPIST-1, though individual events will likely be close to the noise level. I developed a framework to jointly analyze multiple PPO events to extract information about the day/night temperature contrast of the TRAPPIST-1 planets and their orbital configuration. For a future far-infrared surveyor such as the Origins Space Telescope (OST), I find that single PPOs are detectable in TRAPPIST-1 at $> 8\sigma$. Future observations of these events may therefore paint the most complete picture yet of the three-dimensional architecture of an exoplanet system and of the surfaces of its planets.

PPOs likely occur in other compact multi-planet systems, such as Kepler-444; TESS and PLATO may reveal other nearby systems that can be explored with JWST using this technique. Finally, I suggest that PPOs may be a way to detect near-edge-on systems of planets around white dwarfs and other compact objects, given the low geometric probability of transit in these systems.

All code used to generate the figures in this paper is open source and available at `https:`

//github.com/rodluger/planetplanet. Complete documentation is available at the url listed above.

Table 9: Orbital parameters assumed for the TRAPPIST-1 planets

Property	b	c	d	e	f	g	h
P [days]	1.5108708 ± 0.0000006	2.4218233 ± 0.0000017	4.049610 ± 0.000063	6.099615 ± 0.000011	9.206690 ± 0.000015	12.35294 ± 0.00012	18.767 ± 0.004
t_0 [BJD $-$ 2,450,000]	7322.51736 ± 0.00010	7282.80728 ± 0.00019	7670.14165 ± 0.00035	7660.37859 ± 0.00038	7671.39767 ± 0.00023	7665.34937 ± 0.00021	7662.55284 ± 0.00037
I [deg]	89.65 ± 0.245	89.67 ± 0.17	89.75 ± 0.16	89.86 ± 0.11	89.680 ± 0.034	89.710 ± 0.025	89.80 ± 0.075
e	0.0005 ± 0.0001	0.004 ± 0.001	0.0004 ± 0.0003	0.007 ± 0.0005	0.009 ± 0.001	0.004 ± 0.001	0.003 ± 0.001
ω [deg]	Unif$[0, 2\pi)$	Unif$[0, 2\pi)$	Unif$[0, 2\pi)$	Unif$[0, 2\pi)$	Unif$[0, 2\pi)$	Unif$[0, 2\pi)$	Unif$[0, 2\pi)$
Ω [deg]	MC	MC	MC	MC	MC	MC	MC
m [$M_\oplus$]	0.85 ± 0.72	1.38 ± 0.61	0.41 ± 0.27	0.62 ± 0.58	0.68 ± 0.18	1.34 ± 0.88	0.4 ± 0.3
Transit depth [%]	0.7266 ± 0.0088	0.687 ± 0.010	0.367 ± 0.017	0.519 ± 0.026	0.673 ± 0.023	0.782 ± 0.027	0.346 ± 0.032

Note. — Periods, transit times, inclinations, masses, and transit depths from Gillon et al. (2017) and Luger et al. (2017). The mass of h is drawn from a distribution based on that of d, whose radius is the most similar to that of h. Eccentricities based on the tidal/migration simulations in Luger et al. (2017). Longitudes of ascending node drawn from the distribution derived by the Monte Carlo method described in §10.4.1.1.

Conclusions

Conclusions

Over the past three decades, the number of known exoplanets has grown exponentially with a doubling time of about 27 months, a trend informally known as "Mamajek's Law," after Eric Mamajek, who first recognized the trend (see Figure 110). With the advent of transiting exoplanet "factories" such as TESS and PLATO, and with improvements to ground-based radial velocity instruments, this growth rate shows no signs of stopping. If the current rate continues, by 2034 we will have discovered one million exoplanets, and one *billion* by 2057. These figures are of course improbably high, but they capture the notion that we truly live in an amazing time for exoplanet discovery.

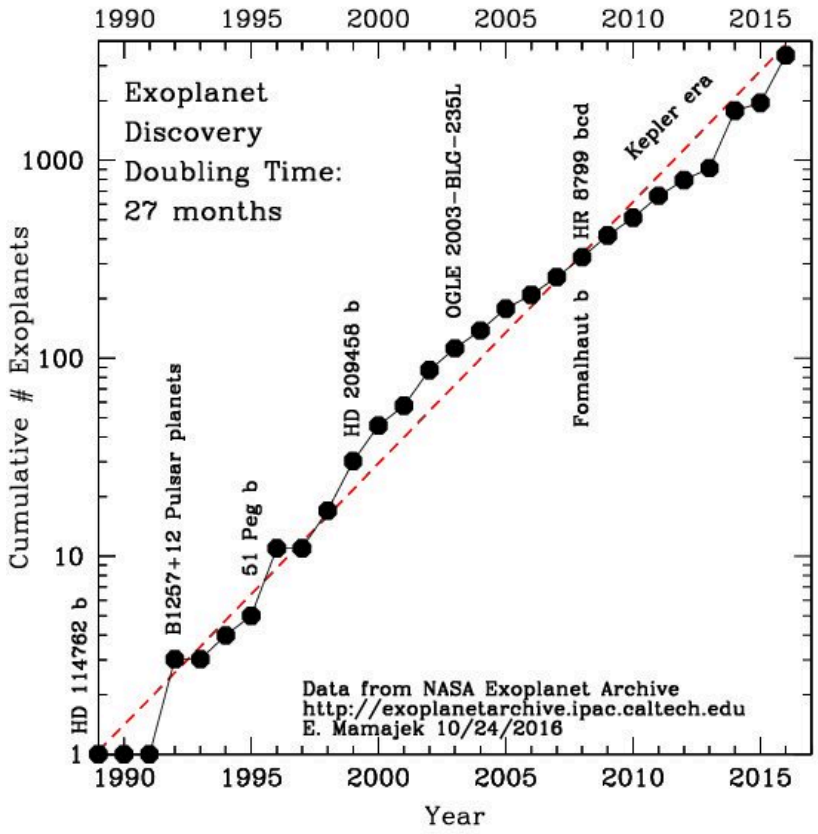

Figure 110: Number of known exoplanets as a function of time. This number has grown exponentially since 1990, with a doubling time of approximately 27 months. *Credit:* Eric Mamajek. Reproduced with permission.

As our ability to detect planets improves, so will our ability to characterize and understand

them. Once dismissed as mere "stamp collecting," (Jones, 2016) exoplanet science is anything but that. We are now at the point where detailed characterization of exoplanet atmospheres is possible and the search for biosignatures in the atmospheres of planets in the habitable zone is within reach. Direct imaging (Guyon and Martinache, 2013) and surface mapping (Cowan and Fujii, 2017) of these planets are just around the corner, and soon we will be able to place the Earth and its habitability in a galactic context.

In the meantime, however, there is much work to be done. If we wish to understand the habitability of exoplanets, target selection for observations with JWST and future planet characterization telescopes must prioritize those planets (and systems of planets) that (1) offer the best potential for habitability according to current data, priors, and models; and (2) are the most amenable to characterization using the tools at our disposal. In this dissertation, I have made progress on both fronts, identifying evolutionary pathways that may hinder or help the habitability of M dwarf planets and developing new techniques to detect and characterize these planets.

In Chapters 1–3, I showed how the prolonged pre-main sequence phase of M dwarfs can lead to vigorous atmospheric escape from planets in the habitable zone. I argued in Chapter 2 that for water-bearing terrestrial planets, this can lead to the loss of several Earth oceans of water from their surfaces, potentially causing their complete desiccation. The escape process can further lead to severe oxidation of the surface of these planets and the possible accumulation of oxygen in their atmospheres, which could at the same time hinder their habitability *and* fool us into believing they are inhabited via the spectroscopic detection of O_2 in those planets' atmospheres. This last point merits some discussion, as recent work has called into question the possibility of atmospheric buildup of O_2 by this mechanism. Zahnle and Catling (2017) argued that photochemical processes could act to ensure that hydrogen and oxygen escape in stoichiometric proportions, preventing the buildup of any O_2 in planetary atmospheres as a result of hydrodynamic escape. My model did not account for photochemistry beyond the photolysis of H_2O by stellar XUV irradiation, so my O_2 buildup rates should be viewed as strict upper limits. Recently, Bolmont et al. (2017) showed that water loss from the TRAPPIST-1 planets may be photolysis-limited, which could further reduce both water loss and O_2 accumulation rates. Moreover, the buildup of ozone and the presence of other atmospheric species could in principle also reduce the escape rate by shielding the atmosphere from XUV photons. A detailed hydrodynamical/photochemical treatment of the

escape process is therefore critical to determining the actual water loss rates and understanding the fate of the oxygen. As I argued previously, if oxygen does remain in the atmosphere following escape, its uptake by a surface magma ocean would be efficient; studies such as that of Hamano et al. (2013) assume that *all* atmospheric O_2 can be absorbed by the magma. Recently, Schaefer et al. (2016) used a detailed geophysical model to calculate O_2 absorption by a magma ocean on the hot super-Earth GJ 1132b. The authors found that a tenuous atmosphere with a few bars of O_2 remained in most of their simulations, with massive, several-hundred bars possible in some cases. It is important to note that GJ 1132b is interior to the habitable zone, meaning its runaway greenhouse/magma ocean phase could have lasted significantly longer than that of the habitable zone planets I consider here. With a shorter-lived magma ocean phase, one might expect planets in the HZ of M dwarfs to build up more atmospheric O_2 than GJ 1132b, provided photochemical processes allow O_2 retention in the first place.

In Chapter 3, I focused on a diametrically opposite consequence of vigorous hydrodynamic escape: I showed how the escape process could potentially help gas-rich super-Earths to shed their hydrogen/helium envelopes, resulting in a novel pathway to the habitability of planets around M dwarfs. I posited that super-Earths migrating from beyond the snow line could have abundant water and sufficient hydrogen/helium to prevent that water from escaping to space. Under certain conditions, this could lead to the formation of "habitable evaporated cores," which are likely very different from Earth but could in principle support life today. Recent work has also challenged this claim; Owen and Mohanty (2016) argued that a transition from hydrodynamic escape to Jeans escape after ~ 100 Myr effectively shuts down the runaway escape of hydrogen, preventing the formation of evaporated cores in many cases. Moreover, they argued that the default escape efficiency I assumed ($\epsilon_{\rm XUV} = 0.3$) is unrealistically high, a fact that finds support in statistical studies of evaporation among the *Kepler* planets (Lopez and Fortney, 2013; Owen and Wu, 2017), which suggest $\epsilon_{\rm XUV} \approx 0.1$ is more realistic. In their study, Owen and Mohanty (2016) found that habitable evaporated cores cannot form from planets with cores more massive than ~ 0.9 $M_\oplus$ or with H/He envelopes with mass fraction greater than about one percent. This is in contrast to what I argued in Luger et al. (2015b) and Chapter 3, that cores less massive than ~ 2 $M_\oplus$ with *any* hydrogen fraction could be evaporated. While some of my assumptions were, in fact, unrealistic, there are two key points that must be considered when comparing the two works. First, Owen and

Mohanty (2016) did not consider Roche lobe overflow (RLO), which for planets as low in mass as 1–2 $M_\oplus$ can strip the envelope down to $\sim 1\%$ in mass (or less; Luger et al., 2015b). However, as I argued in Chapter 3, it is unlikely that such small planets form with such large envelopes in the first place. Discounting escape due to RLO, my model agrees with that of Owen and Mohanty (2016) that envelopes more massive than about $\sim 1\%$ cannot be lost under most circumstances. Second, Owen and Mohanty (2016) considered the escape process for a single fiducial M3.5 dwarf of mass 0.4 $M_\odot$, for which the XUV flux in the HZ drops below the level required for hydrodynamic escape in less than ~ 100 Myr. For lower mass M dwarfs, however, the XUV flux likely remains elevated above this cutoff for several more hundred Myr, so the results of Owen and Mohanty (2016) are *underestimates* for late-type M dwarfs. After a few hundred Myr, the escape rates are negligible in my model, so the fact that I do not model the transition from hydrodynamic to Jeans escape does not significantly affect the escape rates for the lowest mass M dwarfs, which include Proxima Centauri and TRAPPIST-1. For these stars, the actual core mass below which habitable evaporated cores can form is therefore likely to be between my estimate ($\sim 2M_\oplus$) and that of Owen and Mohanty (2016) ($\sim 0.9M_\oplus$). On a related note, recent observational work has identified the evaporation of mini-Neptunes as a key process in explaining the population of planets observed by *Kepler*. Fulton et al. (2017) found a significant gap in the radius distribution of exoplanets at $\sim 1.8R_\oplus$, which is readily explained by hydrodynamic evaporation (Owen and Wu, 2017; Lehmer and Catling, 2017). These studies focus on hot mini-Neptunes and super-Earths around solar-type stars, but suggest that a similar process could occur for M dwarf planets in the HZ, which experience similar XUV fluxes at early times.

In Chapters 4–7, I discussed my efforts to improve the detection of small transiting planets, with application to planets in the habitable zone. In Chapters 5 and 6, I introduced `EVEREST`, a pipeline for de-trending data collected by the *K2* mission. *K2* is the second photometric mission of the *Kepler* telescope, whose primary goal is to detect transiting planets, with particular focus on those orbiting nearby, low mass stars. However, because of the two failed reaction wheels onboard *Kepler*, raw *K2* data lacks the photometric precision to detect the small transit signals of these planets. `EVEREST` uses machine learning techniques to remove instrumental noise from *K2* light curves, restoring the photometric precision of *Kepler* for the majority of stars observed by the mission. In Chapter 7, I applied `EVEREST` to TRAPPIST-1, a nearby ultracool dwarf hosting seven

transiting planets, three of which are in the habitable zone. I used `EVEREST` to detect four additional transits of TRAPPIST-1h, which had previously only been seen to transit once. I constrained its orbital period and used it to study the complex resonant structure of the planetary system.

Finally, in Chapters 8–10, I introduced two novel methods to detect and characterize exoplanets, particularly those in the habitable zone. Chapter 9 discussed the detection of auroral emission from planets orbiting nearby M dwarfs, with particular focus on Proxima Centauri b, a potentially terrestrial planet in the habitable zone of the Sun's closest neighbor. While auroral emission from Earth is relatively weak, emission from Proxima Centauri b could be 100 to 10,000 times stronger if that planet possessed an atmosphere and a magnetic field. This is due to the extreme magnetic activity and vigorous particle radiation of its host star, which could generate $\gtrsim 1$ terawatt auroral emission from the planet. I argue that high resolution spectroscopy with next-generation telescopes could detect it, and observations at different orbital phases could be used to constrain the three-dimensional orbit of the planet and derive an exact planet mass via the Doppler effect. Line strength calculations could further be used to learn about the planet's atmospheric composition and potentially the presence of a magnetic field.

Chapter 10 discussed the detectability of planet-planet occultations (PPOs), arguing for their use as a way to probe the three-dimensional structure of multi-planet systems and to generate crude surface maps of short-period planets. During a PPO, a planet occults another planet in the same system, an event that is typically rare, short-lived, and orders of magnitude below current detection thresholds. However, nearby systems of multiple short-period, extremely coplanar exoplanets such as TRAPPIST-1 are a game changer in this respect. I showed how planet-planet occultations in this system are frequent, long-lived, and potentially detectable by JWST and future space-based telescopes. PPOs thus open a new door to the characterization of nearby exoplanets, including those in the habitable zone.

Though I did not know it until recently, in many ways this dissertation has been about TRAPPIST-1 all along. The planets in TRAPPIST-1 likely formed quickly (Raymond et al., 2007) and migrated from far out to their current close-in orbits (Chapter 7), suggesting they may have accreted hydrogen and helium from the protoplanetary disk in the form of substantial envelopes. Given their low masses and likely history of extreme XUV irradiation (Chapter 1), it is unlikely they retain any today, a statement that is beginning to find support in observations

(e.g., de Wit et al., 2016). Moreover, transit timing variation studies suggest that all planets but TRAPPIST-1c are significantly less dense than Earth (Gillon et al., 2017; Wang et al., 2017), which would suggest they may have a significant water or ice mass fraction. The TRAPPIST-1 planets are therefore the best examples to date of potential habitable evaporated cores (Chapter 3). Planets b–g could have also experienced vigorous water loss. While they may not have been completely desiccated (assuming they are in fact extremely volatile-rich), substantial amounts of oxygen may have accumulated in their atmospheres (Chapter 2), which could be detected with upcoming JWST observations. My efforts to de-trend *K2* light curves (Chapters 4–6) culminated in the application of `EVEREST` to TRAPPIST-1, revealing the period of planet h, the rich resonant structure of the system, and information about the rotation rate and activity level of the host star. Recently, I developed the framework to extract even more information from the system with planet-planet occultations (Chapter 10). Finally, while the detection of exoauroral emission (Chapter 9) will be challenging for systems farther than a few parsecs, one day we may be able to probe the aurorae of the inner TRAPPIST-1 planets.

TRAPPIST-1 is indeed the gift that keeps on giving. The richness of information that can be extracted from TRAPPIST-1, both photometric and spectroscopically, is unprecedented. I will continue to develop techniques to characterize the TRAPPIST-1 planets and other systems like it, which are likely to be discovered in the coming years. In particular, the dynamics of PPOs merit a much more rigorous treatment, given their potential to capture the full three-dimensional architecture of a multi-planet system. The surface mapping potential of PPOs and their unique ability to probe the night side of planets is also especially intriguing, and I will continue to investigate it and to plan observations with JWST and future telescopes. On the de-trending front, I aim to adapt `EVEREST` to TESS and future photometric missions. In principle, the `EVEREST` algorithm could also be applied to spectroscopy, particularly when the name of the game is the detection of signals at the few ppm level. I will develop similar tools for de-trending JWST spectra with the aim to maximize our ability to probe the atmospheres of nearby potentially habitable planets. Finally, I will also work towards developing more realistic codes to model these planets' atmospheric evolution, as this remains one of the main hurdles for their habitability. I will build on the results of Chapters 2 and 3 with numerical hydrodynamic, photochemical, and geophysical codes to model the escape of hydrogen and the fate of the oxygen in these planets' atmospheres.

References

Y. Abe, A. Abe-Ouchi, N. H. Sleep, and K. J. Zahnle. Habitable Zone Limits for Dry Planets. *Astrobiology*, 11:443–460, June 2011. doi: 10.1089/ast.2010.0545.

M. H. Acuña, J. E. P. Connerney, P. Wasilewski, R. P. Lin, D. Mitchell, K. A. Anderson, C. W. Carlson, J. McFadden, H. Rème, C. Mazelle, et al. Magnetic field of mars: Summary of results from the aerobraking and mapping orbits. *Journal of Geophysical Research: Planets*, 106(E10): 23403–23417, 2001.

E. Agol. Rounding up the wanderers: optimizing coronagraphic searches for extrasolar planets. *MNRAS*, 374:1271–1289, February 2007. doi: 10.1111/j.1365-2966.2006.11232.x.

E. Agol, N. B. Cowan, H. A. Knutson, D. Deming, J. H. Steffen, G. W. Henry, and D. Charbonneau. The Climate of HD 189733b from Fourteen Transits and Eclipses Measured by Spitzer. *ApJ*, 721:1861–1877, October 2010. doi: 10.1088/0004-637X/721/2/1861.

S. Aigrain, S. T. Hodgkin, M. J. Irwin, J. R. Lewis, and S. J. Roberts. Precise time series photometry for the Kepler-2.0 mission. *MNRAS*, 447:2880–2893, March 2015. doi: 10.1093/mnras/stu2638.

S. Aigrain, H. Parviainen, and B. J. S. Pope. K2SC: flexible systematics correction and detrending of K2 light curves using Gaussian process regression. *MNRAS*, 459:2408–2419, July 2016. doi: 10.1093/mnras/stw706.

S.-I. Akasofu. Energy coupling between the solar wind and the magnetosphere. *Space Science Reviews*, 28(2):121–190, 1981.

S. C. Albers. Mutual Occultations of Planets: 1557 to 2230. *Sky & Telescope*, 57, March 1979.

M. E. Alexander. The Weak Friction Approximation and Tidal Evolution in Close Binary Systems. *Ap&SS*, 23:459–510, August 1973. doi: 10.1007/BF00645172.

Y. Alibert. On the radius of habitable planets. *A&A*, 561:A41, January 2014. doi: 10.1051/0004-6361/201322293.

S. Ambikasaran, D. Foreman-Mackey, L. Greengard, D. W. Hogg, and M. ONeil. Fast direct methods for gaussian processes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(2):252–265, Feb 2016. ISSN 0162-8828. doi: 10.1109/TPAMI.2015.2448083.

A. D. Anbar, Y. Duan, T. W. Lyons, G. L. Arnold, B. Kendall, R. A. Creaser, A. J. Kaufman, G. W. Gordon, C. Scott, J. Garvin, and R. Buick. A Whiff of Oxygen Before the Great Oxidation Event? *Science*, 317:1903–, September 2007. doi: 10.1126/science.1140325.

G. Anglada-Escudé, M. Tuomi, E. Gerlach, R. Barnes, R. Heller, J. S. Jenkins, S. Wende, S. S. Vogt, R. P. Butler, A. Reiners, and H. R. A. Jones. A dynamically-packed planetary system

around GJ 667C with three super-Earths in its habitable zone. *A&A*, 556:A126, August 2013. doi: 10.1051/0004-6361/201321331.

G. Anglada-Escudé, P. J. Amado, J. Barnes, Z. M. Berdiñas, R. P. Butler, G. A. L. Coleman, I. de La Cueva, S. Dreizler, M. Endl, B. Giesers, S. V. Jeffers, J. S. Jenkins, H. R. A. Jones, M. Kiraga, M. Kürster, M. J. López-González, C. J. Marvin, N. Morales, J. Morin, R. P. Nelson, J. L. Ortiz, A. Ofir, S.-J. Paardekooper, A. Reiners, E. Rodríguez, C. Rodríguez-López, L. F. Sarmiento, J. P. Strachan, Y. Tsapras, M. Tuomi, and M. Zechmeister. A terrestrial planet candidate in a temperate orbit around Proxima Centauri. *Nature*, 536:437–440, August 2016. doi: 10.1038/nature19106.

D. J. Armstrong, J. Kirk, K. W. F. Lam, J. McCormac, S. R. Walker, D. J. A. Brown, H. P. Osborn, D. L. Pollacco, and J. Spake. K2 Variable Catalogue: Variable stars and eclipsing binaries in K2 campaigns 1 and 0. *A&A*, 579:A19, July 2015. doi: 10.1051/0004-6361/201525889.

D. J. Armstrong, J. Kirk, K. W. F. Lam, J. McCormac, H. P. Osborn, J. Spake, S. Walker, D. J. A. Brown, M. H. Kristiansen, D. Pollacco, R. West, and P. J. Wheatley. K2 variable catalogue - II. Machine learning classification of variable stars and eclipsing binaries in K2 fields 0-4. *MNRAS*, 456:2260–2272, February 2016. doi: 10.1093/mnras/stv2836.

G. Arney, V. Meadows, D. Crisp, S. J. Schmidt, J. Bailey, and T. Robinson. Spatially resolved measurements of H_2O, HCl, CO, OCS, SO_2, cloud opacity, and acid concentration in the Venus near-infrared spectral windows. *Journal of Geophysical Research (Planets)*, 119:1860–1891, August 2014. doi: 10.1002/2014JE004662.

P. Arras and A. Socrates. Thermal tides in fluid extrasolar planets. *The Astrophysical Journal*, 714(1):1–12, apr 2010. doi: 10.1088/0004-637x/714/1/1. URL `https://doi.org/10.1088/0004-637x/714/1/1`.

M. Aumer and J. J. Binney. Kinematics and history of the solar neighbourhood revisited. *MNRAS*, 397:1286–1301, August 2009. doi: 10.1111/j.1365-2966.2009.15053.x.

S. Ballard, D. Charbonneau, F. Fressin, G. Torres, J. Irwin, J.-M. Desert, E. Newton, A. W. Mann, D. R. Ciardi, J. R. Crepp, C. E. Henze, S. T. Bryson, S. B. Howell, E. P. Horch, M. E. Everett, and A. Shporer. Exoplanet Characterization by Proxy: A Transiting 2.15 $R_\oplus$ Planet near the Habitable Zone of the Late K Dwarf Kepler-61. *ApJ*, 773:98, August 2013. doi: 10.1088/0004-637X/773/2/98.

I. Baraffe, G. Chabrier, F. Allard, and P. H. Hauschildt. Evolutionary models for solar metallicity low-mass stars: mass-magnitude relationships and color-magnitude diagrams. *A&A*, 337:403–412, September 1998.

R. Barnes and R. Heller. Habitable Planets Around White and Brown Dwarfs: The Perils of a Cooling Primary. *Astrobiology*, 13:279–291, March 2013. doi: 10.1089/ast.2012.0867.

R. Barnes, S. N. Raymond, B. Jackson, and R. Greenberg. Tides and the Evolution of Planetary Habitability. *Astrobiology*, 8:557–568, June 2008. doi: 10.1089/ast.2007.0204.

R. Barnes, K. Mullins, C. Goldblatt, V. S. Meadows, J. F. Kasting, and R. Heller. Tidal Venuses: Triggering a Climate Catastrophe via Tidal Heating. *Astrobiology*, 13:225–250, March 2013. doi: 10.1089/ast.2012.0851.

R. Barnes, R. Deitrick, R. Luger, P. E. Driscoll, T. R. Quinn, D. P. Fleming, B. Guyer, D. V. McDonald, V. S. Meadows, G. Arney, D. Crisp, S. D. Domagal-Goldman, A. Lincowski, J. Lustig-Yaeger, and E. Schwieterman. The Habitability of Proxima Centauri b I: Evolutionary Scenarios. *ArXiv e-prints, arXiv:1608.06919*, August 2016.

O. Barragán, D. Gandolfi, A. M. S. Smith, H. J. Deeg, M. C. V. Fridlund, C. M. Persson, P. Donati, M. Endl, S. Csizmadia, S. Grziwa, D. Nespral, A. P. Hatzes, W. D. Cochran, L. Fossati, S. S. Brems, J. Cabrera, F. Cusano, P. Eigmüller, C. Eiroa, A. Erikson, E. Guenther, J. Korth, L. Mancini, M. Pätzold, J. Prieto-Arranz, H. Rauer, I. Rebollido, J. Saario, and O. V. Zakhozhay. EPIC 218916923 b: a low-mass warm Jupiter on a 29-day orbit transiting an active K0 V star. *ArXiv e-prints*, February 2017.

J. K. Barstow and P. G. J. Irwin. Habitable worlds with JWST: transit spectroscopy of the TRAPPIST-1 system? *Monthly Notices of the Royal Astronomical Society*, 461:L92–L96, September 2016. doi: 10.1093/mnrasl/slw109.

M. Barthelemy and G. Cessateur. Sensitivity of upper atmospheric emissions calculations to solar/stellar uv flux. *Journal of Space Weather and Space Climate*, 4:A35, 2014.

C. Baruteau, A. Crida, S.-J. Paardekooper, F. Masset, J. Guilet, B. Bitsch, R. Nelson, W. Kley, and J. Papaloizou. Planet-Disk Interactions and Early Evolution of Planetary Systems. *Protostars and Planets VI*, pages 667–689, 2014. doi: 10.2458/azu_uapress_9780816531240-ch029.

T. S. Bastian, G. A. Dulk, and Y. Leblanc. A Search for Radio Emission from Extrasolar Planets. *ApJ*, 545:1058–1063, December 2000. doi: 10.1086/317864.

K. Batygin and A. Morbidelli. Dissipative Divergence of Resonant Orbits. *AJ*, 145:1, January 2013. doi: 10.1088/0004-6256/145/1/1.

K. Batygin, P. Bodenheimer, and G. Laughlin. Determination of the Interior Structure of Transiting Planets in Multiple-Planet Systems. *ApJL*, 704:L49–L53, October 2009. doi: 10.1088/0004-637X/704/1/L49.

G. F. Benedict, B. McArthur, E. Nelan, D. Story, A. L. Whipple, P. J. Shelus, W. H. Jefferys, P. D. Hemenway, O. G. Franz, L. H. Wasserman, R. L. Duncombe, W. van Altena, and L. W. Fredrick. Photometry of Proxima Centauri and Barnard's Star Using Hubble Space Telescope Fine Guidance Sensor 3: A Search for Periodic Variations. *AJ*, 116:429–439, July 1998. doi: 10.1086/300420.

D. Bernard, J. Lilensten, M. Barthélemy, and G. Gronoff. Can hydrogen coronae be inferred around a CO_2-dominated exoplanetary atmosphere? *Icarus*, 239:23–31, September 2014. doi: 10.1016/j.icarus.2014.05.043.

Z. K. Berta-Thompson, J. Irwin, D. Charbonneau, E. R. Newton, J. A. Dittmann, N. Astudillo-Defru, X. Bonfils, M. Gillon, E. Jehin, A. A. Stark, B. Stalder, F. Bouchy, X. Delfosse, T. Forveille, C. Lovis, M. Mayor, V. Neves, F. Pepe, N. C. Santos, S. Udry, and A. Wünsche. A rocky planet transiting a nearby low-mass star. *Nature*, 527:204–207, November 2015. doi: 10.1038/nature15762.

J.-L. Bertaux, F. Leblanc, O. Witasse, E. Quemerais, J. Lilensten, S.A. Stern, B. Sandel, and Oleg. Korablev. Discovery of an aurora on mars. *Nature*, 435(7043):790–794, 2005.

J.-L. Beuzit, M. Feldt, K. Dohlen, D. Mouillet, P. Puget, F. Wildi, L. Abe, J. Antichi, A. Baruffolo, P. Baudoz, A. Boccaletti, M. Carbillet, J. Charton, R. Claudi, M. Downing, C. Fabron, P. Feautrier, E. Fedrigo, T. Fusco, J.-L. Gach, R. Gratton, T. Henning, N. Hubin, F. Joos, M. Kasper, M. Langlois, R. Lenzen, C. Moutou, A. Pavlov, C. Petit, J. Pragt, P. Rabou, F. Rigal, R. Roelfsema, G. Rousset, M. Saisse, H.-M. Schmid, E. Stadler, C. Thalmann, M. Turatto, S. Udry, F. Vakili, and R. Waters. SPHERE: a planet finder instrument for the VLT. In *Ground-based and Airborne Instrumentation for Astronomy II*. SPIE-Intl Soc Optical Eng, jul 2008. doi: 10.1117/12.790120. URL `http://dx.doi.org/10.1117/12.790120`.

A. Bhardwaj and G. R. Gladstone. Auroral emissions of the giant planets. *Reviews of Geophysics*, 38(3):295–353, 2000. ISSN 1944-9208. doi: 10.1029/1998RG000046. URL `http://dx.doi.org/10.1029/1998RG000046`.

P. Bodenheimer and J. J. Lissauer. Accretion and Evolution of ˜2.5 M Planets with Voluminous H/He Envelopes. *ApJ*, 791:103, August 2014. doi: 10.1088/0004-637X/791/2/103.

E. Bolmont, S. N. Raymond, J. Leconte, and S. P. Matt. Effect of the stellar spin history on the tidal evolution of close-in planets. *A&A*, 544:A124, August 2012. doi: 10.1051/0004-6361/201219645.

E. Bolmont, F. Selsis, S. N. Raymond, J. Leconte, F. Hersant, A.-S. Maurin, and J. Pericaud. Tidal dissipation and eccentricity pumping: Implications for the depth of the secondary eclipse of 55 Cancri e. *A&A*, 556:A17, August 2013. doi: 10.1051/0004-6361/201220837.

E. Bolmont, S. N. Raymond, J. Leconte, F. Hersant, and A. C. M. Correia. Mercury-T: A new code to study tidally evolving multi-planet systems. Applications to Kepler-62. *A&A*, 583:A116, November 2015. doi: 10.1051/0004-6361/201525909.

E. Bolmont, F. Selsis, J. E. Owen, I. Ribas, S. N. Raymond, J. Leconte, and M. Gillon. Water loss from terrestrial planets orbiting ultracool dwarfs: implications for the planets of TRAPPIST-1. *MNRAS*, 464:3728–3741, January 2017. doi: 10.1093/mnras/stw2578.

W. J. Borucki and A. L. Summers. The photometric method of detecting other planetary systems. *Icarus*, 58:121–134, April 1984. doi: 10.1016/0019-1035(84)90102-7.

W. J. Borucki, D. Koch, G. Basri, N. Batalha, T. Brown, D. Caldwell, J. Caldwell, J. Christensen-Dalsgaard, W. D. Cochran, E. DeVore, E. W. Dunham, A. K. Dupree, T. N. Gautier, J. C. Geary, R. Gilliland, A. Gould, S. B. Howell, J. M. Jenkins, Y. Kondo, D. W. Latham, G. W. Marcy, S. Meibom, H. Kjeldsen, J. J. Lissauer, D. G. Monet, D. Morrison, D. Sasselov, J. Tarter, A. Boss, D. Brownlee, T. Owen, D. Buzasi, D. Charbonneau, L. Doyle, J. Fortney, E. B. Ford, M. J. Holman, S. Seager, J. H. Steffen, W. F. Welsh, J. Rowe, H. Anderson, L. Buchhave, D. Ciardi, L. Walkowicz, W. Sherry, E. Horch, H. Isaacson, M. E. Everett, D. Fischer, G. Torres, J. A. Johnson, M. Endl, P. MacQueen, S. T. Bryson, J. Dotson, M. Haas, J. Kolodziejczak, J. Van Cleve, H. Chandrasekaran, J. D. Twicken, E. V. Quintana, B. D. Clarke, C. Allen, J. Li, H. Wu, P. Tenenbaum, E. Verner, F. Bruhweiler, J. Barnes, and A. Prsa. Kepler Planet-Detection Mission: Introduction and First Results. *Science*, 327:977, February 2010. doi: 10.1126/science.1185402.

W. J. Borucki, D. G. Koch, G. Basri, N. Batalha, T. M. Brown, S. T. Bryson, D. Caldwell, J. Christensen-Dalsgaard, W. D. Cochran, E. DeVore, E. W. Dunham, T. N. Gautier, III, J. C. Geary, R. Gilliland, A. Gould, S. B. Howell, J. M. Jenkins, D. W. Latham, J. J. Lissauer, G. W. Marcy, J. Rowe, D. Sasselov, A. Boss, D. Charbonneau, D. Ciardi, L. Doyle, A. K. Dupree, E. B. Ford, J. Fortney, M. J. Holman, S. Seager, J. H. Steffen, J. Tarter, W. F. Welsh, C. Allen, L. A. Buchhave, J. L. Christiansen, B. D. Clarke, S. Das, J.-M. Désert, M. Endl, D. Fabrycky, F. Fressin, M. Haas, E. Horch, A. Howard, H. Isaacson, H. Kjeldsen, J. Kolodziejczak, C. Kulesa, J. Li, P. W. Lucas, P. Machalek, D. McCarthy, P. MacQueen, S. Meibom, T. Miquel, A. Prsa, S. N. Quinn, E. V. Quintana, D. Ragozzine, W. Sherry, A. Shporer, P. Tenenbaum, G. Torres, J. D. Twicken, J. Van Cleve, L. Walkowicz, F. C. Witteborn, and M. Still. Characteristics of Planetary Candidates Observed by Kepler. II. Analysis of the First Four Months of Data. *ApJ*, 736:19, July 2011. doi: 10.1088/0004-637X/736/1/19.

W. J. Borucki, E. Agol, F. Fressin, L. Kaltenegger, J. Rowe, H. Isaacson, D. Fischer, N. Batalha, J. J. Lissauer, G. W. Marcy, D. Fabrycky, J.-M. Désert, S. T. Bryson, T. Barclay, F. Bastien, A. Boss, E. Brugamyer, L. A. Buchhave, C. Burke, D. A. Caldwell, J. Carter, D. Charbonneau, J. R. Crepp, J. Christensen-Dalsgaard, J. L. Christiansen, D. Ciardi, W. D. Cochran, E. DeVore, L. Doyle, A. K. Dupree, M. Endl, M. E. Everett, E. B. Ford, J. Fortney, T. N. Gautier, J. C. Geary, A. Gould, M. Haas, C. Henze, A. W. Howard, S. B. Howell, D. Huber, J. M. Jenkins, H. Kjeldsen, R. Kolbl, J. Kolodziejczak, D. W. Latham, B. L. Lee, E. Lopez, F. Mullally, J. A. Orosz, A. Prsa, E. V. Quintana, R. Sanchis-Ojeda, D. Sasselov, S. Seader, A. Shporer, J. H. Steffen, M. Still, P. Tenenbaum, S. E. Thompson, G. Torres, J. D. Twicken, W. F. Welsh, and J. N. Winn. Kepler-62: A Five-Planet System with Planets of 1.4 and 1.6 Earth Radii in the Habitable Zone. *Science*, 340:587–590, May 2013. doi: 10.1126/science.1234702.

W.J. Borucki et al. Characteristics of Planetary Candidates Observed by Kepler. II. Analysis of the First Four Months of Data. *ApJ*, 736:19, July 2011. doi: 10.1088/0004-637X/736/1/19.

P. Bouchet, M. García-Marín, P.-O. Lagage, J. Amiaux, J.-L. Auguéres, E. Bauwens, J. A. D. L. Blommaert, C. H. Chen, Ö. H. Detre, D. Dicken, D. Dubreuil, P. Galdemard, R. Gastaud,

A. Glasse, K. D. Gordon, F. Gougnaud, P. Guillard, K. Justtanont, O. Krause, D. Leboeuf, Y. Longval, L. Martin, E. Mazy, V. Moreau, G. Olofsson, T. P. Ray, J.-M. Rees, E. Renotte, M. E. Ressler, S. Ronayette, S. Salasca, S. Scheithauer, J. Sykes, M. P. Thelen, M. Wells, D. Wright, and G. S. Wright. The Mid-Infrared Instrument for the James Webb Space Telescope, III: MIRIM, The MIRI Imager. *PASP*, 127:612, July 2015. doi: 10.1086/682254.

J. Brakensiek and D. Ragozzine. Efficient Geometric Probabilities of Multi-Transiting Exoplanetary Systems from CORBITS. *ApJ*, 821:47, April 2016. doi: 10.3847/0004-637X/821/1/47.

M. Brogi, I. A. G. Snellen, R. J. de Kok, S. Albrecht, J. Birkby, and E. J. W. de Mooij. The signature of orbital motion from the dayside of the planet τ Boötis b. *Nature*, 486:502–504, June 2012. doi: 10.1038/nature11161.

M. Brogi, R. J. de Kok, J. L. Birkby, H. Schwarz, and I. A. G. Snellen. Carbon monoxide and water vapor in the atmosphere of the non-transiting exoplanet HD 179949 b. *A&A*, 565:A124, May 2014. doi: 10.1051/0004-6361/201423537.

B. Brugger, O. Mousis, M. Deleuil, and J. I. Lunine. Possible Internal Structures and Compositions of Proxima Centauri b. *ApJl*, 831:L16, November 2016. doi: 10.3847/2041-8205/831/2/L16.

A. J. Burgasser and E. E. Mamajek. On the Age of the TRAPPIST-1 System. *ArXiv e-prints*, June 2017.

A. J. Burgasser, S. E. Logsdon, J. Gagné, J. J. Bochanski, J. K. Faherty, A. A. West, E. E. Mamajek, S. J. Schmidt, and K. L. Cruz. *AJ*, art. 18. doi: 10.1088/0067-0049/220/1/18.

J. Cabannes and J. Dufay. Mesures interfrentielles des longueurs d'onde des raies interdits de oi dans le ciel nocturne et crpusculaire. *J. Phys. Radium*, 16:129–130, 1955.

J. Cabrera and J. Schneider. Detecting companions to extrasolar planets using mutual events. *A&A*, 464:1133–1138, March 2007. doi: 10.1051/0004-6361:20066111.

T. L. Campante, T. Barclay, J. J. Swift, D. Huber, V. Z. Adibekyan, W. Cochran, C. J. Burke, H. Isaacson, E. V. Quintana, G. R. Davies, V. Silva Aguirre, D. Ragozzine, R. Riddle, C. Baranec, S. Basu, W. J. Chaplin, J. Christensen-Dalsgaard, T. S. Metcalfe, T. R. Bedding, R. Handberg, D. Stello, J. M. Brewer, S. Hekker, C. Karoff, R. Kolbl, N. M. Law, M. Lundkvist, A. Miglio, J. F. Rowe, N. C. Santos, C. Van Laerhoven, T. Arentoft, Y. P. Elsworth, D. A. Fischer, S. D. Kawaler, H. Kjeldsen, M. N. Lund, G. W. Marcy, S. G. Sousa, A. Sozzetti, and T. R. White. An Ancient Extrasolar System with Five Sub-Earth-size Planets. *ApJ*, 799:170, February 2015. doi: 10.1088/0004-637X/799/2/170.

J. M. Carpenter, E. E. Mamajek, L. A. Hillenbrand, and M. R. Meyer. Evidence for Mass-dependent Circumstellar Disk Evolution in the 5 Myr Old Upper Scorpius OB Association. *ApJL*, 651:L49–L52, November 2006. doi: 10.1086/509121.

J. A. Carter, E. Agol, W. J. Chaplin, S. Basu, T. R. Bedding, L. A. Buchhave, J. Christensen-Dalsgaard, K. M. Deck, Y. Elsworth, D. C. Fabrycky, E. B. Ford, J. J. Fortney, S. J. Hale, R. Handberg, S. Hekker, M. J. Holman, D. Huber, C. Karoff, S. D. Kawaler, H. Kjeldsen, J. J. Lissauer, E. D. Lopez, M. N. Lund, M. Lundkvist, T. S. Metcalfe, A. Miglio, L. A. Rogers, D. Stello, W. J. Borucki, S. Bryson, J. L. Christiansen, W. D. Cochran, J. C. Geary, R. L. Gilliland, M. R. Haas, J. Hall, A. W. Howard, J. M. Jenkins, T. Klaus, D. G. Koch, D. W. Latham, P. J. MacQueen, D. Sasselov, J. H. Steffen, J. D. Twicken, and J. N. Winn. Kepler-36: A Pair of Planets with Neighboring Orbits and Dissimilar Densities. *Science*, 337:556, August 2012. doi: 10.1126/science.1223269.

J. C. Carter-Bond, D. P. O'Brien, and S. N. Raymond. The Compositional Diversity of Extrasolar Terrestrial Planets. II. Migration Simulations. *ApJ*, 760:44, November 2012. doi: 10.1088/0004-637X/760/1/44.

D. C. Catling. 6.7 - The Great Oxidation Event Transition. In *Treatise on Geochemistry (Second Edition)*, pages pp. 177 – 195. Elsevier, Oxford, 2014. ISBN 978-0-08-098300-4. doi: http://dx.doi.org/10.1016/B978-0-08-095975-7.01307-3.

D. C. Catling and J. F. Kasting. *Atmospheric Evolution on Inhabited and Lifeless Worlds*. Cambridge University Press, New York, 2017.

D. C. Catling, K. J. Zahnle, and C. P. McKay. Biogenic Methane, Hydrogen Escape, and the Irreversible Oxidation of Early Earth. *Science*, 293:839–843, August 2001. doi: 10.1126/science.1061976.

G. Chabrier and I. Baraffe. Structure and evolution of low-mass stars. *A&A*, 327:1039–1053, November 1997.

J. W. Chamberlain. *Physics of the aurora and airglow*. International Geophysics Series, New York: Academic Press, 1961.

J. E. Chambers. Planetary accretion in the inner Solar System. *Earth and Planetary Science Letters*, 223:241–252, July 2004. doi: 10.1016/j.epsl.2004.04.031.

S.-H. Chang, P.-G. Gu, and P. H. Bodenheimer. Tidal and Magnetic Interactions Between a Hot Jupiter and its Host Star in the Magnetospheric Cavity of a Protoplanetary Disk. *ApJ*, 708:1692–1702, January 2010. doi: 10.1088/0004-637X/708/2/1692.

D. Charbonneau, L. E. Allen, S. T. Megeath, G. Torres, R. Alonso, T. M. Brown, R. L. Gilliland, D. W. Latham, G. Mandushev, F. T. O'Donovan, and A. Sozzetti. Detection of thermal emission from an extrasolar planet. *The Astrophysical Journal*, 626(1):523–529, jun 2005. doi: 10.1086/429991. URL `https://doi.org/10.1086/429991`.

E. Chassefière. Hydrodynamic escape of hydrogen from a hot water-rich atmosphere: The case of Venus. *J. Geophys. Res.*, 101:26039–26056, November 1996a. doi: 10.1029/96JE01951.

E. Chassefière. Hydrodynamic Escape of Oxygen from Primitive Atmospheres: Applications to the Cases of Venus and Mars. *Icarus*, 124:537–552, December 1996b. doi: 10.1006/icar.1996.0229.

E. Chassefière. NOTE: Loss of Water on the Young Venus: The Effect of a Strong Primitive Solar Wind. *Icarus*, 126:229–232, March 1997. doi: 10.1006/icar.1997.5677.

E. Chassefière, R. Wieler, B. Marty, and F. Leblanc. The evolution of Venus: Present state of knowledge and future exploration. *Planet. Space Sci.*, 63:15–23, April 2012. doi: 10.1016/j.pss.2011.04.007.

A. F. Cheng. Triton torus and Neptune aurora. *Geophysical Research Letters*, 17(10):1669–1672, 1990. ISSN 1944-8007. doi: 10.1029/GL017i010p01669. URL `http://dx.doi.org/10.1029/GL017i010p01669`.

J. Choi, C. McCarthy, G. W. Marcy, A. W. Howard, D. A. Fischer, J. A. Johnson, H. Isaacson, and J. T. Wright. Precise Doppler Monitoring of Barnard's Star. *ApJ*, 764:131, February 2013. doi: 10.1088/0004-637X/764/2/131.

M. Choukroun and O. Grasset. Thermodynamic model for water and high-pressure ices up to 2.2 GPa and down to the metastable domain. *J. Chem. Phys.*, 127(12):124506, September 2007. doi: 10.1063/1.2768957.

J. L. Christiansen, J. M. Jenkins, D. A. Caldwell, C. J. Burke, P. Tenenbaum, S. Seader, S. E. Thompson, T. S. Barclay, B. D. Clarke, J. Li, J. C. Smith, M. C. Stumpe, J. D. Twicken, and J. Van Cleve. The Derivation, Properties, and Value of Kepler's Combined Differential Photometric Precision. *PASP*, 124:1279–1287, December 2012. doi: 10.1086/668847.

A. Claret. A new non-linear limb-darkening law for LTE stellar atmosphere models. *A&A*, 363:1081–1190, November 2000.

A. Claret and S. Bloemen. Gravity and limb-darkening coefficients for the Kepler, CoRoT, Spitzer, uvby, UBVRIJHK, and Sloan photometric systems. *A&A*, 529:75, January 2011.

O. Cohen, J. J. Drake, A. Glocer, C. Garraffo, K. Poppenhaeger, J. M. Bell, A. J. Ridley, and T. I. Gombosi. Magnetospheric Structure and Atmospheric Joule Heating of Habitable Planets Orbiting M-dwarf Stars. *ApJ*, 790:57, July 2014. doi: 10.1088/0004-637X/790/1/57.

J. M. Collins, H. R. A. Jones, and J. R. Barnes. Calculations of periodicity from H-alpha profiles of Proxima Centauri. *ArXiv e-prints, arXiv:1608.07834*, 2016.

B. A. Cook, P. K. G. Williams, and E. Berger. Trends in Ultracool Dwarf Magnetism. II. The Inverse Correlation Between X-Ray Activity and Rotation as Evidence for a Bimodal Dynamo. *ApJ*, 785:10, April 2014. doi: 10.1088/0004-637X/785/1/10.

A. R. Cooray and Origins Space Telescope Study Team. Origins Space Telescope. In *American Astronomical Society Meeting Abstracts*, volume 229 of *American Astronomical Society Meeting Abstracts*, page 405.01, January 2017.

A. C. M. Correia and J. Laskar. *Tidal Evolution of Exoplanets*, pages 239–266. 2011.

C. Cossou, S. N. Raymond, and A. Pierens. Convergence zones for Type I migration: an inward shift for multiple planet systems. *A&A*, 553:L2, May 2013. doi: 10.1051/0004-6361/201220853.

C. Cossou, S. N. Raymond, F. Hersant, and A. Pierens. Hot super-Earths and giant planet cores from different migration histories. *A&A*, 569:A56, September 2014. doi: 10.1051/0004-6361/201424157.

N. B. Cowan and E. Agol. Inverting phase functions to map exoplanets. *The Astrophysical Journal*, 678(2):L129–L132, apr 2008. doi: 10.1086/588553. URL `https://doi.org/10.1086/588553`.

N. B. Cowan and Y. Fujii. Mapping Exoplanets. *ArXiv e-prints*, April 2017.

N. B. Cowan, E. Agol, and D. Charbonneau. Hot nights on extrasolar planets: mid-infrared phase variations of hot Jupiters. *MNRAS*, 379:641–646, August 2007. doi: 10.1111/j.1365-2966.2007.11897.x.

N. B. Cowan, E. Agol, V. S. Meadows, T. Robinson, T. A. Livengood, D. Deming, C. M. Lisse, M. F. A'Hearn, D. D. Wellnitz, S. Seager, and D. Charbonneau. Alien Maps of an Ocean-Bearing World. *The Astrophysical Journal*, 700(2):915–923, jul 2009. doi: 10.1088/0004-637x/700/2/915. URL `https://doi.org/10.1088/0004-637x/700/2/915`.

I. J. M. Crossfield, B. M. S. Hansen, J. Harrington, J. Y.-K. Cho, D. Deming, K. Menou, and S. Seager. A New 24 μm Phase Curve for Upsilon Andromedae b. *ApJ*, 723:1436–1446, November 2010. doi: 10.1088/0004-637X/723/2/1436.

I. J. M. Crossfield, E. Petigura, J. E. Schlieder, A. W. Howard, B. J. Fulton, K. M. Aller, D. R. Ciardi, S. Lépine, T. Barclay, I. de Pater, K. de Kleer, E. V. Quintana, J. L. Christiansen, E. Schlafly, L. Kaltenegger, J. R. Crepp, T. Henning, C. Obermeier, N. Deacon, L. M. Weiss, H. T. Isaacson, B. M. S. Hansen, M. C. Liu, T. Greene, S. B. Howell, T. Barman, and C. Mordasini. A Nearby M Star with Three Transiting Super-Earths Discovered by K2. *ApJ*, 804:10, May 2015. doi: 10.1088/0004-637X/804/1/10.

J. Dalcanton, S. Seager, S. Aigrain, S. Battel, N. Brandt, C. Conroy, L. Feinberg, S. Gezari, O. Guyon, W. Harris, C. Hirata, J. Mather, M. Postman, D. Redding, D. Schiminovich, H. P. Stahl, and J. Tumlinson. From Cosmic Birth to Living Earths: The Future of UVOIR Space Astronomy. *ArXiv e-prints, arXiv:1507.04779*, July 2015.

G. H. Darwin. On the Secular Changes in the Elements of the Orbit of a Satellite Revolving about a Tidally Distorted Planet. *Royal Society of London Philosophical Transactions Series I*, 171: 713–891, 1880.

J. R. A. Davenport, S. L. Hawley, L. Hebb, J. P. Wisniewski, A. F. Kowalski, E. C. Johnson, M. Malatesta, J. Peraza, M. Keil, S. M. Silverberg, T. C. Jansen, M. S. Scheffler, J. R. Berdis, D. M. Larsen, and E. J. Hilton. Kepler Flares. II. The Temporal Morphology of White-light Flares on GJ 1243. *ApJ*, 797:122, December 2014. doi: 10.1088/0004-637X/797/2/122.

J. R. A. Davenport, D. M. Kipping, D. Sasselov, J. M. Matthews, and C. Cameron. MOST Observations of Our Nearest Neighbor: Flares on Proxima Centauri. *ApJl*, 829:L31, October 2016. doi: 10.3847/2041-8205/829/2/L31.

J. H. Davies and D. R. Davies. Earth's surface heat flux. *Solid Earth*, 1(1):5–24, 2010. doi: 10.5194/se-1-5-2010. URL `http://www.solid-earth.net/1/5/2010/`.

K. de Kleer, M. Skrutskie, J. Leisenring, A. G. Davies, A. Conrad, I. de Pater, A. Resnick, V. Bailey, D. Defrère, P. Hinz, A. Skemer, E. Spalding, A. Vaz, C. Veillet, and C. E. Woodward. Multi-phase volcanic resurfacing at Loki Patera on Io. *Nature*, 545:199–202, May 2017. doi: 10.1038/nature22339.

J. de Wit, M. Gillon, B.-O. Demory, and S. Seager. Towards consistent mapping of distant worlds: secondary-eclipse scanning of the exoplanet HD 189733b. *A&A*, 548:A128, December 2012. doi: 10.1051/0004-6361/201219060.

J. de Wit, H. R. Wakeford, M. Gillon, N. K. Lewis, J. A. Valenti, B.-O. Demory, A. J. Burgasser, A. Burdanov, L. Delrez, E. Jehin, S. M. Lederer, D. Queloz, A. H. M. J. Triaud, and V. Van Grootel. A combined transmission spectrum of the Earth-sized exoplanets TRAPPIST-1 b and c. *Nature*, 537:69–72, September 2016. doi: 10.1038/nature18641.

K. M. Deck and E. Agol. Measurement of Planet Masses with Transit Timing Variations due to Synodic Chopping Effects. *ApJ*, 802(2):116, apr 2015. doi: 10.1088/0004-637x/802/2/116. URL `http://dx.doi.org/10.1088/0004-637X/802/2/116`.

L. Delchambre. Weighted principal component analysis: a weighted covariance eigendecomposition approach. *MNRAS*, 446:3545–3555, February 2015. doi: 10.1093/mnras/stu2219.

D. Deming, H. Knutson, J. Kammer, B. J. Fulton, J. Ingalls, S. Carey, A. Burrows, J. J. Fortney, K. Todorov, E. Agol, N. Cowan, J.-M. Desert, J. Fraine, J. Langton, C. Morley, and A. P. Showman. Spitzer Secondary Eclipses of the Dense, Modestly-irradiated, Giant Exoplanet HAT-P-20b Using Pixel-level Decorrelation. *ApJ*, 805:132, June 2015. doi: 10.1088/0004-637X/805/2/132.

B.-O. Demory, M. Gillon, J. de Wit, N. Madhusudhan, E. Bolmont, K. Heng, T. Kataria, N. Lewis, R. Hu, J. Krick, V. Stamenković, B. Benneke, S. Kane, and D. Queloz. A map of the large day-night temperature gradient of a super-Earth exoplanet. *Nature*, 532:207–209, April 2016a. doi: 10.1038/nature17169.

B.-O. Demory, D. Queloz, Y. Alibert, E. Gillen, and M. Gillon. Probing TRAPPIST-1-like Systems with K2. *ApJL*, 825:L25, July 2016b.

J. T. Dempsey, J. W. V. Storey, and A. Phillips. Auroral Contribution to Sky Brightness for Optical Astronomy on the Antarctic Plateau. *PASA*, 22:91–104, 2005. doi: 10.1071/AS04036.

M. D. Desch and M. L. Kaiser. Predictions for Uranus from a radiometric Bode's law. *Nature*, 310:755–757, August 1984. doi: 10.1038/310755a0.

S. D. Domagal-Goldman, A. Segura, M. W. Claire, T. D. Robinson, and V. S. Meadows. Abiotic Ozone and Oxygen in Atmospheres Similar to Prebiotic Earth. *ApJ*, 792:90, September 2014. doi: 10.1088/0004-637X/792/2/90.

S. D. Domagal-Goldman, A. Roberge, G. N. Arney, A. M. Mandell, R. K. Kopparapu, and Luvoir Sci. Tech. Definition Team. The Next Generation of Observations of Planets Beyond Our Solar System. In *Planetary Science Vision 2050 Workshop*, volume 1989 of *LPI Contributions*, page 8189, February 2017.

J.-P. Dou, D.-Q. Ren, and Y.-T. Zhu. High-contrast coronagraph for ground-based imaging of Jupiter-like planets. *Research in Astronomy and Astrophysics*, 10(2):189, 2010. URL `http://stacks.iop.org/1674-4527/10/i=2/a=010`.

C. D. Dressing and D. Charbonneau. The Occurrence Rate of Small Planets around Small Stars. *ApJ*, 767:95, April 2013. doi: 10.1088/0004-637X/767/1/95.

C. D. Dressing and D. Charbonneau. The Occurrence of Potentially Habitable Planets Orbiting M Dwarfs Estimated from the Full Kepler Dataset and an Empirical Measurement of the Detection Sensitivity. *ApJ*, 807:45, July 2015. doi: 10.1088/0004-637X/807/1/45.

P. Driscoll and D. Bercovici. Divergent evolution of Earth and Venus: Influence of degassing, tectonics, and magnetic fields. *Icarus*, 226:1447–1464, November 2013. doi: 10.1016/j.icarus.2013.07.025.

P. Driscoll and P. Olson. Optimal dynamos in the cores of terrestrial exoplanets: Magnetic field generation and detectability. *Icarus*, 213:12–23, May 2011. doi: 10.1016/j.icarus.2011.02.010.

P. E. Driscoll and R. Barnes. Tidal heating of earth-like exoplanets around m stars: thermal, magnetic, and orbital evolutions. *Astrobiology*, 15(9):739–760, 2015.

B. R. Durney, D. S. De Young, and I. W. Roxburgh. On the generation of the large-scale and turbulent magnetic fields in solar-type stars. *Sol. Phys.*, 145:207–225, June 1993. doi: 10.1007/BF00690652.

P. P. Eggleton, L. G. Kiseleva, and P. Hut. The Equilibrium Tide Model for Tidal Friction. *ApJ*, 499:853–870, May 1998. doi: 10.1086/305670.

L. T. Elkins-Tanton. Linked magma ocean solidification and atmospheric growth for Earth and Mars. *Earth and Planetary Science Letters*, 271:181–191, July 2008. doi: 10.1016/j.epsl.2008.03.062.

L. T. Elkins-Tanton. Formation of early water oceans on rocky planets. *Ap&SS*, 332:359–364, April 2011. doi: 10.1007/s10509-010-0535-3.

L. T. Elkins-Tanton. Magma Oceans in the Inner Solar System. *Annual Review of Earth and Planetary Sciences*, 40:113–139, May 2012. doi: 10.1146/annurev-earth-042711-105503.

L. T. Elkins-Tanton and S. Seager. Ranges of Atmospheric Mass and Composition of Super-Earth Exoplanets. *ApJ*, 685:1237–1246, October 2008. doi: 10.1086/591433.

S. G. Engle and E. F. Guinan. Red Dwarf Stars: Ages, Rotation, Magnetic Dynamo Activity and the Habitability of Hosted Planets. In S. Qain, K. Leung, L. Zhu, and S. Kwok, editors, *Astronomical Society of the Pacific Conference Series*, volume 451 of *Astronomical Society of the Pacific Conference Series*, page 285, December 2011.

N. V. Erkaev, Y. N. Kulikov, H. Lammer, F. Selsis, D. Langmayr, G. F. Jaritz, and H. K. Biernat. Roche lobe effects on the atmospheric loss from "Hot Jupiters". *A&A*, 472:329–334, September 2007. doi: 10.1051/0004-6361:20066929.

N. V. Erkaev, H. Lammer, P. Odert, Y. N. Kulikov, K. G. Kislyakova, M. L. Khodachenko, M. Güdel, A. Hanslmeier, and H. Biernat. XUV-Exposed, Non-Hydrostatic Hydrogen-Rich Upper Atmospheres of Terrestrial Planets. Part I: Atmospheric Expansion and Thermal Escape. *Astrobiology*, 13:1011–1029, November 2013. doi: 10.1089/ast.2012.0957.

N. V. Erkaev, H. Lammer, L. T. Elkins-Tanton, A. Stökl, P. Odert, E. Marcq, E. A. Dorfi, K. G. Kislyakova, Y. N. Kulikov, M. Leitzinger, and M. Güdel. Escape of the martian protoatmosphere and initial water inventory. *Planetary and Space Science*, 98:106–119, August 2014. doi: 10.1016/j.pss.2013.09.008.

D. C. Fabrycky, E. B. Ford, J. H. Steffen, J. F. Rowe, J. A. Carter, A. V. Moorhead, N. M. Batalha, W. J. Borucki, S. Bryson, L. A. Buchhave, J. L. Christiansen, D. R. Ciardi, W. D. Cochran, M. Endl, M. N. Fanelli, D. Fischer, F. Fressin, J. Geary, M. R. Haas, J. R. Hall, M. J. Holman, J. M. Jenkins, D. G. Koch, D. W. Latham, J. Li, J. J. Lissauer, P. Lucas, G. W. Marcy, T. Mazeh, S. McCauliff, S. Quinn, D. Ragozzine, D. Sasselov, and A. Shporer. Transit Timing Observations from Kepler. IV. Confirmation of Four Multiple-planet Systems by Simple Physical Models. *ApJ*, 750:114, May 2012. doi: 10.1088/0004-637X/750/2/114.

D. C. Fabrycky, J. J. Lissauer, D. Ragozzine, J. F. Rowe, J. H. Steffen, E. Agol, T. Barclay, N. Batalha, W. Borucki, D. R. Ciardi, E. B. Ford, T. N. Gautier, J. C. Geary, M. J. Holman, J. M. Jenkins, J. Li, R. C. Morehead, R. L. Morris, A. Shporer, J. C. Smith, M. Still, and J. Van Cleve. Architecture of Kepler's Multi-transiting Systems. II. New Investigations with Twice as Many Candidates. *ApJ*, 790:146, August 2014. doi: 10.1088/0004-637X/790/2/146.

S. Faigler and T. Mazeh. Photometric detection of non-transiting short-period low-mass companions through the beaming, ellipsoidal and reflection effects in Kepler and CoRoT light curves. *MNRAS*, 415:3921–3928, August 2011. doi: 10.1111/j.1365-2966.2011.19011.x.

S. Ferraz-Mello, A. Rodríguez, and H. Hussmann. Tidal friction in close-in satellites and exoplanets: The Darwin theory re-visited. *Celestial Mechanics and Dynamical Astronomy*, 101:171–201, May 2008. doi: 10.1007/s10569-008-9133-x.

E. B. Ford and F. A. Rasio. Origins of Eccentric Extrasolar Planets: Testing the Planet-Planet Scattering Model. *ApJ*, 686:621–636, October 2008. doi: 10.1086/590926.

D. Foreman-Mackey. corner.py: Scatterplot matrices in python. *The Journal of Open Source Software*, 24, 2016. doi: 10.21105/joss.00024. URL `http://dx.doi.org/10.5281/zenodo.45906`.

D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman. emcee: The MCMC Hammer. *PASP*, 125:306, March 2013. doi: 10.1086/670067.

D. Foreman-Mackey, D. W. Hogg, and T. D. Morton. Exoplanet Population Inference and the Abundance of Earth Analogs from Noisy, Incomplete Catalogs. *ApJ*, 795:64, November 2014. doi: 10.1088/0004-637X/795/1/64.

D. Foreman-Mackey, B. T. Montet, D. W. Hogg, T. D. Morton, D. Wang, and B. Schölkopf. A Systematic Search for Transiting Planets in the K2 Data. *ApJ*, 806:215, June 2015. doi: 10.1088/0004-637X/806/2/215.

J. J. Fortney, M. S. Marley, and J. W. Barnes. Planetary Radii across Five Orders of Magnitude in Mass and Stellar Insolation: Application to Transits. *ApJ*, 659:1661–1672, April 2007. doi: 10.1086/512120.

L. Fossati, D. Bisikalo, H. Lammer, B. Shustov, and M. Sachkov. Major prospects of exoplanet astronomy with the World Space ObservatoryUltraViolet mission. *Ap&SS*, 354:9–19, November 2014. doi: 10.1007/s10509-014-2027-3.

K. France, J. T. Stocke, H. Yang, J. L. Linsky, B. C. Wolven, C. S. Froning, J. C. Green, and S. N. Osterman. Searching for Far-ultraviolet Auroral/Dayglow Emission from HD 209458b. *ApJ*, 712:1277–1286, April 2010. doi: 10.1088/0004-637X/712/2/1277.

M. Fridlund, E. Gaidos, O. Barragán, C. M. Persson, D. Gandolfi, J. Cabrera, T. Hirano, M. Kuzuhara, S. Csizmadia, G. Nowak, M. Endl, S. Grziwa, J. Korth, J. Pfaff, B. Bitsch, A. Johansen, A. J. Mustill, M. B. Davies, H. Deeg, E. Palle, W. D. Cochran, P. Eigmüller, A. Erikson, E. Guenther, A. P. Hatzes, A. Kiilerich, T. Kudo, P. MacQueen, N. Narita, D. Nespral, M. Pätzold, J. Prieto-Arranz, H. Rauer, and V. Van Eylen. EPIC 210894022b - A short period super-Earth transiting a metal poor, evolved old star. *ArXiv e-prints*, April 2017.

B. J. Fulton, E. A. Petigura, A. W. Howard, H. Isaacson, G. W. Marcy, P. A. Cargile, L. Hebb, L. M. Weiss, J. A. Johnson, T. D. Morton, E. Sinukoff, I. J. M. Crossfield, and L. A. Hirsch. The California-Kepler Survey. III. A Gap in the Radius Distribution of Small Planets. *ArXiv e-prints*, March 2017.

J. P. Gardner, J. C. Mather, M. Clampin, R. Doyon, M. A. Greenhouse, H. B. Hammel, J. B. Hutchings, P. Jakobsen, S. J. Lilly, K. S. Long, J. I. Lunine, M. J. McCaughrean, M. Mountain, J. Nella, G. H. Rieke, M. J. Rieke, H.-W. Rix, E. P. Smith, G. Sonneborn, M. Stiavelli, H. S.

Stockman, R. A. Windhorst, and G. S. Wright. The James Webb Space Telescope. *SSR*, 123: 485–606, April 2006. doi: 10.1007/s11214-006-8315-7.

C. Garraffo, J. J. Drake, and O. Cohen. The Space Weather of Proxima Centauri b. *ApJl*, 833:L4, December 2016. doi: 10.3847/2041-8205/833/1/L4.

G. Gatewood and H. Eichhorn. An unsuccessful search for a planetary companion of Barnard's star BD +4 3561. *AJ*, 78:769–776, October 1973. doi: 10.1086/111480.

R. L. Gilliland, J. M. Jenkins, W. J. Borucki, S. T. Bryson, D. A. Caldwell, B. D. Clarke, J. L. Dotson, M. R. Haas, J. Hall, T. Klaus, D. Koch, S. McCauliff, E. V. Quintana, J. D. Twicken, and J. E. van Cleve. Initial Characteristics of Kepler Short Cadence Data. *The Astrophysical Journal Letters*, 713(2):L160, 2010. URL `http://stacks.iop.org/2041-8205/713/i=2/a=L160`.

R. L. Gilliland, W. J. Chaplin, E. W. Dunham, V. S. Argabright, W. J. Borucki, G. Basri, S. T. Bryson, D. L. Buzasi, D. A. Caldwell, Y. P. Elsworth, J. M. Jenkins, D. G. Koch, J. Kolodziejczak, A. Miglio, J. van Cleve, L. M. Walkowicz, and W. F. Welsh. Kepler Mission Stellar and Instrument Noise Properties. *ApJS*, 197:6, November 2011. doi: 10.1088/0067-0049/197/1/6.

C. Gillmann, E. Chassefière, and P. Lognonné. A consistent picture of early hydrodynamic escape of Venus atmosphere explaining present Ne and Ar isotopic ratios and low oxygen atmospheric content. *Earth and Planetary Science Letters*, 286:503–513, September 2009. doi: 10.1016/j.epsl.2009.07.016.

M. Gillon, A. H. M. J. Triaud, J. J. Fortney, B.-O. Demory, E. Jehin, M. Lendl, P. Magain, P. Kabath, D. Queloz, R. Alonso, D. R. Anderson, A. Collier Cameron, A. Fumel, L. Hebb, C. Hellier, A. Lanotte, P. F. L. Maxted, N. Mowlavi, and B. Smalley. The TRAPPIST survey of southern transiting planets. I. Thirty eclipses of the ultra-short period planet WASP-43 b. *A&A*, 542:A4, June 2012.

M. Gillon, E. Jehin, S. M. Lederer, L. Delrez, J. de Wit, A. Burdanov, V. Van Grootel, A. J. Burgasser, A. H. M. J. Triaud, C. Opitom, B.-O. Demory, D. K. Sahu, D. Bardalez Gagliuffi, P. Magain, and D. Queloz. Temperate Earth-sized planets transiting a nearby ultracool dwarf star. *Nature*, 533:221–224, May 2016. doi: 10.1038/nature17448.

M. Gillon, A. H. M. J. Triaud, B.-O. Demory, E. Jehin, E. Agol, K. M. Deck, S. M. Lederer, J. de Wit, A. Burdanov, J. G. Ingalls, E. Bolmont, J. Leconte, S. N. Raymond, F. Selsis, M. Turbet, K. Barkaoui, A. Burgasser, M. R. Burleigh, S. J. Carey, A. Chaushev, C. M. Copperwheat, L. Delrez, C. S. Fernandes, D. L. Holdsworth, E. J. Kotze, V. Van Grootel, Y. Almleaky, Z. Benkhaldoun, P. Magain, and D. Queloz. Seven temperate terrestrial planets around the nearby ultracool dwarf star TRAPPIST-1. *Nature*, 542:456–460, February 2017. doi: 10.1038/nature21360.

A. Giménez. Equations for the analysis of the light curves of extra-solar planetary transits. *Astronomy & Astrophysics*, 450(3):1231–1237, apr 2006. doi: 10.1051/0004-6361:20054445. URL https://doi.org/10.1051/0004-6361:20054445.

A. Glasse, G. H. Rieke, E. Bauwens, M. García-Marín, M. E. Ressler, S. Rost, T. V. Tikkanen, B. Vandenbussche, and G. S. Wright. The Mid-Infrared Instrument for the James Webb Space Telescope, IX: Predicted Sensitivity. *PASP*, 127:686, July 2015. doi: 10.1086/682259.

P. Goldreich. Final spin states of planets and satellites. *AJ*, 71:1, February 1966. doi: 10.1086/109844.

P. Goldreich and S. Soter. Q in the Solar System. *Icarus*, 5:375–389, 1966. doi: 10.1016/0019-1035(66)90051-0.

G. H. Golub and C. F. Van Loan. *Matrix Computations (3rd ed.; Baltimore: Johns Hopkins Univ. Press)*. 1996.

R. Gomes, H. F. Levison, K. Tsiganis, and A. Morbidelli. Origin of the cataclysmic Late Heavy Bombardment period of the terrestrial planets. *Nature*, 435:466–469, May 2005. doi: 10.1038/nature03676.

W. D. Gonzalez, J. A. Joselyn, Y. Kamide, H. W. Kroehl, G. Rostoker, B. T. Tsurutani, and V. M. Vasyliunas. What is a geomagnetic storm? *Journal of Geophysical Research: Space Physics*, 99 (A4):5771–5792, 1994.

W. D. Gonzalez, B. T. Tsurutani, and A. L. C. De Gonzalez. Interplanetary origin of geomagnetic storms. *Space Science Reviews*, 88(3-4):529–562, 1999.

N. Gopalswamy, S. Yashiro, G. Michalek, G. Stenborg, A. Vourlidas, S. Freeland, and R. Howard. The SOHO/LASCO CME Catalog. *Earth, Moon, and Planets*, 104(1-4):295–313, 2009.

K. Goździewski, C. Migaszewski, F. Panichi, and E. Szuszkiewicz. The Laplace resonance in the Kepler-60 planetary system. *MNRAS*, 455:L104–L108, January 2016. doi: 10.1093/mnrasl/slv156.

C. L. Gray, N. J. Chanover, T. G. Slanger, and K. Molaverdikhani. The effect of solar flares, coronal mass ejections, and solar wind streams on Venus 5577Å oxygen green line. *Icarus*, 233: 342–347, 2014.

R. Greenberg. Frequency Dependence of Tidal q. *ApJL*, 698:L42–L45, June 2009. doi: 10.1088/0004-637X/698/1/L42.

R. G. H. Greer, D. P. Murtagh, I. C. McDade, P. H. G. Dickinson, L. Thomas, D. B. Jenkins, J. Stegman, E. J. Llewellyn, G. Witt, D. J. Mackinnon, et al. ETON 1: A data base pertinent to the study of energy transfer in the oxygen nightglow. *Planetary and space science*, 34(9): 771–788, 1986.

J.-M. Grießmeier. Detection methods and relevance of exoplanetary magnetic fields. In *Characterizing stellar and exoplanetary environments*, pages 213–237. Springer, 2015.

J.-M. Grießmeier, P. Zarka, and H. Spreeuw. Predicting low-frequency radio fluxes of known extrasolar planets. *A&A*, 475:359–368, November 2007. doi: 10.1051/0004-6361:20077397.

P.-G. Gu, D. N. C. Lin, and P. H. Bodenheimer. The Effect of Tidal Inflation Instability on the Mass and Dynamical Evolution of Extrasolar Planets with Ultrashort Periods. *ApJ*, 588: 509–534, May 2003. doi: 10.1086/373920.

O. Guyon and F. Martinache. How Extremely Large Telescopes (ELTs) will Acquire the First Spectra of Rocky Habitable Planets. In *American Astronomical Society Meeting Abstracts #221*, volume 221 of *American Astronomical Society Meeting Abstracts*, page 419.05, January 2013.

O. Guyon, F. Martinache, E. J. Cady, R. Belikov, K. Balasubramanian, D. Wilson, C. S. Clergeon, and M. Mateen. How ELTs will acquire the first spectra of rocky habitable planets , 2012. URL `http://dx.doi.org/10.1117/12.927181`.

J. B. S. Haldane. The origin of life. *The Rationalist Annal.*, pages 3–10, 1929.

K. Hamano, Y. Abe, and H. Genda. Emergence of two types of terrestrial planet on solidification of magma ocean. *Nature*, 497:607–610, May 2013. doi: 10.1038/nature12163.

E. Han, S. X. Wang, J. T. Wright, Y. K. Feng, M. Zhao, O. Fakhouri, J. I. Brown, and C. Hancock. Exoplanet Orbit Database. II. Updates to Exoplanets.org. *PASP*, 126:827, September 2014. doi: 10.1086/678447.

M. H. Hart. Habitable Zones about Main Sequence Stars. *Icarus*, 37:351–357, January 1979. doi: 10.1016/0019-1035(79)90141-6.

P. Hartogh, D. C. Lis, D. Bockelée-Morvan, M. de Val-Borro, N. Biver, M. Küppers, M. Emprechtinger, E. A. Bergin, J. Crovisier, M. Rengel, R. Moreno, S. Szutowicz, and G. A. Blake. Ocean-like water in the Jupiter-family comet 103P/Hartley 2. *Nature*, 478:218–220, October 2011. doi: 10.1038/nature10519.

G. L. Hashimoto, Y. Abe, and S. Sugita. The chemical composition of the early terrestrial atmosphere: Formation of a reducing atmosphere from CI-like material. *Journal of Geophysical Research (Planets)*, 112:E05010, May 2007. doi: 10.1029/2006JE002844.

C. Hayashi. Stellar evolution in early phases of gravitational contraction. *PASJ*, 13:450–452, 1961.

C. Hayashi, K. Nakazawa, and Y. Nakagawa. Formation of the solar system. In D. C. Black and M. S. Matthews, editors, *Protostars and Planets II*, pages 1100–1153, 1985.

L. E. Hays. NASA Astrobiology Strategy Document. 2015. URL `https://astrobiology.nasa.gov/research/astrobiology-at-nasa/astrobiology-strategy/`.

M. Y. He, A. H. M. J. Triaud, and M. Gillon. First limits on the occurrence rate of short-period planets orbiting brown dwarfs. *MNRAS*, 464:2687–2697, January 2017. doi: 10.1093/mnras/stw2391.

R. Heller and R. Barnes. Runaway greenhouse effect on exomoons due to irradiation from hot, young giant planets. *ArXiv e-prints*, November 2013.

R. Heller, J. Leconte, and R. Barnes. Tidal obliquity evolution of potentially habitable planets. *A&A*, 528:A27, April 2011. doi: 10.1051/0004-6361/201015809.

J. L. Hershey. Astrometric analysis of the field of AC +65 6955 from plates taken with the Sproul 24-inch refractor. *AJ*, 78:421–425, June 1973. doi: 10.1086/111436.

S. L. G. Hess and P. Zarka. Modeling the radio signature of the orbital parameters, rotation, and magnetic field of exoplanets. *A&A*, 531:A29, July 2011. doi: 10.1051/0004-6361/201116510.

E. J. Hilton, S. L. Hawley, A. F. Kowalski, and J. Holtzman. The Galactic M Dwarf Flare Rate. In C. Johns-Krull, M. K. Browning, and A. A. West, editors, *16th Cambridge Workshop on Cool Stars, Stellar Systems, and the Sun*, volume 448 of *Astronomical Society of the Pacific Conference Series*, page 197, December 2011.

J. L. Hilton, P. K. Seidelmann, and C. Liu. Analysis of ancient Chinese records of occultations between planets and stars. *AJ*, 96:1482–1493, October 1988. doi: 10.1086/114900.

T. Hirano, N. Narita, B. Sato, Y. H. Takahashi, K. Masuda, Y. Takeda, W. Aoki, M. Tamura, and Y. Suto. Planet-planet eclipse and the Rossiter-McLaughlin effect of a multiple transiting system: joint analysis of the SUBARU spectroscopy and the Kepler photometry. *The Astrophysical Journal*, 759(2):L36, oct 2012. doi: 10.1088/2041-8205/759/2/l36. URL `https://doi.org/10.1088/2041-8205/759/2/l36`.

A. W. Howard, G. W. Marcy, S. T. Bryson, J. M. Jenkins, J. F. Rowe, N. M. Batalha, W. J. Borucki, D. G. Koch, E. W. Dunham, T. N. Gautier, III, J. Van Cleve, W. D. Cochran, D. W. Latham, J. J. Lissauer, G. Torres, T. M. Brown, R. L. Gilliland, L. A. Buchhave, D. A. Caldwell, J. Christensen-Dalsgaard, D. Ciardi, F. Fressin, M. R. Haas, S. B. Howell, H. Kjeldsen, S. Seager, L. Rogers, D. D. Sasselov, J. H. Steffen, G. S. Basri, D. Charbonneau, J. Christiansen, B. Clarke, A. Dupree, D. C. Fabrycky, D. A. Fischer, E. B. Ford, J. J. Fortney, J. Tarter, F. R. Girouard, M. J. Holman, J. A. Johnson, T. C. Klaus, P. Machalek, A. V. Moorhead, R. C. Morehead, D. Ragozzine, P. Tenenbaum, J. D. Twicken, S. N. Quinn, H. Isaacson, A. Shporer, P. W. Lucas, L. M. Walkowicz, W. F. Welsh, A. Boss, E. Devore, A. Gould, J. C. Smith, R. L. Morris, A. Prsa, T. D. Morton, M. Still, S. E. Thompson, F. Mullally, M. Endl, and P. J. MacQueen. Planet Occurrence within 0.25 AU of Solar-type Stars from Kepler. *ApJs*, 201:15, August 2012. doi: 10.1088/0067-0049/201/2/15.

S. B. Howell, C. Sobeck, M. Haas, M. Still, T. Barclay, F. Mullally, J. Troeltzsch, S. Aigrain, S. T. Bryson, D. Caldwell, W. J. Chaplin, W. D. Cochran, D. Huber, G. W. Marcy, A. Miglio, J. R.

Najita, M. Smith, J. D. Twicken, and J. J. Fortney. The K2 Mission: Characterization and Early results. *ArXiv e-prints*, February 2014.

C. X. Huang, K. Penev, J. D. Hartman, G. Á. Bakos, W. Bhatti, I. Domsa, and M. de Val-Borro. High-precision photometry for K2 Campaign 1. *MNRAS*, 454:4159–4171, December 2015. doi: 10.1093/mnras/stv2257.

B. Hubert, J.-C. Grard, D. S. Evans, M. Meurant, S. B. Mende, H. U. Frey, and T. J. Immel. Total electron and proton energy input during auroral substorms: Remote sensing with image-fuv. *Journal of Geophysical Research: Space Physics*, 107(A8):SMP 15–1–SMP 15–12, 2002. ISSN 2156-2202. doi: 10.1029/2001JA009229. URL `http://dx.doi.org/10.1029/2001JA009229`.

G. B. Hughes and M. Chraibi. Calculating ellipse overlap areas. *ArXiv e-prints*, June 2011.

D. M. Hunten. Some photometric observations of auroral spectra. *Journal of Atmospheric and Terrestrial Physics*, 7:141–151, 1955. doi: 10.1016/0021-9169(55)90121-5.

D. M. Hunten. The Escape of Light Gases from Planetary Atmospheres. *Journal of Atmospheric Sciences*, 30:1481–1494, November 1973. doi: 10.1175/1520-0469(1973)030⟨1481:TEOLGF⟩2.0.CO;2.

D. M. Hunten. Thermal and nonthermal escape mechanisms for terrestrial bodies. *Planet. Space Sci.*, 30:773–783, August 1982. doi: 10.1016/0032-0633(82)90110-6.

D. M. Hunten, H. N. Rundle, G. G. Shepherd, and A. Vallance Jones. Optical upper atmospheric investigations at the University of Saskatchewan. *ApOpt*, 6:1609, October 1967. doi: 10.1364/AO.6.001609.

D. M. Hunten, R. O. Pepin, and J. C. G. Walker. Mass fractionation in hydrodynamic escape. *Icarus*, 69:532–549, March 1987. doi: 10.1016/0019-1035(87)90022-4.

S. Ida and D. N. C. Lin. Toward a Deterministic Model of Planetary Formation. IV. Effects of Type I Migration. *ApJ*, 673:487–501, January 2008a. doi: 10.1086/523754.

S. Ida and D. N. C. Lin. Toward a Deterministic Model of Planetary Formation. V. Accumulation Near the Ice Line and Super-Earths. *ApJ*, 685:584–595, September 2008b. doi: 10.1086/590401.

A. P. Ingersoll. The Runaway Greenhouse: A History of Water on Venus. *Journal of Atmospheric Sciences*, 26:1191–1198, November 1969. doi: 10.1175/1520-0469(1969)026⟨1191:TRGAHO⟩2.0.CO;2.

A. Izidoro, M. Ogihara, S. N. Raymond, A. Morbidelli, A. Pierens, B. Bitsch, C. Cossou, and F. Hersant. Breaking the Chains: Hot Super-Earth systems from migration and disruption of compact resonant chains. *ArXiv e-prints*, March 2017.

A. P. Jackson, T. A. Davis, and P. J. Wheatley. The coronal X-ray-age relation and its implications for the evaporation of exoplanets. *MNRAS*, 422:2024–2043, May 2012. doi: 10.1111/j.1365-2966.2012.20657.x.

B. Jackson, R. Barnes, and R. Greenberg. Tidal heating of terrestrial extrasolar planets and implications for their habitability. *MNRAS*, 391:237–245, November 2008. doi: 10.1111/j.1365-2966.2008.13868.x.

B. Jackson, R. Greenberg, and R. Barnes. Tidal evolution of close-in extrasolar planets. *The Astrophysical Journal*, 678(2):1396–1406, may 2008a. doi: 10.1086/529187. URL `https://doi.org/10.1086/529187`.

B. Jackson, R. Greenberg, and R. Barnes. Tidal heating of extrasolar planets. *The Astrophysical Journal*, 681(2):1631–1638, jul 2008b. doi: 10.1086/587641. URL `https://doi.org/10.1086/587641`.

B. Jackson, N. Miller, R. Barnes, S. N. Raymond, J. J. Fortney, and R. Greenberg. The roles of tidal evolution and evaporative mass loss in the origin of CoRoT-7 b. *MNRAS*, 407:910–922, September 2010. doi: 10.1111/j.1365-2966.2010.17012.x.

J. H. Jeans. *The dynamical theory of gases*. University Press, 1925. URL `http://books.google.com/books?id=hkt-AAAAIAAJ`.

M. Johns, P. McCarthy, K. Raybould, A. Bouchez, A. Farahani, J. Filgueira, G. Jacoby, S. Shectman, and M. Sheehan. Giant Magellan Telescope: overview. In *Ground-based and Airborne Telescopes IV*, volume 8444 of *Proc. SPIE*, page 84441H, September 2012. doi: 10.1117/12.926716.

J. A. Johnson. JWST Exoplanet Characterization: Big Opportunities for Small Planets Around Small Stars. In *American Astronomical Society Meeting Abstracts*, volume 223 of *American Astronomical Society Meeting Abstracts*, page #314.05, January 2014.

C. Jones. Bodies in Transit. *New York Times Magazine*, page MM91, December 2016. URL `https://nyti.ms/2jEGgU7`.

M. I. Jones, R. Brahm, N. Espinoza, A. Jordan, F. Rojas, M. Rabus, H. Drass, A. Zapata, M. G. Soto, J. S. Jenkins, M. Vuckovic, S. Ciceri, and P. Sarkis. A hot Saturn on an eccentric orbit around the giant star EPIC228754001. *ArXiv e-prints*, July 2017.

L. Kaltenegger, D. Sasselov, and S. Rugheimer. Water-planets in the Habitable Zone: Atmospheric Chemistry, Observable Features, and the Case of Kepler-62e and -62f. *ApJL*, 775:L47, October 2013. doi: 10.1088/2041-8205/775/2/L47.

S. R. Kane and D. M. Gelino. Decoupling Phase Variations in Multi-planet Systems. *ApJ*, 762:129, January 2013. doi: 10.1088/0004-637X/762/2/129.

S. R. Kane, M. L. Hill, J. F. Kasting, R. K. Kopparapu, E. V. Quintana, T. Barclay, N. M. Batalha, W. J. Borucki, D. R. Ciardi, N. Haghighipour, N. R. Hinkel, L. Kaltenegger, F. Selsis, and G. Torres. A Catalog of Kepler Habitable Zone Exoplanet Candidates. *ApJ*, 830:1, October 2016. doi: 10.3847/0004-637X/830/1/1.

J. F. Kasting. Runaway and moist greenhouse atmospheres and the evolution of earth and Venus. *Icarus*, 74:472–494, June 1988. doi: 10.1016/0019-1035(88)90116-9.

J. F. Kasting and J. B. Pollack. Loss of water from Venus. I - Hydrodynamic escape of hydrogen. *Icarus*, 53:479–508, March 1983. doi: 10.1016/0019-1035(83)90212-9.

J. F. Kasting, J. B. Pollack, and T. P. Ackerman. Response of earth's atmosphere to increases in solar flux and implications for loss of water from Venus. *Icarus*, 57:335–355, March 1984. doi: 10.1016/0019-1035(84)90122-2.

J. F. Kasting, D. P. Whitmire, and R. T. Reynolds. Habitable Zones around Main Sequence Stars. *Icarus*, 101:108–128, January 1993. doi: 10.1006/icar.1993.1010.

H. Kawahara, N. Murakami, T. Matsuo, and T. Kotani. Spectroscopic Coronagraphy for Planetary Radial Velocimetry of Exoplanets. *ApJs*, 212:27, June 2014. doi: 10.1088/0067-0049/212/2/27.

W. Kent. Barnard's Wobble. *Swarthmore College Bulletin*, XCVIII(4):28–31, March 2001. URL `http://bulletin.swarthmore.edu/bulletin-issue-archive/wp-content/archived_issues_pdf/Bulletin_2001_03.pdf`.

M. L. Khodachenko, I. Ribas, H. Lammer, J.-M. Grießmeier, M. Leitner, F. Selsis, C. Eiroa, A. Hanslmeier, H. K. Biernat, C. J. Farrugia, et al. Coronal mass ejection (CME) activity of low mass M stars as an important factor for the habitability of terrestrial exoplanets. I. CME impact on expected magnetospheres of Earth-like exoplanets in close-in habitable zones. *Astrobiology*, 7(1):167–184, 2007.

D. M. Kipping. LUNA: An algorithm for generating dynamic planet-moon transits. *Monthly Notices of the Royal Astronomical Society*, pages no–no, jul 2011. doi: 10.1111/j.1365-2966.2011.19086.x. URL `https://doi.org/10.1111/j.1365-2966.2011.19086.x`.

K. G. Kislyakova, H. Lammer, M. Holmström, M. Panchenko, P. Odert, N. V. Erkaev, M. Leitzinger, M. L. Khodachenko, Y. N. Kulikov, M. Güdel, and A. Hanslmeier. XUV-Exposed, Non-Hydrostatic Hydrogen-Rich Upper Atmospheres of Terrestrial Planets. Part II: Hydrogen Coronae and Ion Escape. *Astrobiology*, 13:1030–1048, November 2013. doi: 10.1089/ast.2012.0958.

K. G. Kislyakova, C. P. Johnstone, P. Odert, N. V. Erkaev, H. Lammer, T. Lüftinger, M. Holmström, M. L. Khodachenko, and M. Güdel. Stellar wind interaction and pick-up ion escape of the Kepler-11 "super-Earths". *A&A*, 562:A116, February 2014. doi: 10.1051/0004-6361/201322933.

M. G. Kivelson and C. T. Russell. *Introduction to Space Physics*. April 1995.

T. Kleine, M. Touboul, B. Bourdon, F. Nimmo, K. Mezger, H. Palme, S. B. Jacobsen, Q.-Z. Yin, and A. N. Halliday. Hf-W chronology of the accretion and early evolution of asteroids and terrestrial planets. *Geochim. Cosmochim. Acta*, 73:5150–5188, September 2009. doi: 10.1016/j.gca.2008.11.047.

H. A. Knutson, D. Charbonneau, L. E. Allen, J. J. Fortney, E. Agol, N. B. Cowan, A. P. Showman, C. S. Cooper, and S. T. Megeath. A map of the day–night contrast of the extrasolar planet HD 189733b. *Nature*, 447(7141):183–186, may 2007. doi: 10.1038/nature05782. URL `http://dx.doi.org/10.1038/nature05782`.

H. A. Knutson, D. Charbonneau, L. E. Allen, A. Burrows, and S. T. Megeath. The 3.6-8.0 μm Broadband Emission Spectrum of HD 209458b: Evidence for an Atmospheric Temperature Inversion. *ApJ*, 673:526-531, January 2008. doi: 10.1086/523894.

R. K. Kopparapu. A Revised Estimate of the Occurrence Rate of Terrestrial Planets in the Habitable Zones around Kepler M-dwarfs. *ApJL*, 767:L8, April 2013. doi: 10.1088/2041-8205/767/1/L8.

R. K. Kopparapu, R. Ramirez, J. F. Kasting, V. Eymet, T. D. Robinson, S. Mahadevan, R. C. Terrien, S. Domagal-Goldman, V. Meadows, and R. Deshpande. Habitable Zones around Main-sequence Stars: New Estimates. *ApJ*, 765:131, March 2013. doi: 10.1088/0004-637X/765/2/131.

R. K. Kopparapu, R. M. Ramirez, J. SchottelKotte, J. F. Kasting, S. Domagal-Goldman, and V. Eymet. Habitable Zones around Main-sequence Stars: Dependence on Planetary Mass. *ApJL*, 787:L29, June 2014. doi: 10.1088/2041-8205/787/2/L29.

R. K. Kopparapu, E. T. Wolf, J. Haqq-Misra, J. Yang, J. F. Kasting, V. Meadows, R. Terrien, and S. Mahadevan. The Inner Edge of the Habitable Zone for Synchronously Rotating Planets around Low-mass Stars Using General Circulation Models. *ApJ*, 819:84, March 2016. doi: 10.3847/0004-637X/819/1/84.

T. T. Koskinen, M. J. Harris, R. V. Yelle, and P. Lavvas. The escape of heavy atoms from the ionosphere of HD209458b. I. A photochemical-dynamical model of the thermosphere. *Icarus*, 226:1678–1694, November 2013a. doi: 10.1016/j.icarus.2012.09.027.

T. T. Koskinen, R. V. Yelle, M. J. Harris, and P. Lavvas. The escape of heavy atoms from the ionosphere of HD209458b. II. Interpretation of the observations. *Icarus*, 226:1695–1708, November 2013b. doi: 10.1016/j.icarus.2012.09.026.

C. Kouveliotou, E. Agol, N. Batalha, J. Bean, M. Bentz, N. Cornish, A. Dressler, E. Figueroa-Feliciano, S. Gaudi, O. Guyon, D. Hartmann, J. Kalirai, M. Niemack, F. Ozel, C. Reynolds, A. Roberge, K. S. A. Straughn, D. Weinberg, and J. Zmuidzinas. Enduring Quests-Daring Visions (NASA Astrophysics in the Next Three Decades). *ArXiv e-prints, arXiv:1401.3741*, January 2014.

G. Kovács, S. Zucker, and T. Mazeh. A box-fitting algorithm in the search for periodic transits. *A&A*, 391:369, January 2002.

L. Kreidberg. batman: BAsic Transit Model cAlculatioN in Python. *PASP*, 127:1161, November 2015. doi: 10.1086/683602.

L. Kreidberg and A. Loeb. Prospects for Characterizing the Atmosphere of Proxima Centauri b. *ArXiv e-prints, arXiv:1608.07345*, August 2016.

E. Kruse, R. Luger, and E. Agol. Detection of New Planets With *K2*. 2017, in prep.

M. J. Kuchner. Volatile-rich Earth-Mass Planets in the Habitable Zone. *ApJL*, 596:L105–L108, October 2003. doi: 10.1086/378397.

Y. N. Kulikov, H. Lammer, H. I. M. Lichtenegger, N. Terada, I. Ribas, C. Kolb, D. Langmayr, R. Lundin, E. F. Guinan, S. Barabash, and H. K. Biernat. Atmospheric and water loss from early Venus. *Planet. Space Sci.*, 54:1425–1444, November 2006. doi: 10.1016/j.pss.2006.04.021.

J. R. Kulow, K. France, J. Linsky, and R. O. P. Loyd. Lyα Transit Spectroscopy and the Neutral Hydrogen Tail of the Hot Neptune GJ 436b. *ApJ*, 786:132, May 2014. doi: 10.1088/0004-637X/786/2/132.

M. Kürster, M. Endl, F. Rouesnel, S. Els, A. Kaufer, S. Brillant, A. P. Hatzes, S. H. Saar, and W. D. Cochran. The low-level radial velocity variability in Barnard's star (= GJ 699). Secular acceleration, indications for convective redshift, and planet mass limits. *A&A*, 403:1077–1087, June 2003. doi: 10.1051/0004-6361:20030396.

W. S. Kurth, D. A. Gurnett, J. T. Clarke, P. Zarka, M. D. Desch, M. L. Kaiser, B. Cecconi, A. Lecacheux, W. M. Farrell, P. Galopeau, J.-C. Gérard, D. Grodent, R. Prangé, M. K. Dougherty, and F. J. Crary. An Earth-like correspondence between Saturn's auroral features and radio emission. *Nature*, 433:722–725, February 2005. doi: 10.1038/nature03334.

H. Lammer. *Origin and Evolution of Planetary Atmospheres*. Springer, Berlin, 2013. doi: 10.1007/978-3-642-32087-3.

H. Lammer, F. Selsis, I. Ribas, E. F. Guinan, S. J. Bauer, and W. W. Weiss. Atmospheric Loss of Exoplanets Resulting from Stellar X-Ray and Extreme-Ultraviolet Heating. *ApJL*, 598:L121–L124, December 2003. doi: 10.1086/380815.

H. Lammer, H. I. M. Lichtenegger, H. K. Biernat, N. V. Erkaev, I. L. Arshukova, C. Kolb, H. Gunell, A. Lukyanov, M. Holmstrom, S. Barabash, T. L. Zhang, and W. Baumjohann. Loss of hydrogen and oxygen from the upper atmosphere of Venus. *Planet. Space Sci.*, 54:1445–1456, November 2006. doi: 10.1016/j.pss.2006.04.022.

H. Lammer, H. I. M. Lichtenegger, Y. N. Kulikov, J.-M. Grießmeier, N. Terada, N. V. Erkaev, H. K. Biernat, M. L. Khodachenko, I. Ribas, T. Penz, and F. Selsis. Coronal Mass Ejection (CME) Activity of Low Mass M Stars as An Important Factor for The Habitability of Terrestrial Exoplanets. II. CME-Induced Ion Pick Up of Earth-like Exoplanets in Close-In Habitable Zones. *Astrobiology*, 7:185–207, February 2007. doi: 10.1089/ast.2006.0128.

H. Lammer, P. Odert, M. Leitzinger, M. L. Khodachenko, M. Panchenko, Y. N. Kulikov, T. L. Zhang, H. I. M. Lichtenegger, N. V. Erkaev, G. Wuchterl, G. Micela, T. Penz, H. K. Biernat,

J. Weingrill, M. Steller, H. Ottacher, J. Hasiba, and A. Hanslmeier. Determining the mass loss limit for close-in exoplanets: what can we learn from transit observations? *A&A*, 506:399–410, October 2009. doi: 10.1051/0004-6361/200911922.

H. Lammer, V. Eybl, K. G. Kislyakova, J. Weingrill, M. Holmström, M. L. Khodachenko, Y. N. Kulikov, A. Reiners, M. Leitzinger, P. Odert, M. Xiang Grüß, B. Dorner, M. Güdel, and A. Hanslmeier. UV transit observations of EUV-heated expanded thermospheres of Earth-like exoplanets around M-stars: testing atmosphere evolution scenarios. *Ap&SS*, 335:39–50, September 2011a. doi: 10.1007/s10509-011-0646-5.

H. Lammer, K. G. Kislyakova, P. Odert, M. Leitzinger, R. Schwarz, E. Pilat-Lohinger, Y. N. Kulikov, M. L. Khodachenko, M. Güdel, and A. Hanslmeier. Pathways to Earth-Like Atmospheres. Extreme Ultraviolet (EUV)-Powered Escape of Hydrogen-Rich Protoatmospheres. *Origins of Life and Evolution of the Biosphere*, 41:503–522, December 2011b. doi: 10.1007/s11084-012-9264-7.

H. Lammer, N. V. Erkaev, P. Odert, K. G. Kislyakova, M. Leitzinger, and M. L. Khodachenko. Probing the blow-off criteria of hydrogen-rich 'super-Earths'. *MNRAS*, 430:1247–1256, April 2013. doi: 10.1093/mnras/sts705.

H. Lammer, A. Stökl, N. V. Erkaev, E. A. Dorfi, P. Odert, M. Güdel, Y. N. Kulikov, K. G. Kislyakova, and M. Leitzinger. Origin and loss of nebula-captured hydrogen envelopes from 'sub'- to 'super-Earths' in the habitable zone of Sun-like stars. *MNRAS*, 439:3225–3238, April 2014. doi: 10.1093/mnras/stu085.

T. Lebrun, H. Massol, E. ChassefièRe, A. Davaille, E. Marcq, P. Sarda, F. Leblanc, and G. Brandeis. Thermal evolution of an early magma ocean in interaction with the atmosphere. *Journal of Geophysical Research (Planets)*, 118:1155–1176, June 2013. doi: 10.1002/jgre.20068.

A. Lecavelier Des Etangs, D. Ehrenreich, A. Vidal-Madjar, G. E. Ballester, J.-M. Désert, R. Ferlet, G. Hébrard, D. K. Sing, K.-O. Tchakoumegni, and S. Udry. Evaporation of the planet HD 189733b observed in H I Lyman-α. *A&A*, 514:A72, May 2010. doi: 10.1051/0004-6361/200913347.

J. Leconte, G. Chabrier, I. Baraffe, and B. Levrard. Is tidal heating sufficient to explain bloated exoplanets? Consistent calculations accounting for finite initial eccentricity. *A&A*, 516:A64, June 2010. doi: 10.1051/0004-6361/201014337.

J. Leconte, F. Forget, B. Charnay, R. Wordsworth, and A. Pottier. Increased insolation threshold for runaway greenhouse processes on Earth-like planets. *Nature*, 504:268–271, December 2013. doi: 10.1038/nature12827.

J. Leconte, H. Wu, K. Menou, and N. Murray. Asynchronous rotation of earth-mass planets in the habitable zone of lower-mass stars. *Science*, 347(6222):632–635, jan 2015. doi: 10.1126/science.1258686. URL `https://doi.org/10.1126/science.1258686`.

C. Lécuyer and Y. Ricard. Long-term fluxes and budget of ferric iron: implication for the redox states of the Earth's mantle and atmosphere. *Earth and Planetary Science Letters*, 165:197–211, January 1999. doi: 10.1016/S0012-821X(98)00267-2.

A. Léger, F. Selsis, C. Sotin, T. Guillot, D. Despois, D. Mawet, M. Ollivier, A. Labèque, C. Valette, F. Brachet, B. Chazelas, and H. Lammer. A new family of planets? "Ocean-Planets". *Icarus*, 169:499–504, June 2004. doi: 10.1016/j.icarus.2004.01.001.

Owen R. Lehmer and David C. Catling. Rocky worlds limited to 1.8 earth radii by atmospheric escape during a star's extreme uv saturation. *ApJ*, 845(2):130, 2017. URL `http://stacks.iop.org/0004-637X/845/i=2/a=130`.

M. Leitzinger, P. Odert, Y. N. Kulikov, H. Lammer, G. Wuchterl, T. Penz, M. G. Guarcello, G. Micela, M. L. Khodachenko, J. Weingrill, A. Hanslmeier, H. K. Biernat, and J. Schneider. Could CoRoT-7b and Kepler-10b be remnants of evaporated gas or ice giants? *Planet. Space Sci.*, 59:1472–1481, October 2011. doi: 10.1016/j.pss.2011.06.003.

A. Levi, D. Sasselov, and M. Podolak. Volatile Transport inside Super-Earths by Entrapment in the Water-ice Matrix. *ApJ*, 769:29, May 2013. doi: 10.1088/0004-637X/769/1/29.

S. Levitus. *Climatological atlas of the world ocean*. U.S. Dept. of Commerce, National Oceanic and Atmospheric Administration, Rockville, MD, 1982.

H. I. M. Lichtenegger, H. Lammer, J.-M. Grießmeier, Y. N. Kulikov, P. von Paris, W. Hausleitner, S. Krauss, and H. Rauer. Aeronomical evidence for higher CO_2 levels during Earth's Hadean epoch. *Icarus*, 210:1–7, November 2010. doi: 10.1016/j.icarus.2010.06.042.

J. Lilensten, D. Bernard, M. Barthélémy, G. Gronoff, C. Simon Wedlund, and A. Opitz. Prediction of blue, red and green aurorae at Mars. *P&SS*, 115:48–56, September 2015. doi: 10.1016/j.pss.2015.04.015.

J. J. Lissauer. Planets Formed in Habitable Zones of M Dwarf Stars Probably Are Deficient in Volatiles. *ApJL*, 660:L149–L152, May 2007. doi: 10.1086/518121.

J. J. Lissauer, D. Ragozzine, D. C. Fabrycky, J. H. Steffen, E. B. Ford, J. M. Jenkins, A. Shporer, M. J. Holman, J. F. Rowe, E. V. Quintana, N. M. Batalha, W. J. Borucki, S. T. Bryson, D. A. Caldwell, J. A. Carter, D. Ciardi, E. W. Dunham, J. J. Fortney, T. N. Gautier, III, S. B. Howell, D. G. Koch, D. W. Latham, G. W. Marcy, R. C. Morehead, and D. Sasselov. Architecture and Dynamics of Kepler's Candidate Multiple Transiting Planet Systems. *ApJs*, 197:8, November 2011. doi: 10.1088/0067-0049/197/1/8.

Y. Lithwick and Y. Wu. Resonant Repulsion of Kepler Planet Pairs. *ApJL*, 756:L11, September 2012. doi: 10.1088/2041-8205/756/1/L11.

Y. Lithwick, J. Xie, and Y. Wu. Extracting Planet Mass and Eccentricity from TTV Data. *ApJ*, 761:122, December 2012. doi: 10.1088/0004-637X/761/2/122.

E. D. Lopez and J. J. Fortney. The Role of Core Mass in Controlling Evaporation: The Kepler Radius Distribution and the Kepler-36 Density Dichotomy. *ApJ*, 776:2, October 2013. doi: 10.1088/0004-637X/776/1/2.

E. D. Lopez and J. J. Fortney. Understanding the Mass-Radius Relation for Sub-neptunes: Radius as a Proxy for Composition. *ApJ*, 792:1, September 2014. doi: 10.1088/0004-637X/792/1/1.

E. D. Lopez, J. J. Fortney, and N. Miller. How Thermal Evolution and Mass-loss Sculpt Populations of Super-Earths and Sub-Neptunes: Application to the Kepler-11 System and Beyond. *ApJ*, 761: 59, December 2012. doi: 10.1088/0004-637X/761/1/59.

C. Lovis and D. Fischer. *Radial Velocity Techniques for Exoplanets*, pages 27–53. December 2010.

C. Lovis, I. Snellen, D. Mouillet, F. Pepe, F. Wildi, N. Astudillo-Defru, J.-L. Beuzit, X. Bonfils, A. Cheetham, U. Conod, X. Delfosse, D. Ehrenreich, P. Figueira, T. Forveille, J. H. C. Martins, S. Quanz, N. C. Santos, H.-M. Schmid, D. Ségransan, and S. Udry. Atmospheric characterization of Proxima b by coupling the SPHERE high-contrast imager to the ESPRESSO spectrograph. *ArXiv e-prints, arXiv:1609.03082*, September 2016.

R. Luger and R. Barnes. Extreme Water Loss and Abiotic O2Buildup on Planets Throughout the Habitable Zones of M Dwarfs. *Astrobiology*, 15:119–143, February 2015. doi: 10.1089/ast.2014.1231.

R. Luger, R. Barnes, E. Lopez, J. Fortney, B. Jackson, and V. Meadows. Habitable Evaporated Cores: Transforming Mini-Neptunes into Super-Earths in the Habitable Zones of M Dwarfs. *Astrobiology*, 15:57–88, January 2015a. doi: 10.1089/ast.2014.1215.

R. Luger, R. Barnes, E. Lopez, J. Fortney, B. Jackson, and V. Meadows. Habitable Evaporated Cores: Transforming Mini-Neptunes into Super-Earths in the Habitable Zones of M Dwarfs. *Astrobiology*, 15:57–88, January 2015b. doi: 10.1089/ast.2014.1215.

R. Luger, E. Agol, E. Kruse, R. Barnes, A. Becker, D. Foreman-Mackey, and D. Deming. EVEREST: Pixel Level Decorrelation of K2 Light Curves. *AJ*, 152:100, October 2016. doi: 10.3847/0004-6256/152/4/100.

R. Luger, M. Sestovic, E. Kruse, S. L. Grimm, B.-O. Demory, E. Agol, E. Bolmont, D. Fabrycky, C. S. Fernandes, V. Van Grootel, A. Burgasser, M. Gillon, J. G. Ingalls, E. Jehin, S. N. Raymond, F. Selsis, A. H. M. J. Triaud, T. Barclay, G. Barentsen, S. B. Howell, L. Delrez, J. de Wit, D. Foreman-Mackey, D. L. Holdsworth, J. Leconte, S. Lederer, M. Turbet, Y. Almleaky, Z. Benkhaldoun, P. Magain, B. M. Morris, K. Heng, and D. Queloz. A seven-planet resonant chain in TRAPPIST-1. *Nature Astronomy*, 1:0129, June 2017. doi: 10.1038/s41550-017-0129.

M. N. Lund, R. Handberg, G. R. Davies, W. J. Chaplin, and C. D. Jones. K2P^2: A Photometry Pipeline for the K2 Mission. *ApJ*, 806:30, June 2015. doi: 10.1088/0004-637X/806/1/30.

M. G. MacDonald, D. Ragozzine, D. C. Fabrycky, E. B. Ford, M. J. Holman, H. T. Isaacson, J. J. Lissauer, E. D. Lopez, T. Mazeh, L. Rogers, J. F. Rowe, J. H. Steffen, and G. Torres. A Dynamical Analysis of the Kepler-80 System of Five Transiting Planets. *AJ*, 152:105, October 2016. doi: 10.3847/0004-6256/152/4/105.

M. G. MacDonald, D. Ragozzine, D. C. Fabrycky, E. B. Ford, M. J. Holman, H. T. Isaacson, J. J. Lissauer, E. D. Lopez, T. Mazeh, L. Rogers, J. F. Rowe, J. H. Steffen, and G. Torres. A dynamical analysis of the Kepler-80 system of five transiting planets. *AJ*, 152(4):105, oct 2016. doi: 10.3847/0004-6256/152/4/105. URL `https://doi.org/10.3847%2F0004-6256%2F152%2F4%2F105`.

B. Macintosh, J. R. Graham, P. Ingraham, Q. Konopacky, C. Marois, M. Perrin, L. Poyneer, B. Bauman, T. Barman, A. S. Burrows, A. Cardwell, J. Chilcote, R. J. De Rosa, D. Dillon, R. Doyon, J. Dunn, D. Erikson, M. P. Fitzgerald, D. Gavel, S. Goodsell, M. Hartung, P. Hibon, P. Kalas, J. Larkin, J. Maire, F. Marchis, M. S. Marley, J. McBride, M. Millar-Blanchaer, K. Morzinski, A. Norton, B. R. Oppenheimer, D. Palmer, J. Patience, L. Pueyo, F. Rantakyro, N. Sadakuni, L. Saddlemyer, D. Savransky, A. Serio, R. Soummer, A. Sivaramakrishnan, I. Song, S. Thomas, J. K. Wallace, S. Wiktorowicz, and S. Wolff. First light of the gemini planet imager. *Proceedings of the National Academy of Sciences*, 111(35):12661–12666, may 2014. doi: 10.1073/pnas.1304215111. URL `http://dx.doi.org/10.1073/pnas.1304215111`.

D. J. C. Mackay. *Information Theory, Inference and Learning Algorithms*. October 2003.

S. J. Mackwell, M. E. Zimmerman, and D. L. Kohlstedt. High-temperature deformation of dry diabase with application to tectonics on Venus. *J. Geophys. Res.*, 103:975, January 1998. doi: 10.1029/97JB02671.

C. Majeau, E. Agol, and N. B. Cowan. A two-dimensional infrared map of the extrasolar planet HD 189733b. *The Astrophysical Journal*, 747(2):L20, feb 2012. doi: 10.1088/2041-8205/747/2/l20. URL `https://doi.org/10.1088/2041-8205/747/2/l20`.

K. Mandel and E. Agol. Analytic Light Curves for Planetary Transit Searches. *ApJL*, 580:L171–L175, December 2002. doi: 10.1086/345520.

A. W. Mann, E. Gaidos, A. Vanderburg, A. C. Rizzuto, M. Ansdell, J. V. Medina, G. N. Mace, A. L. Kraus, and K. R. Sokal. Zodiacal Exoplanets in Time (ZEIT). IV. Seven Transiting Planets in the Praesepe Cluster. *AJ*, 153:64, February 2017. doi: 10.1088/1361-6528/aa5276.

R. A. Mardling. Long-term tidal evolution of short-period planets with companions. *MNRAS*, 382: 1768–1790, December 2007. doi: 10.1111/j.1365-2966.2007.12500.x.

F. S. Masset, A. Morbidelli, A. Crida, and J. Ferreira. Disk Surface Density Transitions as Protoplanet Traps. *ApJ*, 642:478–487, May 2006. doi: 10.1086/500967.

K. Masuda. Very Low Density Planets around Kepler-51 Revealed with Transit Timing Variations and an Anomaly Similar to a Planet-Planet Eclipse Event. *ApJ*, 783:53, March 2014. doi: 10.1088/0004-637X/783/1/53.

K. Masuda, T. Hirano, A. Taruya, M. Nagasawa, and Y. Suto. Characterization of the KOI-94 system with transit timing variation analysis: implication for the planet-planet eclipse. *The Astrophysical Journal*, 778(2):185, nov 2013. doi: 10.1088/0004-637x/778/2/185. URL `https://doi.org/10.1088/0004-637x/778/2/185`.

T. Matsui and Y. Abe. Evolution of an impact-induced atmosphere and magma ocean on the accreting Earth. *Nature*, 319:303–305, January 1986. doi: 10.1038/319303a0.

Y. Matsumoto, M. Nagasawa, and S. Ida. The orbital stability of planets trapped in the first-order mean-motion resonances. *Icarus*, 221:624–631, November 2012. doi: 10.1016/j.icarus.2012.08.032.

A. S. Maurin, F. Selsis, F. Hersant, and A. Belu. Thermal phase curves of nontransiting terrestrial exoplanets. II. Characterizing airless planets. *A&A*, 538:A95, February 2012. doi: 10.1051/0004-6361/201117054.

M. Mayor and D. Queloz. A Jupiter-mass companion to a solar-type star. *Nature*, 378:355–359, November 1995. doi: 10.1038/378355a0.

B. A. Mazin, B. Bumble, S. R. Meeker, K. O'Brien, S. McHugh, and E. Langman. A superconducting focal plane array for ultraviolet, optical, and near-infrared astrophysics. *Optics Express*, 20:1503, January 2012. doi: 10.1364/OE.20.001503.

B. A. Mazin, S. Meeker, M. Strader, P. Szypryt, A. Walter, C. Bockstiegel, G. Collura, D. Mawet, R. Jensen-Clem, O. Guyon, N. Jovanovic, R. Oppenheimer, and E. Serabyn. MKIDs for Direct Imaging of Exoplanets. In *AAS/Division for Extreme Solar Systems Abstracts*, volume 3 of *AAS/Division for Extreme Solar Systems Abstracts*, page 104.07, December 2015.

V. S. Meadows, G. N. Arney, E. W. Schwieterman, J. Lustig-Yaeger, A. P. Lincowski, T. Robinson, S. D. Domagal-Goldman, R. K. Barnes, D. P. Fleming, R. Deitrick, R. Luger, P. E. Driscoll, T. R. Quinn, and D. Crisp. The Habitability of Proxima Centauri b: II: Environmental States and Observational Discriminants. *ArXiv e-prints, arXiv:1608.08620*, August 2016.

V. S. Meadows, C. T. Reinhard, G. N. Arney, M. N. Parenteau, E. W. Schwieterman, S. D. Domagal-Goldman, A. P. Lincowski, K. R. Stapelfeldt, H. Rauer, S. DasSarma, S. Hegde, N. Narita, R. Deitrick, T. W. Lyons, N. Siegler, and J. Lustig-Yaeger. Exoplanet Biosignatures: Understanding Oxygen as a Biosignature in the Context of Its Environment. *ArXiv e-prints*, May 2017.

H. Menager, M. Barthélemy, T. Koskinen, J. Lilensten, D. Ehrenreich, and C. D. Parkinson. Calculation of the H Lyman α emission of the hot Jupiters HD 209458b and HD 189733b. *Icarus*, 226:1709–1718, November 2013. doi: 10.1016/j.icarus.2013.02.028.

B. Mennesson, S. Gaudi, S. Seager, K. Cahoy, S. Domagal-Goldman, L. Feinberg, O. Guyon, J. Kasdin, C. Marois, D. Mawet, M. Tamura, D. Mouillet, T. Prusti, A. Quirrenbach, T. Robinson, L. Rogers, P. Scowen, R. Somerville, K. Stapelfeldt, D. Stern, M. Still, M. Turnbull,

J. Booth, A. Kiessling, G. Kuan, and K. Warfield. The Habitable Exoplanet (HabEx) Imaging Mission: preliminary science drivers and technical requirements , 2016. URL http://dx.doi.org/10.1117/12.2240457.

S. M. Mills and D. C. Fabrycky. Mass, Density, and Formation Constraints in the Compact, Sub-Earth Kepler-444 System including Two Mars-mass Planets. *The Astrophysical Journal*, 838(1): L11, mar 2017. doi: 10.3847/2041-8213/aa6543. URL https://doi.org/10.3847/2041-8213/aa6543.

S. M. Mills, D. C. Fabrycky, C. Migaszewski, E. B. Ford, E. Petigura, and H. Isaacson. A resonant chain of four transiting, sub-Neptune planets. *Nature*, 533:509–512, May 2016. doi: 10.1038/nature17445.

H. Mizutani, T. Yamamoto, and A. Fujimura. A new scaling law of the planetary magnetic fields. *Advances in Space Research*, 12(8):265–279, 1992.

B. T. Montet, T. D. Morton, D. Foreman-Mackey, J. A. Johnson, D. W. Hogg, B. P. Bowler, D. W. Latham, A. Bieryla, and A. W. Mann. Stellar and Planetary Properties of K2 Campaign 1 Candidates and Validation of 17 Planets, Including a Planet Receiving Earth-like Insolation. *ApJ*, 809:25, August 2015. doi: 10.1088/0004-637X/809/1/25.

A. Morbidelli, J. Chambers, J. I. Lunine, J. M. Petit, F. Robert, G. B. Valsecchi, and K. E. Cyr. Source regions and time scales for the delivery of water to Earth. *Meteoritics and Planetary Science*, 35:1309–1320, November 2000. doi: 10.1111/j.1945-5100.2000.tb01518.x.

C. Mordasini, Y. Alibert, C. Georgy, K.-M. Dittkrist, H. Klahr, and T. Henning. Characterization of exoplanets from their formation. II. The planetary mass-radius relationship. *A&A*, 547:A112, November 2012a. doi: 10.1051/0004-6361/201118464.

C. Mordasini, Y. Alibert, H. Klahr, and T. Henning. Characterization of exoplanets from their formation. I. Models of combined planet formation and evolution. *A&A*, 547:A111, November 2012b. doi: 10.1051/0004-6361/201118457.

L. Moresi and V. Solomatov. Mantle convection with a brittle lithosphere: thoughts on the global tectonic styles of the Earth and Venus. *Geophysical Journal International*, 133:669–682, June 1998. doi: 10.1046/j.1365-246X.1998.00521.x.

P. S. Muirhead, J. A. Johnson, K. Apps, J. A. Carter, T. D. Morton, D. C. Fabrycky, J. S. Pineda, M. Bottom, B. Rojas-Ayala, E. Schlawin, K. Hamren, K. R. Covey, J. R. Crepp, K. G. Stassun, J. Pepper, L. Hebb, E. N. Kirby, A. W. Howard, H. T. Isaacson, G. W. Marcy, D. Levitan, T. Diaz-Santos, L. Armus, and J. P. Lloyd. Characterizing the Cool KOIs. III. KOI 961: A Small Star with Large Proper Motion and Three Small Planets. *ApJ*, 747:144, March 2012. doi: 10.1088/0004-637X/747/2/144.

G. D. Mulders, I. Pascucci, and D. Apai. An Increase in the Mass of Planetary Systems around Lower-mass Stars. *ApJ*, 814:130, December 2015. doi: 10.1088/0004-637X/814/2/130.

B. C. Murray, R. L. Wildey, and J. A. Westphal. Infrared Photometric Mapping of Venus through the 8- to 14-Micron Atmospheric Window. *JGR*, 68:4813, August 1963. doi: 10.1029/JZ068i016p04813.

R. A. Murray-Clay, E. I. Chiang, and N. Murray. Atmospheric Escape From Hot Jupiters. *ApJ*, 693:23–42, March 2009. doi: 10.1088/0004-637X/693/1/23.

S. Nakajima, Y.-Y. Hayashi, and Y. Abe. A study on the 'runaway greenhouse effect' with a one-dimensional radiative-convective equilibrium model. *Journal of Atmospheric Sciences*, 49: 2256–2266, December 1992. doi: 10.1175/1520-0469(1992)049⟨2256:ASOTGE⟩2.0.CO;2.

O. Neron de Surgy and J. Laskar. On the long term evolution of the spin of the Earth. *A&A*, 318: 975–989, February 1997.

D. Nesvorný and D. Vokrouhlický. The Effect of Conjunctions on the Transit Timing Variations of Exoplanets. *ApJ*, 790:58, July 2014. doi: 10.1088/0004-637X/790/1/58.

E. R. Newton, J. Irwin, D. Charbonneau, Z. K. Berta-Thompson, J. A. Dittmann, and A. A. West. The Rotation and Galactic Kinematics of Mid M Dwarfs in the Solar Neighborhood. *ApJ*, 821: 93, April 2016. doi: 10.3847/0004-637X/821/2/93.

M. Ogihara and S. Ida. N-Body Simulations of Planetary Accretion Around M Dwarf Stars. *ApJ*, 699:824–838, July 2009. doi: 10.1088/0004-637X/699/1/824.

J. T. O'Malley-James and L. Kaltenegger. Biofluorescent Worlds: Biological fluorescence as a temporal biosignature for flare star worlds. *ArXiv e-prints, arXiv:1608.06930*, August 2016.

A. Oparin. *The Origin of Life*. Moscow, 1924. Tr. in J. D. Bernal, The Origin of Life, Cleveland: World, 1967.

E. J. Öpik. Selective Escape of Gases? *Geophysical Journal International*, 7:490–506, May 1963. doi: 10.1111/j.1365-246X.1963.tb07091.x.

J. E. Owen and A. P. Jackson. Planetary evaporation by UV & X-ray radiation: basic hydrodynamics. *MNRAS*, 425:2931–2947, October 2012. doi: 10.1111/j.1365-2966.2012.21481.x.

J. E. Owen and S. Mohanty. Habitability of terrestrial-mass planets in the HZ of M Dwarfs - I. H/He-dominated atmospheres. *MNRAS*, 459:4088–4108, July 2016. doi: 10.1093/mnras/stw959.

J. E. Owen and Y. Wu. The evaporation valley in the Kepler planets. *ArXiv e-prints*, May 2017.

A. Pál. Light-curve modelling for mutual transits. *MNRAS*, 420:1630–1635, February 2012. doi: 10.1111/j.1365-2966.2011.20151.x.

E. N. Parker. Hydromagnetic Dynamo Models. *ApJ*, 122:293, September 1955. doi: 10.1086/146087.

E. N. Parker. Dynamics of the interplanetary gas and magnetic fields. *The Astrophysical Journal*, 128:664, 1958.

E. N. Parker. Dynamical Theory of the Solar Wind. *Space Sci. Rev.*, 4:666–708, September 1965. doi: 10.1007/BF00216273.

I. Pascucci, D. Apai, K. Luhman, T. Henning, J. Bouwman, M. R. Meyer, F. Lahuis, and A. Natta. The Different Evolution of Gas and Dust in Disks around Sun-Like and Cool Stars. *ApJ*, 696: 143–159, May 2009. doi: 10.1088/0004-637X/696/1/143.

I. Pascucci, L. Testi, G. J. Herczeg, F. Long, C. F. Manara, N. Hendler, G. D. Mulders, S. Krijt, F. Ciesla, T. Henning, S. Mohanty, E. Drabek-Maunder, D. Apai, L. Szűcs, G. Sacco, and J. Olofsson. A Steeper than Linear Disk Mass-Stellar Mass Scaling Relation. *ApJ*, 831:125, November 2016. doi: 10.3847/0004-637X/831/2/125.

S. J. Peale. A primordial origin of the laplace relation among the galilean satellites. *Science*, 298(5593):593–597, oct 2002. doi: 10.1126/science.1076557. URL `https://doi.org/10.1126%2Fscience.1076557`.

K. Penev, B. Jackson, F. Spada, and N. Thom. Constraining Tidal Dissipation in Stars from the Destruction Rates of Exoplanets. *ApJ*, 751:96, June 2012. doi: 10.1088/0004-637X/751/2/96.

P. Perreault and S. I. Akasofu. A study of geomagnetic storms. *Geophysical Journal International*, 54(3):547–573, 1978.

E. A. Petigura and G. W. Marcy. Identification and Removal of Noise Modes in Kepler Photometry. *PASP*, 124:1073–1082, October 2012. doi: 10.1086/668291.

E. A. Petigura, A. W. Howard, and G. W. Marcy. Prevalence of Earth-size planets orbiting Sun-like stars. *Proceedings of the National Academy of Science*, 110:19273–19278, November 2013. doi: 10.1073/pnas.1319909110.

R. Pierrehumbert and E. Gaidos. Hydrogen Greenhouse Planets Beyond the Habitable Zone. *ApJ*, 734:L13, June 2011. doi: 10.1088/2041-8205/734/1/L13.

R. T. Pierrehumbert. *Principles of Planetary Climate*. University Press, Cambridge, UK, December 2010.

N. Pizzolato, A. Maggio, G. Micela, S. Sciortino, P. Ventura, and F. D'Antona. Determination of convective turnover times in young stars. In R. Pallavicini, G. Micela, and S. Sciortino, editors, *Stellar Clusters and Associations: Convection, Rotation, and Dynamos*, volume 198 of *Astronomical Society of the Pacific Conference Series*, page 71, 2000.

J. B. Pollack, O. Hubickyj, P. Bodenheimer, J. J. Lissauer, M. Podolak, and Y. Greenzweig. Formation of the Giant Planets by Concurrent Accretion of Solids and Gas. *Icarus*, 124:62–85, November 1996. doi: 10.1006/icar.1996.0190.

W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical recipes in C. The art of scientific computing*. Cambridge University Press, 1992.

S. Preusse, A. Kopp, J. Büchner, and U. Motschmann. MHD simulation scenarios of the stellar wind interaction with Hot Jupiter magnetospheres. *Planetary and Space Science*, 55(5):589–597, 2007.

E. V. Quintana, J. M. Jenkins, B. D. Clarke, H. Chandrasekaran, J. D. Twicken, S. D. McCauliff, M. T. Cote, T. C. Klaus, C. Allen, D. A. Caldwell, and S. T. Bryson. Pixel-level calibration in the Kepler Science Operations Center pipeline. In *Software and Cyberinfrastructure for Astronomy*, volume 7740 of *Proceedings of SPIE*, page 77401X, July 2010.

E. V. Quintana, T. Barclay, S. N. Raymond, J. F. Rowe, E. Bolmont, D. A. Caldwell, S. B. Howell, S. R. Kane, D. Huber, J. R. Crepp, J. J. Lissauer, D. R. Ciardi, J. L. Coughlin, M. E. Everett, C. E. Henze, E. Horch, H. Isaacson, E. B. Ford, F. C. Adams, M. Still, R. C. Hunter, B. Quarles, and F. Selsis. An Earth-Sized Planet in the Habitable Zone of a Cool Star. *Science*, 344:277–280, April 2014. doi: 10.1126/science.1249403.

R. R. Rafikov. Atmospheres of Protoplanetary Cores: Critical Mass for Nucleated Instability. *ApJ*, 648:666–682, September 2006. doi: 10.1086/505695.

R. Ragazzoni, D. Magrin, H. Rauer, I. Pagano, V. Nascimbeni, G. Piotto, D. Piazza, P. Levacher, M. Schweitzer, S. Basso, T. Bandy, W. Benz, M. Bergomi, F. Biondi, A. Boerner, F. Borsa, A. Brandeker, M. Brändli, G. Bruno, J. Cabrera, S. Chinellato, T. De Roche, M. Dima, A. Erikson, J. Farinato, M. Munari, M. Ghigo, D. Greggio, M. Gullieuszik, M. Klebor, L. Marafatto, V. Mogulsky, G. Peter, M. Rieder, D. Sicilia, D. Spiga, V. Viotto, M. Wieser, A. M. Heras, P. Gondoin, P. Bodin, and C. Catala. PLATO: a multiple telescope spacecraft for exo-planets hunting. In *Space Telescopes and Instrumentation 2016: Optical, Infrared, and Millimeter Wave*, volume 9904 of *Proc SPIE*, page 990428, July 2016. doi: 10.1117/12.2236094.

D. Ragozzine and M. J. Holman. The Value of Systems with Multiple Transiting Planets. *ArXiv e-prints*, June 2010.

R. M. Ramirez and L. Kaltenegger. A Volcanic Hydrogen Habitable Zone. *ApJL*, 837:L4, March 2017. doi: 10.3847/2041-8213/aa60c8.

C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning*. 2006.

H. Rauer, C. Catala, C. Aerts, T. Appourchaux, W. Benz, A. Brandeker, J. Christensen-Dalsgaard, M. Deleuil, L. Gizon, M.-J. Goupil, M. Güdel, E. Janot-Pacheco, M. Mas-Hesse, I. Pagano, G. Piotto, D. Pollacco, Ċ. Santos, A. Smith, J.-C. Suárez, R. Szabó, S. Udry, V. Adibekyan, Y. Alibert, J.-M. Almenara, P. Amaro-Seoane, M. A.-v. Eiff, M. Asplund, E. Antonello, S. Barnes, F. Baudin, K. Belkacem, M. Bergemann, G. Bihain, A. C. Birch, X. Bonfils, I. Boisse, A. S. Bonomo, F. Borsa, I. M. Brandão, E. Brocato, S. Brun, M. Burleigh, R. Burston, J. Cabrera, S. Cassisi, W. Chaplin, S. Charpinet, C. Chiappini, R. P. Church, S. Csizmadia,

M. Cunha, M. Damasso, M. B. Davies, H. J. Deeg, R. F. Díaz, S. Dreizler, C. Dreyer, P. Eggenberger, D. Ehrenreich, P. Eigmüller, A. Erikson, R. Farmer, S. Feltzing, F. de Oliveira Fialho, P. Figueira, T. Forveille, M. Fridlund, R. A. García, P. Giommi, G. Giuffrida, M. Godolt, J. Gomes da Silva, T. Granzer, J. L. Grenfell, A. Grotsch-Noels, E. Günther, C. A. Haswell, A. P. Hatzes, G. Hébrard, S. Hekker, R. Helled, K. Heng, J. M. Jenkins, A. Johansen, M. L. Khodachenko, K. G. Kislyakova, W. Kley, U. Kolb, N. Krivova, F. Kupka, H. Lammer, A. F. Lanza, Y. Lebreton, D. Magrin, P. Marcos-Arenal, P. M. Marrese, J. P. Marques, J. Martins, S. Mathis, S. Mathur, S. Messina, A. Miglio, J. Montalban, M. Montalto, M. J. P. F. G. Monteiro, H. Moradi, E. Moravveji, C. Mordasini, T. Morel, A. Mortier, V. Nascimbeni, R. P. Nelson, M. B. Nielsen, L. Noack, A. J. Norton, A. Ofir, M. Oshagh, R.-M. Ouazzani, P. Pápics, V. C. Parro, P. Petit, B. Plez, E. Poretti, A. Quirrenbach, R. Ragazzoni, G. Raimondo, M. Rainer, D. R. Reese, R. Redmer, S. Reffert, B. Rojas-Ayala, I. W. Roxburgh, S. Salmon, A. Santerne, J. Schneider, J. Schou, S. Schuh, H. Schunker, A. Silva-Valio, R. Silvotti, I. Skillen, I. Snellen, F. Sohl, S. G. Sousa, A. Sozzetti, D. Stello, K. G. Strassmeier, M. Švanda, G. M. Szabó, A. Tkachenko, D. Valencia, V. Van Grootel, S. D. Vauclair, P. Ventura, F. W. Wagner, N. A. Walton, J. Weingrill, S. C. Werner, P. J. Wheatley, and K. Zwintz. The PLATO 2.0 mission. *Experimental Astronomy*, 38:249–330, November 2014. doi: 10.1007/s10686-014-9383-4.

E. Rauscher, K. Menou, S. Seager, D. Deming, J. Y.-K. Cho, and B. M. S. Hansen. Toward Eclipse Mapping of Hot Jupiters. *ApJ*, 664:1199–1209, August 2007. doi: 10.1086/519213.

S. N. Raymond and C. Cossou. No universal minimum-mass extrasolar nebula: evidence against in situ accretion of systems of hot super-Earths. *MNRAS*, 440:L11–L15, May 2014. doi: 10.1093/mnrasl/slu011.

S. N. Raymond, T. Quinn, and J. I. Lunine. High-resolution simulations of the final assembly of Earth-like planets I. Terrestrial accretion and dynamics. *Icarus*, 183:265–282, August 2006. doi: 10.1016/j.icarus.2006.03.011.

S. N. Raymond, J. Scalo, and V. S. Meadows. A Decreased Probability of Habitable Planet Formation around Low-Mass Stars. *ApJ*, 669:606–614, November 2007. doi: 10.1086/521587.

S. N. Raymond, R. Barnes, and A. M. Mandell. Observable consequences of planet formation models in systems with close-in terrestrial planets. *MNRAS*, 384:663–674, February 2008. doi: 10.1111/j.1365-2966.2007.12712.x.

S. N. Raymond, E. Kokubo, A. Morbidelli, R. Morishima, and K. J. Walsh. Terrestrial Planet Formation at Home and Abroad. *ArXiv e-prints*, December 2013.

I. N. Reid and S. L. Hawley. *New Light on Dark Stars: Red Dwarfs, Low-Mass Stars, Brown Stars*. 2005. doi: 10.1007/3-540-27610-6.

I. N. Reid and S. L. Hawley. *New Light on Dark Stars: Red Dwarfs, Low-Mass Stars, Brown Dwarfs*. Springer Praxis Books. Springer London, 2013. ISBN 9781447136637.

H. Rein and S.-F. Liu. REBOUND: an open-source multi-purpose N-body code for collisional dynamics. *A&A*, 537:A128, January 2012. doi: 10.1051/0004-6361/201118085.

H. Rein and D. S. Spiegel. IAS15: a fast, adaptive, high-order integrator for gravitational dynamics, accurate to machine precision over a billion orbits. *MNRAS*, 446:1424–1437, January 2015. doi: 10.1093/mnras/stu2164.

H. Rein and D. Tamayo. WHFAST: a fast and unbiased implementation of a symplectic Wisdom-Holman integrator for long-term gravitational simulations. *MNRAS*, 452:376–388, September 2015. doi: 10.1093/mnras/stv1257.

A. Reiners and G. Basri. The moderate magnetic field of the flare star Proxima Centauri. *A&A*, 489:L45–L48, October 2008. doi: 10.1051/0004-6361:200810491.

P. Riaud and J. Schneider. Improving Earth-like planets' detection with an ELT: the differential radial velocity experiment. *A&A*, 469:355–361, July 2007. doi: 10.1051/0004-6361:20077085.

Á. Ribas, B. Merín, H. Bouy, and L. T. Maud. Disk evolution in the solar neighborhood. I. Disk frequencies from 1 to 100 Myr. *A&A*, 561:A54, January 2014. doi: 10.1051/0004-6361/201322597.

I. Ribas, E. F. Guinan, M. Güdel, and M. Audard. Evolution of the Solar Activity over Time and Effects on Planetary Atmospheres. I. High-Energy Irradiances (1-1700 Å). *ApJ*, 622:680–694, March 2005. doi: 10.1086/427977.

I. Ribas, E. Bolmont, F. Selsis, A. Reiners, J. Leconte, S. N. Raymond, S. G. Engle, E. F. Guinan, J. Morin, M. Turbet, F. Forget, and G. Anglada-Escudé. The habitability of Proxima Centauri b. I. Irradiation, rotation and volatile inventory from formation to the present. *A&A*, 596:A111, December 2016. doi: 10.1051/0004-6361/201629576.

G. R. Ricker, J. N. Winn, R. Vanderspek, D. W. Latham, G. Á. Bakos, J. L. Bean, Z. K. Berta-Thompson, T. M. Brown, L. Buchhave, N. R. Butler, R. P. Butler, W. J. Chaplin, D. Charbonneau, J. Christensen-Dalsgaard, M. Clampin, D. Deming, J. Doty, N. De Lee, C. Dressing, E. W. Dunham, M. Endl, F. Fressin, J. Ge, T. Henning, M. J. Holman, A. W. Howard, S. Ida, J. M. Jenkins, G. Jernigan, J. A. Johnson, L. Kaltenegger, N. Kawai, H. Kjeldsen, G. Laughlin, A. M. Levine, D. Lin, J. J. Lissauer, P. MacQueen, G. Marcy, P. R. McCullough, T. D. Morton, N. Narita, M. Paegert, E. Palle, F. Pepe, J. Pepper, A. Quirrenbach, S. A. Rinehart, D. Sasselov, B. Sato, S. Seager, A. Sozzetti, K. G. Stassun, P. Sullivan, A. Szentgyorgyi, G. Torres, S. Udry, and J. Villasenor. Transiting Exoplanet Survey Satellite (TESS). *Journal of Astronomical Telescopes, Instruments, and Systems*, 1(1):014003, January 2015. doi: 10.1117/1.JATIS.1.1.014003.

H. Ritter. Turning on and off mass transfer in cataclysmic binaries. *A&A*, 202:93–100, August 1988.

E. J. Rivera, G. Laughlin, R. P. Butler, S. S. Vogt, N. Haghighipour, and S. Meschiari. The Lick-Carnegie Exoplanet Survey: A Uranus-mass fourth planet for GJ 876 in an extrasolar

Laplace configuration. *ApJ*, 719(1):890–899, jul 2010. doi: 10.1088/0004-637x/719/1/890. URL https://doi.org/10.1088%2F0004-637x%2F719%2F1%2F890.

P. Robertson, S. Mahadevan, M. Endl, and A. Roy. Stellar activity masquerading as planets in the habitable zone of the M dwarf Gliese 581. *Science*, 345:440–444, July 2014. doi: 10.1126/science.1253253.

T. D. Robinson, K. R. Stapelfeldt, and M. S. Marley. Characterizing Rocky and Gaseous Exoplanets with 2 m Class Space-based Coronagraphs. *PASP*, 128(2):025003, February 2016. doi: 10.1088/1538-3873/128/960/025003.

L. A. Rogers. Most 1.6 Earth-radius Planets are Not Rocky. *ApJ*, 801:41, March 2015. doi: 10.1088/0004-637X/801/1/41.

L. A. Rogers, P. Bodenheimer, J. J. Lissauer, and S. Seager. Formation and Structure of Low-density exo-Neptunes. *ApJ*, 738:59, September 2011. doi: 10.1088/0004-637X/738/1/59.

J. Rosenqvist and E. Chassefière. Inorganic chemistry of O $_2$ in a dense primitive atmosphere. *Planet. Space Sci.*, 43:3–10, February 1995. doi: 10.1016/0032-0633(94)00202-3.

M. Sato and H. Asada. Effects of Mutual Transits by Extrasolar Planet-Companion Systems on Light Curves. *PASJ*, 61:L29–L34, August 2009. doi: 10.1093/pasj/61.4.L29.

M. Sato and H. Asada. Transiting Extrasolar Planet with a Companion: Effects of Orbital Eccentricity and Inclination. *PASJ*, 62:1203–1213, October 2010. doi: 10.1093/pasj/62.5.1203.

J. Saur, F. M. Neubauer, D. F. Strobel, and M. E. Summers. Io's ultraviolet aurora: Remote sensing of Io's interaction. *Geophysical research letters*, 27(18):2893–2896, 2000.

J. Saur, F. M. Neubauer, J. E. P. Connerney, P. Zarka, and M. G. Kivelson. Plasma interaction of Io with its plasma torus. *Jupiter: The Planet, Satellites and Magnetosphere*, 1:537–560, 2004.

A. Savitzky and M. J. E. Golay. Smoothing and differentiation of data by simplified least squares procedures. *Analytical Chemistry*, 36(8):1627–1639, 1964. doi: 10.1021/ac60214a047. URL http://dx.doi.org/10.1021/ac60214a047.

J. Scalo, L. Kaltenegger, A. G. Segura, M. Fridlund, I. Ribas, Y. N. Kulikov, J. L. Grenfell, H. Rauer, P. Odert, M. Leitzinger, F. Selsis, M. L. Khodachenko, C. Eiroa, J. Kasting, and H. Lammer. M Stars as Targets for Terrestrial Exoplanet Searches And Biosignature Detection. *Astrobiology*, 7:85–166, February 2007. doi: 10.1089/ast.2006.0125.

L. Schaefer, R. D. Wordsworth, Z. Berta-Thompson, and D. Sasselov. Predictions of the Atmospheric Composition of GJ 1132b. *ApJ*, 829:63, October 2016. doi: 10.3847/0004-637X/829/2/63.

M. A. Schield. Pressure balance between solar wind and magnetosphere. *Journal of Geophysical Research*, 74(5):1275–1286, 1969.

T. L. Schindler and J. F. Kasting. Synthetic Spectra of Simulated Terrestrial Atmospheres Containing Possible Biomarker Gases. *Icarus*, 145:262–271, May 2000. doi: 10.1006/icar.2000.6340.

K. C. Schlaufman, D. N. C. Lin, and S. Ida. A Population of Very Hot Super-Earths in Multiple-planet Systems Should be Uncovered by Kepler. *ApJL*, 724:L53–L58, November 2010. doi: 10.1088/2041-8205/724/1/L53.

J. William Schopf, Anatoliy B. Kudryavtsev, Andrew D. Czaja, and Abhishek B. Tripathi. Evidence of archean life: Stromatolites and microfossils. *Precambrian Research*, 158(34):141 – 155, 2007. ISSN 0301-9268. doi: http://dx.doi.org/10.1016/j.precamres.2007.04.009.

A. Segura, L. M. Walkowicz, V. Meadows, J. Kasting, and S. Hawley. The Effect of a Strong Stellar Flare on the Atmospheric Chemistry of an Earth-like Planet Orbiting an M Dwarf. *Astrobiology*, 10:751–771, September 2010. doi: 10.1089/ast.2009.0376.

F. Selsis, B. Chazelas, P. Bordé, M. Ollivier, F. Brachet, M. Decaudin, F. Bouchy, D. Ehrenreich, J.-M. Grießmeier, H. Lammer, C. Sotin, O. Grasset, C. Moutou, P. Barge, M. Deleuil, D. Mawet, D. Despois, J. F. Kasting, and A. Léger. Could we identify hot ocean-planets with CoRoT, Kepler and Doppler velocimetry? *Icarus*, 191:453–468, November 2007. doi: 10.1016/j.icarus.2007.04.010.

F. Selsis, R. D. Wordsworth, and F. Forget. Thermal phase curves of nontransiting terrestrial exoplanets. I. Characterizing atmospheres. *A&A*, 532:A1, August 2011. doi: 10.1051/0004-6361/201116654.

F. Selsis, A.-S. Maurin, F. Hersant, J. Leconte, E. Bolmont, S. N. Raymond, and M. Delbo'. The effect of rotation and tidal heating on the thermal lightcurves of super mercuries. *Astronomy & Astrophysics*, 555:A51, jul 2013. doi: 10.1051/0004-6361/201321661. URL `https://doi.org/10.1051/0004-6361/201321661`.

J. F. Sepinsky, B. Willems, and V. Kalogera. Equipotential Surfaces and Lagrangian Points in Nonsynchronous, Eccentric Binary and Planetary Systems. *ApJ*, 660:1624–1635, May 2007. doi: 10.1086/513736.

J. F. Sepinsky, B. Willems, V. Kalogera, and F. A. Rasio. Interacting Binaries with Eccentric Orbits. II. Secular Orbital Evolution due to Non-conservative Mass Transfer. *ApJ*, 702:1387–1392, September 2009. doi: 10.1088/0004-637X/702/2/1387.

S. P. Seth, S. A. Haider, and K. I. Oyama. Photoelectron flux and nightglow emissions of 5577 and 6300 Å due to solar wind electron precipitation in Martian atmosphere. *Journal of Geophysical Research: Space Physics*, 107(A10), 2002.

V. I. Shematovich, D. E. Ionov, and H. Lammer. Heating efficiency in hydrogen-dominated upper atmospheres. *ArXiv e-prints*, September 2014.

E. Sinukoff, A. W. Howard, E. A. Petigura, J. E. Schlieder, I. J. M. Crossfield, D. R. Ciardi, B. J. Fulton, H. Isaacson, K. M. Aller, C. Baranec, C. A. Beichman, B. M. S. Hansen, H. A. Knutson, N. M. Law, M. C. Liu, and R. Riddle. Ten Multi-planet Systems from K2 Campaigns 1 and 2 and the Masses of Two Hot Super-Earths. *ArXiv e-prints*, November 2015.

W. Skidmore, TMT International Science Development Teams, and T. Science Advisory Committee. Thirty Meter Telescope Detailed Science Case: 2015. *Research in Astronomy and Astrophysics*, 15:1945, December 2015. doi: 10.1088/1674-4527/15/12/001.

A. Skumanich. Time Scales for CA II Emission Decay, Rotational Braking, and Lithium Depletion. *ApJ*, 171:565, February 1972. doi: 10.1086/151310.

T. G. Slanger, P. C. Cosby, D. L. Huestis, and T. A. Bida. Discovery of the Atomic Oxygen Green Line in the Venus Night Airglow. *Science*, 291:463–465, January 2001. doi: 10.1126/science.291.5503.463.

T. G. Slanger, D. L. Huestis, P. C. Cosby, N. J. Chanover, and T. A. Bida. The Venus nightglow: Ground-based observations and chemical mechanisms. *Icarus*, 182:1–9, May 2006. doi: 10.1016/j.icarus.2005.12.007.

D. S. Smith, J. Scalo, and J. C. Wheeler. Transport of ionizing radiation in terrestrial-like exoplanet atmospheres. *Icarus*, 171:229–253, September 2004. doi: 10.1016/j.icarus.2004.04.009.

J. C. Smith, M. C. Stumpe, J. E. Van Cleve, J. M. Jenkins, T. S. Barclay, M. N. Fanelli, F. R. Girouard, J. J. Kolodziejczak, S. D. McCauliff, R. L. Morris, and J. D. Twicken. Kepler Presearch Data Conditioning II - A Bayesian Approach to Systematic Error Correction. *PASP*, 124:1000–1014, September 2012. doi: 10.1086/667697.

I. Snellen, R. de Kok, J. L. Birkby, B. Brandl, M. Brogi, C. Keller, M. Kenworthy, H. Schwarz, and R. Stuik. Combining high-dispersion spectroscopy with high contrast imaging: Probing rocky planets around our nearest neighbors. *A&A*, 576:A59, April 2015. doi: 10.1051/0004-6361/201425018.

C. Sotin, O. Grasset, and A. Mocquet. Mass radius curve for extrasolar Earth-like planets and ocean planets. *Icarus*, 191:337–351, November 2007. doi: 10.1016/j.icarus.2007.04.006.

W. B. Sparks and H. C. Ford. Imaging Spectroscopy for Extrasolar Planet Detection. *ApJ*, 578:543–564, October 2002. doi: 10.1086/342401.

J. R. Spencer, K. L. Jessup, M. A. McGrath, G. E. Ballester, and R. Yelle. Discovery of Gaseous S_2 in Io's Pele Plume. *Science*, 288:1208–1210, May 2000. doi: 10.1126/science.288.5469.1208.

J. R. Stauffer, J.-P. Caillault, M. Gagne, C. F. Prosser, and L. W. Hartmann. A deep imaging survey of the Pleiades with ROSAT. *ApJs*, 91:625–657, April 1994. doi: 10.1086/191951.

D. P. Steele and D. J. McEwen. Electron auroral excitation efficiencies and intensity ratios. *Journal of Geophysical Research: Space Physics*, 95(A7):10321–10336, 1990. ISSN 2156-2202. doi: 10.1029/JA095iA07p10321. URL http://dx.doi.org/10.1029/JA095iA07p10321.

B. Stelzer, A. Marino, G. Micela, J. López-Santiago, and C. Liefke. The UV and X-ray activity of the M dwarfs within 10 pc of the Sun. *MNRAS*, 431:2063–2079, May 2013. doi: 10.1093/mnras/stt225.

D. J. Stevenson. Planetary magnetic fields. *Reports on Progress in Physics*, 46(5):555, 1983.

D. J. Stevenson. Life-sustaining planets in interstellar space? *Nature*, 400:32, July 1999. doi: 10.1038/21811.

K. B. Stevenson, J.-M. Désert, M. R. Line, J. L. Bean, J. J. Fortney, A. P. Showman, T. Kataria, L. Kreidberg, P. R. McCullough, G. W. Henry, D. Charbonneau, A. Burrows, S. Seager, N. Madhusudhan, M. H. Williamson, and D. Homeier. Thermal structure of an exoplanet atmosphere from phase-resolved emission spectroscopy. *Science*, 346:838–841, November 2014. doi: 10.1126/science.1256758.

R. Storn and K. Price. Differential evolution – a simple and efficient heuristic for global optimization over continuous spaces. *J. Global Optim.*, 11(4):341–359, 1997. ISSN 1573-2916. doi: 10.1023/A:1008202821328. URL http://dx.doi.org/10.1023/A:1008202821328.

D. J. Strickland, J. H. Hecht, A. B. Christensen, and D. J. McEwen. Thermospheric disturbance recorded by photometers onboard the aria ii rocket. *Journal of Geophysical Research: Space Physics*, 105(A2):2461–2475, 2000. ISSN 2156-2202. doi: 10.1029/1999JA900473. URL http://dx.doi.org/10.1029/1999JA900473.

K. M. Strom, S. E. Strom, S. Edwards, S. Cabrit, and M. F. Skrutskie. Circumstellar material associated with solar-type pre-main-sequence stars - A possible constraint on the timescale for planet building. *AJ*, 97:1451–1470, May 1989. doi: 10.1086/115085.

W. Stubbs. *Historical Works of Gervase of Canterbury*, volume 1. Longman & Company, 1879.

M. C. Stumpe, J. C. Smith, J. E. Van Cleve, J. D. Twicken, T. S. Barclay, M. N. Fanelli, F. R. Girouard, J. M. Jenkins, J. J. Kolodziejczak, S. D. McCauliff, and R. L. Morris. Kepler Presearch Data Conditioning I: Architecture and Algorithms for Error Correction in Kepler Light Curves. *PASP*, 124:985–999, September 2012. doi: 10.1086/667698.

P. W. Sullivan, J. N. Winn, Z. K. Berta-Thompson, D. Charbonneau, D. Deming, C. D. Dressing, D. W. Latham, A. M. Levine, P. R. McCullough, T. Morton, G. R. Ricker, R. Vanderspek, and D. Woods. The Transiting Exoplanet Survey Satellite: Simulations of Planet Detections and Astrophysical False Positives. *ApJ*, 809:77, August 2015. doi: 10.1088/0004-637X/809/1/77.

J. J. Swift, J. A. Johnson, T. D. Morton, J. R. Crepp, B. T. Montet, D. C. Fabrycky, and P. S. Muirhead. Characterizing the Cool KOIs. IV. Kepler-32 as a Prototype for the Formation

of Compact Planetary Systems throughout the Galaxy. *ApJ*, 764:105, February 2013. doi: 10.1088/0004-637X/764/1/105.

O. Tamuz, T. Mazeh, and S. Zucker. Correcting systematic effects in a large set of photometric light curves. *MNRAS*, 356:1466–1470, February 2005. doi: 10.1111/j.1365-2966.2004.08585.x.

C. Terquem and J. C. B. Papaloizou. Migration and the Formation of Systems of Hot Super-Earths and Neptunes. *ApJ*, 654:1110–1120, January 2007. doi: 10.1086/509497.

P. Testa, J. J. Drake, and G. Peres. The density of coronal plasma in active stellar coronae. *The Astrophysical Journal*, 617(1):508, 2004.

F. Tian. Thermal Escape from Super Earth Atmospheres in the Habitable Zones of M Stars. *ApJ*, 703:905–909, September 2009. doi: 10.1088/0004-637X/703/1/905.

F. Tian, O. B. Toon, A. A. Pavlov, and H. De Sterck. A Hydrogen-Rich Early Earth Atmosphere. *Science*, 308:1014–1017, May 2005. doi: 10.1126/science.1106983.

F. Tian, J. F. Kasting, H.-L. Liu, and R. G. Roble. Hydrodynamic planetary thermosphere model: 1. Response of the Earth's thermosphere to extreme solar EUV conditions and the significance of adiabatic cooling. *Journal of Geophysical Research (Planets)*, 113:E05008, May 2008. doi: 10.1029/2007JE002946.

F. Tian, K. France, J. L. Linsky, P. J. D. Mauas, and M. C. Vieytes. High stellar FUV/NUV ratio and oxygen contents in the atmospheres of potentially habitable planets. *Earth and Planetary Science Letters*, 385:22–27, January 2014. doi: 10.1016/j.epsl.2013.10.024.

D. E. Trilling, W. Benz, T. Guillot, J. I. Lunine, W. B. Hubbard, and A. Burrows. Orbital Evolution and Migration of Giant Planets: Modeling Extrasolar Planets. *ApJ*, 500:428, June 1998. doi: 10.1086/305711.

M. Tuomi, H. R. A. Jones, J. R. Barnes, G. Anglada-Escudé, and J. S. Jenkins. Bayesian search for low-mass planets around nearby M dwarfs - estimates for occurrence rate based on global detectability statistics. *MNRAS*, 441:1545–1569, June 2014. doi: 10.1093/mnras/stu358.

M. Turbet, J. Leconte, F. Selsis, E. Bolmont, F. Forget, I. Ribas, S. N. Raymond, and G. Anglada-Escudé. The habitability of Proxima Centauri b. II. Possible climates and observability. *A&A*, 596:A112, December 2016. doi: 10.1051/0004-6361/201629577.

S. Udry, C. Lovis, F. Bouchy, A. Collier Cameron, T. Henning, M. Mayor, F. Pepe, N. Piskunov, D. Pollacco, D. Queloz, A. Quirrenbach, H. Rauer, R. Rebolo, N. C. Santos, I. Snellen, and F. Zerbi. Exoplanet Science with the European Extremely Large Telescope. The Case for Visible and Near-IR Spectroscopy at High Resolution. *ArXiv e-prints, arXiv:1412.1048*, November 2014.

P. van de Kamp. Astrometric study of Barnard's star from plates taken with the 24-inch Sproul refractor. *AJ*, 68:515–521, September 1963. doi: 10.1086/109001.

P. van de Kamp. Alternate dynamical analysis of Barnard's star. *AJ*, 74:757–759, August 1969. doi: 10.1086/110852.

P. van de Kamp. Astrometric study of Barnard's star from plates taken with the Sproul 61-cm refractor. *AJ*, 80:658, August 1975. doi: 10.1086/111791.

P. van de Kamp. Dark companions of stars - Astrometric commentary on the lower end of the Main Sequence. *SSR*, 43:211–327, April 1986.

A. Vanderburg. Reduced Light Curves from Campaign 0 of the K2 Mission. *ArXiv e-prints*, December 2014.

A. Vanderburg and J. A. Johnson. A Technique for Extracting Highly Precise Photometry for the Two-Wheeled Kepler Mission. *PASP*, 126:948–958, October 2014. doi: 10.1086/678764.

V. M. Vasyliunas, J. R. Kan, G. L. Siscoe, and S.-I. Akasofu. Scaling relations governing magnetospheric energy transfer. *Planetary and Space Science*, 30(4):359 – 365, 1982. ISSN 0032-0633. doi: http://dx.doi.org/10.1016/0032-0633(82)90041-1. URL `http://www.sciencedirect.com/science/article/pii/0032063382900411`.

D. Veras and E. Breedt. Eclipse, transit and occultation geometry of planetary systems at exo-syzygy. *Monthly Notices of the Royal Astronomical Society*, 468(3):2672–2683, mar 2017. doi: 10.1093/mnras/stx614. URL `https://doi.org/10.1093/mnras/stx614`.

A. Vidal-Madjar, A. Lecavelier des Etangs, J.-M. Désert, G. E. Ballester, R. Ferlet, G. Hébrard, and M. Mayor. An extended upper atmosphere around the extrasolar planet HD209458b. *Nature*, 422:143–146, March 2003. doi: 10.1038/nature01448.

S. S. Vogt, R. P. Butler, E. J. Rivera, N. Haghighipour, G. W. Henry, and M. H. Williamson. The Lick-Carnegie Exoplanet Survey: A 3.1 $M_{\oplus}$ Planet in the Habitable Zone of the Nearby M3V Star Gliese 581. *ApJ*, 723:954–965, November 2010. doi: 10.1088/0004-637X/723/1/954.

F. M. Walter, A. Brown, R. D. Mathieu, P. C. Myers, and F. J. Vrba. X-ray sources in regions of star formation. III - Naked T Tauri stars associated with the Taurus-Auriga complex. *AJ*, 96: 297–325, July 1988. doi: 10.1086/114809.

C. Wang, J. P. Han, H. Li, Z. Peng, and J. D. Richardson. Solar wind-magnetosphere energy coupling function fitting: Results from a global mhd simulation. *Journal of Geophysical Research: Space Physics*, 119(8):6199–6212, 2014. ISSN 2169-9402. doi: 10.1002/2014JA019834. URL `http://dx.doi.org/10.1002/2014JA019834`.

D. Wang, D. W. Hogg, D. Foreman-Mackey, and B. Schölkopf. A Causal, Data-driven Approach to Modeling the Kepler Data. *PASP*, 128(9):094503, September 2016. doi: 10.1088/1538-3873/128/967/094503.

S. Wang, D.-H. Wu, T. Barclay, and G. P. Laughlin. Updated Masses for the TRAPPIST-1 Planets. *ArXiv e-prints*, April 2017.

W. R. Ward. Protoplanet Migration by Nebula Tides. *Icarus*, 126:261–281, April 1997. doi: 10.1006/icar.1996.5647.

D. Q. Wark. Doppler Widths of the Atomic Oxygen Lines in the Airglow. *ApJ*, 131:491, March 1960. doi: 10.1086/146854.

A. J. Watson, T. M. Donahue, and J. C. G. Walker. The dynamics of a rapidly escaping atmosphere - Applications to the evolution of earth and Venus. *Icarus*, 48:150–166, November 1981. doi: 10.1016/0019-1035(81)90101-9.

R. F. Wendlandt. Oxygen diffusion in basalt and andesite melts: experimental results and discussion of chemical versus tracer diffusion. *Contributions to Mineralogy and Petrology*, 108:463–471, October 1991. doi: 10.1007/BF00303450.

A. A. West, S. L. Hawley, J. J. Bochanski, K. R. Covey, I. N. Reid, S. Dhital, E. J. Hilton, and M. Masuda. Constraining the Age-Activity Relation for Cool Stars: The Sloan Digital Sky Survey Data Release 5 Low-Mass Star Spectroscopic Sample. *AJ*, 135:785–795, March 2008. doi: 10.1088/0004-6256/135/3/785.

D. M. Williams and D. Pollard. Earth-like worlds on eccentric orbits: excursions beyond the habitable zone. *International Journal of Astrobiology*, 1:61–69, January 2002. doi: 10.1017/S1473550402001064.

P. K. G. Williams, D. Charbonneau, C. S. Cooper, A. P. Showman, and J. J. Fortney. Resolving the Surfaces of Extrasolar Planets with Secondary Eclipse Light Curves. *ApJ*, 649:1020–1027, October 2006. doi: 10.1086/506468.

J. N. Winn. *Exoplanet Transits and Occultations*, pages 55–77. University of Arizona Press, December 2010.

A. Wolfgang and E. Lopez. How Rocky Are They? The Composition Distribution of Kepler's Sub-Neptune Planet Candidates within 0.15 AU. *ArXiv e-prints*, September 2014.

A. Wolfgang and E. Lopez. How Rocky Are They? The Composition Distribution of Kepler's Sub-Neptune Planet Candidates within 0.15 AU. *ApJ*, 806:183, June 2015. doi: 10.1088/0004-637X/806/2/183.

B. E. Wood, H.-R. Müller, G. P. Zank, V. V. Izmodenov, and J. L. Linsky. The heliospheric hydrogen wall and astrospheres. *Advances in Space Research*, 34(1):66–73, 2004.

B. E. Wood, H.-R. Müller, G. P. Zank, J. L. Linsky, and S. Redfield. New mass-loss measurements from astrospheric Lyα absorption. *The Astrophysical Journal Letters*, 628(2):L143, 2005.

R. Wordsworth and R. Pierrehumbert. Hydrogen-Nitrogen Greenhouse Warming in Earth's Early Atmosphere. *Science*, 339:64, January 2013a. doi: 10.1126/science.1225759.

R. Wordsworth and R. Pierrehumbert. Abiotic Oxygen-dominated Atmospheres on Terrestrial Habitable Zone Planets. *ApJL*, 785:L20, April 2014. doi: 10.1088/2041-8205/785/2/L20.

R. Wordsworth, Y. Kalugina, S. Lokshtanov, A. Vigasin, B. Ehlmann, J. Head, C. Sanders, and H. Wang. Transient reducing greenhouse warming on early Mars. *ArXiv e-prints*, October 2016.

R. D. Wordsworth and R. T. Pierrehumbert. Water Loss from Terrestrial Planets with CO_2-rich Atmospheres. *ApJ*, 778:154, December 2013b. doi: 10.1088/0004-637X/778/2/154.

N. J. Wright, J. J. Drake, E. E. Mamajek, and G. W. Henry. The Stellar-activity-Rotation Relationship and the Evolution of Stellar Dynamos. *ApJ*, 743:48, December 2011. doi: 10.1088/0004-637X/743/1/48.

J. Yang, N. B. Cowan, and D. S. Abbot. Stabilizing Cloud Feedback Dramatically Expands the Habitable Zone of Tidally Locked Planets. *ApJL*, 771:L45, July 2013. doi: 10.1088/2041-8205/771/2/L45.

J. Yang, G. Boué, D. C. Fabrycky, and D. S. Abbot. Strong Dependence of the Inner Edge of the Habitable Zone on Planetary Rotation Rate. *ApJl*, 787:L2, May 2014. doi: 10.1088/2041-8205/787/1/L2.

R. V. Yelle. Aeronomy of extra-solar giant planets at small orbital distances. *Icarus*, 170:167–179, July 2004. doi: 10.1016/j.icarus.2004.02.008.

K. Zahnle, J. B. Pollack, and J. F. Kasting. Mass fractionation of noble gases in diffusion-limited hydrodynamic hydrogen escape. *Icarus*, 84:502–527, April 1990.

K. J. Zahnle and D. C. Catling. The cosmic shoreline: the evidence that escape determines which planets have atmospheres, and what this may mean for Proxima Centauri b. *ArXiv e-prints*, February 2017.

K. J. Zahnle and J. F. Kasting. Mass fractionation during transonic escape and implications for loss of water from Mars and Venus. *Icarus*, 68:462–480, December 1986. doi: 10.1016/0019-1035(86)90051-5.

K. J. Zahnle, J. F. Kasting, and J. B. Pollack. Evolution of a steam atmosphere during earth's accretion. *Icarus*, 74:62–97, April 1988. doi: 10.1016/0019-1035(88)90031-0.

P. Zarka. Hot Jupiters and Magnetized Stars: Giant Analogs of the Satellite-Jupiter System. In H. O. Rucker, W. Kurth, and G. Mann, editors, *Planetary Radio Emissions VI*, page 543, 2006.

P. Zarka. Plasma interactions of exoplanets with their parent star and associated radio emissions. *P&SS*, 55:598–617, April 2007. doi: 10.1016/j.pss.2006.05.045.

R. T. Zellem, N. K. Lewis, H. A. Knutson, C. A. Griffith, A. P. Showman, J. J. Fortney, N. B. Cowan, E. Agol, A. Burrows, D. Charbonneau, D. Deming, G. Laughlin, and J. Langton. The 4.5 μm Full-orbit Phase Curve of the Hot Jupiter HD 209458b. *ApJ*, 790:53, July 2014. doi: 10.1088/0004-637X/790/1/53.

J. I. Zuluaga and S. Bustamante. Geomagnetic properties of Proxima Centauri b analogues. *ArXiv e-prints, arXiv:1609.00707*, September 2016.