



Theses and Dissertations

2025-04-22

Modeling the Kepler Systems With Multiple Transiting Planets

Daniel Kempton Jones
Brigham Young University

Follow this and additional works at: <https://scholarsarchive.byu.edu/etd>



Part of the [Physical Sciences and Mathematics Commons](#)

BYU ScholarsArchive Citation

Jones, Daniel Kempton, "Modeling the Kepler Systems With Multiple Transiting Planets" (2025). *Theses and Dissertations*. 11258.

<https://scholarsarchive.byu.edu/etd/11258>

This Dissertation is brought to you for free and open access by BYU ScholarsArchive. It has been accepted for inclusion in Theses and Dissertations by an authorized administrator of BYU ScholarsArchive. For more information, please contact ellen_amatangelo@byu.edu.

Modeling the Kepler Systems With Multiple Transiting Planets

Daniel Kempton Jones

A dissertation submitted to the faculty of
Brigham Young University
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

Darin A. Ragozzine, Advisor
Denise Stephens
Eric Hintz
David Neilsen
Benjamin Boizelle

Department of Physics and Astronomy
Brigham Young University

Copyright © 2025 Daniel Kempton Jones

All Rights Reserved

ABSTRACT

Modeling the Kepler Systems With Multiple Transiting Planets

Daniel Kempton Jones

Department of Physics and Astronomy, BYU

Doctor of Philosophy

As the study of exoplanets, planets not in our solar system, has developed, new observational constraints have illuminated the formation and evolution of planetary systems. Among these observations, the true distributions of planetary mass, radius, and density are at the forefront. NASA's Kepler Space Telescope has provided an enormous wealth of data on the exoplanet radius distribution as well as a significant fraction of mass measurements through the detection of planet-planet dynamical interactions. The best method for studying these exoplanetary systems is with a photodynamical model which combines an n-body integrator with lightcurve model to generate synthetic lightcurves. To support photodynamical modeling, we have utilized the PhotoDynamical Multiplanet Model (PhoDyMM) which combines a photodynamical model with a Differential Evolution Markov Chain Monte Carlo (DEMCMC) algorithm for Bayesian parameter inference. PhoDyMM can easily work with arbitrary Kepler systems, enabling the self-consistent analysis of all Kepler systems of multiple exoplanets (multis). Using PhoDyMM, we construct converged posterior distributions for all of the physical and orbital parameters for 661 out of the 719 Kepler multis. This catalog of planetary parameters is known as the Kepler Multis Dynamical Catalog (KMDC) and is the largest, most complete database of planetary parameters to date. Due to its homogeneous construction, the KMDC can readily allow for the extension from studying a single system to accessing trends across the population of Kepler multis. One such analysis that we perform is applying an interior modeling software to a subset of 830 planets across 465 systems. This software is able to construct posterior distributions for a four component interior consisting of an H/He atmosphere, water/ice layer, silicate mantle, and iron core. The results of which were then able to show the probable compositional distribution for a considerable subpopulation of the Kepler multis and support current theories such as the cause of the Kepler Radius Gap. Finally, we also demonstrate that by requiring stable configurations of exoplanetary systems, it is possible to further constrain the physical and orbital parameters of the Kepler multis which can therefore enhance the values given in the KMDC.

Keywords: Kepler, Photodynamical modeling, Stability, Interior modeling, Exoplanets

ACKNOWLEDGMENTS

First and foremost, I wish to acknowledge the efforts of Dr. Darin Ragozzine, without his support and mentoring, these projects would have never been completed. Additionally, I would like to recognize the efforts of past and current members of Dr. Ragozzine's research group which include: Abigail Graham, Burke Boyer, Steven Blodgett, Raymond Kelly, and Jacob Guerrette. Each of these students contributed to this project in a meaningful way, whether it was producing output to be analyzed or generating figures for publication. I would also like to express gratitude to the Office of Research Computing at BYU both for the staff's support in answering questions about software and for the supercomputing resources offered through this office. The photodynamical analysis of all of the Kepler systems with multiple transiting planets required over a million CPU days and was only computationally tractable by using the supercomputer located at BYU. Lastly, this work was funded by the National Science Foundation under Award No. 2143195, this funding made all of the research presented in this dissertation possible.

Contents

Table of Contents	iv
1 Introduction	1
1.1 Exoplanets: Who, What, and Why?	1
1.2 Physical and Orbital Parameters	4
1.3 Transiting Planet Method	6
1.4 The Kepler Space Telescope	11
1.5 Systems of Multiple Transiting Planets	13
1.6 Parameter Degeneracies	18
1.7 Photodynamical Modeling	19
1.8 Long-term Stability	21
1.9 Formation and Evolution	24
1.10 High Performance Computing and Book Keeping	27
1.11 Dissertation Outline	30
2 Photodynamically Modeling the Kepler Multis	31
2.1 Introduction	31
2.2 Methods	35
2.2.1 PhoDyMM	36
2.2.2 Priors	38
2.2.3 Planet Selection	39
2.2.4 Lightcurve Preparation	40
2.2.5 PhoDyMM Input Generation	42
2.2.6 PhoDyMM Output	44
2.2.7 Transformation Matrix	45
2.2.8 metaPhoDyMM	46
2.2.9 Restarting Systems	52
2.2.10 Convergence by Visual Inspection	53
2.2.11 Computational Resources	53
2.3 Results	54
2.3.1 Convergence Results	55
2.3.2 Kepler Dynamical Catalog	59

2.4	Discussion	67
2.4.1	Comparison to Other Catalog Level Analyses	68
2.4.2	Individual Systems of Note	69
2.4.3	Systems with TTVs	81
2.4.4	Hidden Planet Systems	83
2.5	Conclusion	86
2.6	Future Work	91
2.7	Acknowledgments	92
3	Enforcing Stability	93
3.1	Introduction	94
3.2	Stability Metrics	97
3.2.1	What is the Goal of a Stability Metric?	97
3.2.2	Hill Stability	100
3.2.3	SPOCK	100
3.2.4	Petrovich	102
3.2.5	Hadden and Lithwick	102
3.2.6	Additional Metrics	103
3.3	Methods	104
3.3.1	Run Initialization	104
3.3.2	Metric Calculation and System Simulation	108
3.3.3	N-body Integrations	109
3.4	Results	109
3.4.1	Metric Comparisons	109
3.4.2	Unconstrained Parameters	112
3.4.3	$e=i=0$	113
3.4.4	N-body results	119
3.4.5	$e=i=\text{small}$	120
3.4.6	$m=i=\text{small}$	122
3.4.7	Mutual inclinations	123
3.4.8	Jontof-Hutter posteriors	124
3.4.9	Constraint Comparisons	126
3.5	Discussion	127
3.6	Conclusion	130
3.7	Acknowledgments	132
4	Interior Modeling	133
4.1	Introduction	133
4.2	Methods	136
4.2.1	ExoMDN	136
4.2.2	Inputs	138
4.2.3	Computational Resources	140

4.3	Results	140
4.3.1	Catalog Compositions - Hot	141
4.3.2	Catalog Compositions - Cold	144
4.3.3	Catalog Compositions - Summary	144
4.3.4	Individual Planet Analysis	146
4.4	Discussion	156
4.5	Conclusion	162
4.6	Future Work	165
4.7	Acknowledgments	166

Bibliography	167
---------------------	------------

Chapter 1

Introduction

The primary purpose of this chapter is to provide a foundation whereby the reader will be able to understand the concepts that will be presented in the following chapters. As such, the topics will start from the most basic areas of discussion, for example, “what is an exoplanet” and build to more complex ideas such as how these planets are modeled. Additionally, a goal is that with the information presented here, it will be possible to contextualize this work and gain an appreciation for the benefit that it provides to the exoplanetary community.

1.1 Exoplanets: Who, What, and Why?

It seems to be part of human nature to wonder about whether or not we are alone in the universe. Before there was any observational evidence for other planets, stories of alien abductions or invasions captivated imaginations in the mid to late 20th century. Today, media franchises such as Star Trek and Star Wars are worth billions of dollars each because they continue to build on ideas such as: What would it be like if there were people on other worlds? What types of other worlds could exist? What would it mean for humanity if we were not alone? and many more thoughts along these lines. While we have known about most of the planets in our solar system for thousands

of years, evidence for extrasolar planets (exoplanets) did not come until the very end of the 20th century.

The first detected exoplanet was found by Wolszczan & Frail (1992) in 1992 around the rapidly rotating neutron star PSR1257 + 12. By using the 305 meter Arecibo radio-telescope to make precise timing measurements of pulses from this neutron star, it was concluded by Wolszczan & Frail that the pulsar is orbited by at least two planet-sized bodies. Several years later, in 1995, the first exoplanet around a sun-like star was found. This planet was discovered after acquiring 18 months of Doppler measurements (radial velocity) of the star 51 Pegasi by Mayor et al. (1995) and was found to be about half the size of Jupiter orbiting with a period of about four days.

These two discoveries opened the floodgates into the detection, exploration, and analysis of exoplanets, and since then, according to the NASA exoplanet archive, there are now over 5800 confirmed exoplanets. Immediately questions began to arise about whether rocky or gaseous planets are more common or whether the structure of our solar system is common or rare. To the point of whether our solar system structure is common or not, initial evidence originally pointed towards it being more unique. Figure 1.1, obtained from the NASA Exoplanet Archive, shows all known exoplanets plotted by their periods and radii with the eight labeled points being the eight planets in our solar system (NASA 2024). As can be seen, all of our planets fall outside of the clusters of known planets primarily because of their longer periods, this led to the possibility that our solar system did not fit the mold. However, as detection baselines have increased, methods have improved, and our overall capability to find exoplanets has been enhanced, this diagram has started to fill in, and the distance from the edge of the cluster of confirmed planets to the solar system planets has decreased. Therefore, there is emerging evidence suggesting that our solar system may not be an abnormal configuration for planets.

Addressing the question of how our solar system fits in the framework of exoplanetary systems as a whole is just one small example of what is currently being looked at within the exoplanet

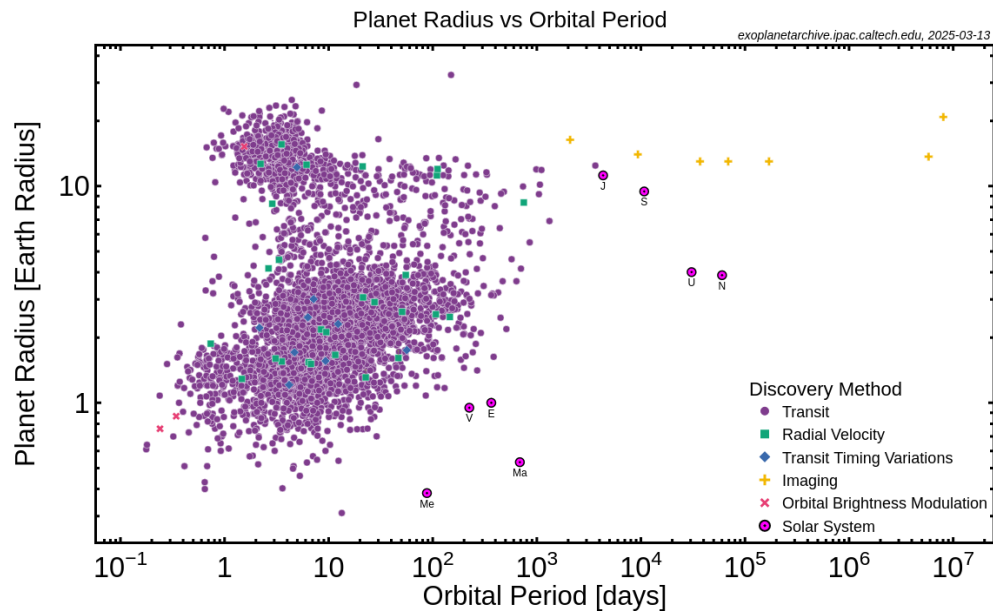


Figure 1.1 Period-radius diagram for all of the currently confirmed exoplanets obtained from the NASA Exoplanet Archive (NASA 2024). What is seen here is that most discovered planets are found with periods between 1-100 days and thus the labeled points indicating solar system planets fall outside of clusters of observed planets. However, as detection methods and capabilities have improved the distance between our solar system planets and these clusters has decreased. In the end, it is very possible that we will find that our solar system is not a rare configuration for planets to be found in.

community. What is most exciting and humbling is that we are currently living in a time where these questions can be addressed scientifically. For this work, the question of primary interest is: 'What are the physical and orbital parameters for all confirmed exoplanets' and along with this, 'how can these parameters be further constrained?' The importance of determining the physical and orbital parameters of exoplanets lies in the ability of these parameters to improve our understanding and ability to model planet formation and evolution. If through this work, we are able to construct a catalog of all the physical and orbital parameters for a large sample size of confirmed exoplanets, then any model that is meant to generate the same sample of exoplanets should be able to accurately reproduce our findings. In addition to this, by combining the mass and radius of a planet, it is possible to obtain the bulk density, which can then be used to constrain the planets composition. Understanding a planet's composition is one of the most important steps in answering the ultimate question of whether or not a planet is habitable.

1.2 Physical and Orbital Parameters

When addressing the question of calculating and constraining the physical and orbital parameters of exoplanetary systems, it is necessary to first define these parameters. For physical parameters, each planet in the system contributes two which are the planets mass and radius. Additionally, the star contributes five parameters which are its mass and radius, two limb darkening parameters, and the stellar dilution. Masses and radii are extremely self-explanatory, the limb darkening parameters c_1 and c_2 are used to describe how the stellar light profile changes when moving from the center of the star to the edge or limb. This change in brightness is due to the fact that when observing the central of a stellar disk you are looking into a deeper, hotter and therefore brighter layer of the star. The parameters c_1 and c_2 are used to fit a quadratic limb darkening profile of the form:

$$I(r) = 1 - c_1(1 - \mu) - c_2(1 - \mu)^2 \quad (1.1)$$

as given in Mandel & Agol (2002). Lastly, stellar dilution sets what fraction of light is believed to be the result of background objects. For a completely isolated stellar system, the dilution would be 0, in reality the dilution is not 0, but it is low enough that it can be set as such without inducing meaningful errors.

For orbital parameters, when modeling a system of two particles (or a star and planet) where the only force acting on the two results from a central potential such as gravity, the system originally begins with six spatial degrees of freedom which are often cast as being the position of the center of mass (\mathbf{R}) and the vector describing the difference in positions between the two particles ($\mathbf{r} = \mathbf{r}_2 - \mathbf{r}_1$). However, when studied in the Lagrangian formalization of the problem, it is found that the parameters for the center of mass are ignorable, as the center of mass is either fixed or moves with constant motion. This reduces the degrees of freedom down to three. An additional simplification comes when considering the spherical symmetry of the problem which requires the angular momentum to be conserved and thus there is only motion in a single plane. This in turn reduces the problem to simply considering a radial component, r , and an angular component, θ . From here it is possible and not entirely too difficult to work out the details for how r and θ change as a function of time as given in Goldstein et al..

While r and θ provide a usable framework for defining the orbit of a planet, they are not the most natural since it turns out that the motion will be an elliptical orbit around the star. By combining this knowledge, with the degrees of freedom for the problem, the constants of motion, and integration constants that are found by solving the second order differential equations, the problem can be solved in terms of orbital elements. Specifically, there are six orbital elements that are used to define the orbit and position of the planet which are: semi-major axis (a), eccentricity (e), inclination (i), longitude of ascending node (Ω), argument of periapsis (ω), and either the true anomaly (f),

eccentric anomaly (E), or mean anomaly (M). All of these are shown in Figure 1.2 (Wikimedia 2007). The first two parameters: a , and e are used to define the size and shape of the orbit. The next three: i , Ω , and ω define the orientation of the orbit. However, to truly orient the orbit, it is necessary to define a reference plane that the inclination and argument of periapsis are measured with respect to, as well as a reference direction to define the zero point for the longitude of the ascending node. When studying exoplanets, it is convenient to set the reference plane as the plane of the sky and the direction is the line connecting the exoplanetary system to the observer. With this definition, an edge on system has an inclination of 90° and a face on system has an inclination of 0° . The last parameter which is either the true anomaly (f), eccentric anomaly (E), or mean anomaly (M) gives the position of the planet within its orbit. Each of these three parameters gives this value, but with a slightly different definition. For example, the true anomaly gives the angle from the reference direction to the planet's location in its orbit with the angle's vertex at the star's position, but the eccentric anomaly gives the angle from the reference direction to the planet's location assuming it is traveling on a fictitious circle of radius equal to the semi major axis a that circumscribes the elliptical orbit. The mean anomaly does not have a physical representation, but is basically the position of the planet based off of its averaged motion throughout its entire orbit.

1.3 Transiting Planet Method

To date, the most successful method for detecting exoplanets has been the transiting planet method. A transit, is a generalized term for an eclipse and the method is to look for planets that eclipse, or transit, in front of their host star. In practice, this is not too difficult to accomplish. One observes a star and collects time series brightness measurements (photometry) for a period of days, months, or years. When plotted as a function of time, these measurements create what is known as a lightcurve.




Figure 1.2 Representation of the six orbital elements used to define a planet's orbit around its host star. The first two parameters: a , and e are used to define the size of the orbit. The next three: i , Ω , and ω are then able to define the orientation of the orbit once combined with a reference plane and direction. The last parameter which can either be the true anomaly (f), the mean anomaly (M), or the eccentric anomaly (E) is used to define the planets position in its orbit (Wikimedia 2007).

If when analyzing the lightcurve there are periodic drops in the star's brightness, this is a sign that a planet is transiting in front of the star.

A model representation of a transit obtained from Figure 1 of Winn (2010) is shown in Figure 1.3. There are several features in this figure that merit explanation. First, and perhaps the most obvious point, is that the reason for the drop in brightness is due to the planet covering up a portion of its host star. Before the transit occurs, the light received from the system is a combination of light from the star and the planet. As the transit begins, the planet covers a growing region of the star which in turn blocks the light from this region. The amount of light blocked reaches a maximum as the planet fully crosses in front of the star, and then reduces as the planet moves out from in front. At this point the light from the system reaches the original baseline. It is also worth noting that there are secondary eclipses (occultations) when the planet passes behind the star, given that this blocks the light from the planet, there is a drop in overall brightness, but this is much less than what is seen from the primary transit and thus they are essentially ignorable. A second feature of transits is their general U-shape. If stars were uniformly bright across their disk, a transit would cause a linear drop in light, reach a max drop once the planet was entirely in front of the star, continue at this level while fully in front, and then linearly increase in brightness again up to the normal level. However, as mentioned previously, because stars are brighter at their cores and dimmer at their limbs, the result in the transit lightcurve is a smoother U-shaped drop and increase in light. There are additional mechanics that can distort this clean transit signal such as a grazing transit. In this case, the planet does not completely pass in front of the star and this causes a V-shaped signal in the data rather than the typical U-shaped transit.

As for planetary parameters, the transiting planet method most readily gives the planet's orbital period, P , since this is simply the time between transit events. If the mass of the star is known, which is often the case thanks to stellar modeling, then the period can be used to calculate the semi-major axis of the planet's orbit, a . The depth of the transit $\delta \approx (R_p/R_*)^2$ which in turn means that when




Figure 1.3 Model representation of a transiting planet obtained from Figure 1 of Winn (2010) As the planet crosses in front of its star, it blocks a portion of light proportional to the planet's radius. The amount of blocked light reaches a maximum as the planet passes through the core of the stellar disk. Finally, as the planet moves out from in front of the star, the measured brightness increases back to the normal level. The U shape of the transit is largely a result of limb darkening which is the observed phenomenon that the edges of the star are not as bright as the core region. Additional variations in the planet's orbital parameters can also affect the shape of the transit in the lightcurve data.

the radius of the star is known, the radius of the planet can also be calculated as explained in Winn (2010). Determining the star's radius is not a trivial problem, as it requires precise luminosity, distance, and temperature measurements, but due to recent missions, such as Gaia which obtained exceptionally accurate stellar distances, stellar parameters tend to be fairly well constrained. The last parameter that is directly measurable from transit data is the mean anomaly, M . The equation for mean anomaly is:

$$M = \frac{2\pi}{P}(t - \tau) \quad (1.2)$$

where P is the period, t is the current time, and τ is the time of pericenter passage (Murray & Dermott 1999). However, the time of pericenter passage is just a reference time. Instead it is possible to use the transit times as the reference and given that the period has already been found, it is clear that every term is available to calculate M . Putting all of this together, using just the lightcurve data itself, and assuming stellar parameters are available, three of the eight physical and orbital parameters for exoplanets can quickly be determined. The remaining orbital parameters do have an effect on the lightcurves, but constraining them requires additional methods beyond just detecting the transit and measuring its depth.

Lastly, another important aspect of the transiting planet method is the probability of detecting a transit. Given that the transit depth depends on both the radius of the star and the planet, it should not be surprising that these play a part in determining the probability of a transit being visible. The second piece is the semi-major axis, a . A planet that orbits further from its star is going to have a lower probability of transiting. Lastly, the eccentricity, e , of the orbit will also play a factor as it affects the distance between the star and planet. Putting this all together, the probability of a transit is

$$p_{tra} = \left(\frac{R_* \pm R_p}{a}\right) \left(\frac{1}{1 - e^2}\right) \quad (1.3)$$

which is again given in Winn. This probability is especially useful in giving a sense of how many stars need to be searched in order to find a planet with a certain set of characteristics. Winn show that if the fraction η of stars that have this type of planet, then you would need to search $N \approx (\eta p_{tra})^{-1}$ to find such a planet that transits.

1.4 The Kepler Space Telescope

The most significant leap forward in the study of exoplanets and the use of the transiting planet method has been from the 0.95-meter Kepler Space Telescope (“Kepler”)(e.g., Catanzarite & Shao 2011; Fressin et al. 2011; Lissauer et al. 2014a; Lithwick & Wu 2012; Swift et al. 2013). Kepler is an exceptionally useful telescope for several reasons which start with it being a space telescope. By being outside of Earth’s atmosphere, it is not affected by the thermal variations, glare, and absorption that limit the abilities of ground-based telescopes. In part due to this, as well as its extremely high-quality construction, Kepler had a photometric precision of ~ 30 parts per million (ppm) in 6.5 hours for 12th magnitude stars (Lissauer et al. 2014b). This unprecedented precision was necessary in order to detect Earth-like worlds around Sun-like stars as the Earth would only block about 80 ppm of the Sun’s light once per year for at most 13 hours as viewed by a distant observer.

As explained by Lissauer et al. (2014b), Kepler’s primary mission was to conduct a statistical census of the abundance of planets as a function of planetary size, period, and stellar type. This statistical assessment was performed by examining nearly 150,000 stars for a period of four years while obtaining a photometric measurement every 30 minutes. By the end of its mission, Kepler identified thousands of exoplanetary candidates that have since been confirmed as true planets. With such a large sample size, it is then possible to look at large scale trends across Kepler’s exoplanetary database. For example, one parameter of interest is the occurrence rate of Earth-like planets, η_{\oplus} .

Zhu & Dong (2021) finds that $\sim 30\%$ of sun-like stars host planets with masses/radii down at the Earth-like level and each of these systems averages about three such planets, meaning that $\eta_{\oplus} \sim 1$. From this it is seen that planet formation processes are generally more effective than expected and Earth-like planets are quite common. Additionally, Kepler found that nearly half of the planet candidates are in system with multiple transiting planets and that configurations with tightly packed, low inclination orbits are quite common. Lastly, another surprising result was the Kepler Radius Gap, Ho & Van Eylen (2023) explains that there is an under-density of planets that have radii of $1.5-2R_{\oplus}$. Understanding this radius gap has been of particular interest as it is tied directly to planet formation and evolution.

While some statistical-level questions have already been answered or are currently close to an answer, one area in need of attention is constructing a homogeneously derived set of physical and orbital parameters for the Kepler planets. The reason why generating such a catalog is challenging is because many of the planetary parameters, as mentioned previously, are not readily available through transit data. To obtain the remaining planetary parameters, it often takes a considerable amount of time and can be computationally expensive. Many individual planets or systems have been studied on their own by various research groups, but the results cannot easily be combined with each other without introducing biases. This is unfortunate because it does not allow one to fully utilize the Kepler database. Due to the size of the Kepler catalog and the fact that it in and of itself was derived in a self-consistent manner, it is the perfect dataset to answer questions such as the true underlying distribution of orbital parameters or how do planetary compositions differ within a system and between systems. These questions will remain largely unanswered until an appropriate database of planetary parameters is available.

It is additionally worth noting that while the Kepler dataset is currently the largest and best catalog of exoplanets, it is not without biases itself. These are primarily introduced due to Kepler using transits to detect exoplanets. The transiting planet method is most sensitive to larger planets

that orbit closer to their host star as these are the planets that are going to have a high probability of transiting and have a clear signal in the lightcurve data. Thus, when looking at the period-radius plot in Figure 1.1 the cluster boundaries at low radii or long periods must not be taken as a physical cutoff in planet formation, but rather an indicator of our limits in detecting small planets or planets with longer periods. In regards to Kepler specifically, it is especially difficult to identify a planet with a year period or longer as there would be a maximum of four transits. Often times for smaller planets especially, this is simply not enough data to clearly identify the transits and prove the planet's membership in the system.

1.5 Systems of Multiple Transiting Planets

A welcomed finding, while searching the Kepler data for transiting exoplanets, was that many have been found to exist in systems with multiple transiting exoplanets and as such these systems are referred to as “multis”. Figure 1.4 shows the data for nearly 1,700 planets across 719 multis. Each point gives the period and radius for an individual planet and is color coded by the number of transiting planets that have been detected within the system. As can be seen, the plot is dominated by two and three planet systems, but there are several higher multiplicity systems including a single eight planet system. The Kepler multis are the main focus of much of this work as they have been shown by Ragozzine & Holman (2010) to be the most information rich out of all of the Kepler systems.

The greatest benefit obtained by looking at the Kepler multis is that there is additional dynamical information encoded within the lightcurves. For a system of just a star and planet, the planet's transit times are perfectly periodic because there are no other forces acting on the planet. However, once an additional planet is added to the system, these planets interact with each other through their mutual gravitational forces which in turn can cause transits to occur early or late. These




Figure 1.4 Period-radius diagram for nearly 1700 exoplanets discovered by the Kepler Space Telescope. These planets are all contained in systems where multiple planets have been observed to transit their host star. Each point is colored by the number of currently known planets in the system. As can be seen there is a significant number of two and three planet systems, but the range extends up to an eight planet system.

deviations are known as transit timing variations or TTVs. The strength of TTVs depends on both the planets' masses as well as their orbital configurations since a sequence of random pushes and pulls throughout an orbit tends to average out to no real effect. If instead the planets are configured such that the gravitational interactions add up coherently, then it is possible to get measurable TTVs. As such by measuring TTVs it is possible to obtain mass measurements for the planets as well as constrain their orbital parameters beyond what is possible for a single transiting system. Unfortunately, TTVs are only statistically measurable in $\sim 20\%$ of systems. For the remaining $\sim 80\%$ of systems, that do not have TTVs, it is expected that it not possible to constrain parameters such as mass or eccentricity at all.

Another observable dynamical effect in the Kepler multis are transit duration variations or TDVs. At a given orbital distance, the transit duration is controlled by the eccentricity and impact parameter. Planet-planet interactions can cause an orbit to precess and while this does not directly change the eccentricity, it does change the orbit location where transits are observed. Due to the

orbital eccentricity, the planet is moving at a variable speed throughout its orbit and therefore, by changing where in the orbit the transit is seen, the duration of the transit is also changed. However, the timescales for this to occur make this particular process unobservable in the Kepler data. The second method of changing durations comes from variations in the planet's inclination which in turn changes the impact parameter. In order for this to occur, orbits need to have significant mutual inclinations, which appear rare since there are only a handful of Kepler systems where measurable TDVs have been found. Using the observed TDVs (or lack thereof), it is possible to constrain orbit configurations which result in some constraints on theories of planet formation and evolution as explained in Ragozzine & Holman. The challenge is when modeling the system, it is necessary to find a full 3D solution as mutual inclinations are different from the planet inclinations relative to the plane of the sky. Given this additional difficulty and the fact that TDVs are only found in ~ 10 Kepler systems, TDVs have not been used much to determine planetary parameters.

When multiple transiting planets are observed, it is also significantly easier to prove that the signals in the data are the result of true exoplanets. Systems with only a single transiting planet were often confirmed via follow up of other detection methods such as radial velocities (RV). However, given that the Kepler database was focused on observing smaller planets around faint stars, it is not generally possible to examine and confirm these planets using RV and therefore confirmation must be accomplished by removing alternative explanations for the transit signal. With multiple transiting planet signals, this process is more straightforward for several reasons. First, in general, there are significantly fewer astronomical processes that could mimic the signal of a multiple transiting planet system compared to a single planet system. This implies that just the lightcurve alone is stronger evidence in the case of the Kepler multis than with the single transiting systems. Additionally, by comparing stellar parameters such as the limb darkening coefficients derived from the transits, it is possible to conclude that the planets are transiting the same star. Finally, since TTVs can provide

mass measurements, it can be shown that the signals are the result of planetary sized objects and not a stellar companion.

Beyond being able to constrain parameters better than systems with only one observed planet, the Kepler multis are also the most interesting systems. As already mentioned, there are dynamical interactions between the planets, but they also have orbital configurations that are not seen in our own solar system. For example, in Figure 1.5 is a diagram of the six planet system Kepler-11 compared to our inner solar system (Pyle 2011). What is fascinating to see is that the 5 out of the 6 planets' orbits are interior to Mercury's orbit in our solar system. This finding is not unique to Kepler-11, in fact, there are many observed planets, within the Kepler database, that orbit in extremely compact configurations. These systems then call into question ideas about planetary formation as well as the long term stability of such a setup. Another feature that has been investigated by groups such as Millholland & Winn (2021) is the "pea in a pod" configuration where it is often seen that the planets in a system tend to be regularly spaced and nearly equal in size. Lastly, many of the planets in the Kepler multis have been found to be just outside of mean motion resonances which are where the orbital period of one planet is an integer ratio of the other. Resonances generate many interesting dynamical interactions within a system such as enhancing TTVs or causing regions of reduced orbital stability. As such the resonant configurations for the Kepler systems is an area of active research by those such as Siegel & Fabrycky (2021). Specifically, the fact that these planets tend to be found just out of resonance also brings up questions about the evolution of the systems.

With all this in mind the argument that the Kepler multis are the most information rich systems within the Kepler dataset is easily justified. By studying them and providing constraints on the physical and orbital parameters of the planets they contain, it is possible to provide a significant enhancement to our overall understanding of exoplanets.




Figure 1.5 Comparison between the Kepler-11 system and the inner planetary orbits of our solar system. Within the Kepler database there are multiple examples of tightly space planetary configurations, such as the one for Kepler-11, where multiple planets have orbits interior to the innermost planets in our solar system (Pyle 2011) In the case of Kepler-11, there are five planets that orbit their star closer than Mercury orbits the Sun. These systems are noteworthy both for their significant contrast to our solar systems and their implications for planet formation and evolution.

1.6 Parameter Degeneracies

As mentioned, there are planetary parameters that are readily obtainable when using the transiting planet method to detect and characterize exoplanets. On the other side of the equation, there are parameters that are particularly difficult to constrain because they are degenerate with each other. This means that instead of finding a unique value for each parameter, what is found is a relationship between the parameters. When analyzing exoplanets that have strong TTVs, one specific degeneracy is the mass-eccentricity degeneracy.

This degeneracy is explored thoroughly in Lithwick et al. (2012), but in summary, an exoplanet's eccentricity can be considered as containing a free and a forced part. The free eccentricity is essentially an intrinsic level of energy within the planet's orbit and overtime can be damped out due to interactions with a planetary disk or through tidal damping. The forced eccentricity is a result of being near a mean motion resonance, which is a configuration that many Kepler systems find themselves in. As long as the planet remains close to the resonance, the forced eccentricity remains, but eventually dissipative interactions can push planets apart and thus further away from the resonance. For the mass-eccentricity, it is the free eccentricity that causes the main challenge. The equation for the amplitude of TTVs after dropping some order-unity coefficients is given by Lithwick et al. and is equal to

$$|V| \sim P \frac{\mu'}{|\Delta|} \left(1 + \frac{|Z_{free}|}{|\Delta|}\right) \quad (1.4)$$

$$|V'| \sim P' \frac{\mu}{|\Delta|} \left(1 + \frac{|Z_{free}|}{|\Delta|}\right) \quad (1.5)$$

Where P is the planet period, μ is the planet to star mass, Δ is the distance from resonance, and Z_{free} is the free eccentricity where the primed parameters are for the outer planet. From this equation, it can be seen that, first, in order to use TTVs to constrain masses, the TTVs need to be seen for

both planets as the equations are coupled. Second, in order to use the observed TTV amplitude to invert the equation and calculate the masses, the free eccentricity term must be known. In this second point lies the degeneracy. Without knowing the free eccentricity, all you are left with is a relationship between mass and free eccentricity.

The mass-eccentricity degeneracy is the primary obstacle when using TTVs to constrain planetary masses. Outside of this particular degeneracy, there are additional challenges that can present themselves when considering more extreme planet configurations. For example, in the case of grazing transits, it becomes possible to allow for significantly inclined orbits if the planet's radius is also increased to account for this and keep the planet transiting. While this effect can largely be constrained by examining the shape of the transit, it again illustrates the difficulty in fitting planetary parameters and the need for robust fitting routines.

1.7 Photodynamical Modeling

From the discussion of planetary parameters in Section 1.2 and the inherent challenges in obtaining unique values for these parameters as explained in Section 1.6, it is clear that a more advanced method must be employed in order to constrain or determine the physical and orbital parameters of the Kepler planets. The best tool for performing such a task is a photodynamical model Ragozzine et al. (in prep.). While many of the particulars of photodynamical modeling will be presented in Chapter 2, especially as it relates to the modeling code PhoDyMM which was used to perform the research presented in that chapter, it is still valuable to give an introduction to photodynamical modeling here.

A photodynamical model consists of two key elements which are a photometric model and an n-body integrator. When given a set of stellar and planetary parameters, the n-body integrator is able to evolve the orbits of the planets by calculating the gravitational forces on each body and

stepping forward in time. The output of such a routine is the position of the planets as a function of time. This can then be passed to a photometric model which has the capability of determining photometric dimming based off of planetary positions. In the end, the result is a model lightcurve that can be compared directly to the Kepler data. Input parameters to the photodynamical model can be adjusted to find the best fit to the Kepler data for a given system of planets.

While there are other methods for determining planetary parameters, such as direct fitting to the TTVs, they are suboptimal for several reasons which include:

1. When using TTVs to infer masses, generally the transit times for each planet are used alone while the lightcurve as a whole is ignored, which precludes a deeper and more self-consistent analysis that would be possible when using the entire lightcurve.
2. Many planets are too small to determine individual transit times and as a result, reliable TTVs cannot be measured.
3. Most TTV analyses only use the 30 min cadence data from Kepler. As mentioned previously, Kepler gathered photometric measurements every 30 minutes (long-cadence data) for all the stars in observational field. However, for some of the time, Kepler obtained photometric data every 1 minute (short cadence data) and these data are available for most of the Kepler multis. Short-cadence data have been shown to \sim double the precision of TTV measurements.
4. Existing analyses tend to focus only on a handful of particularly interesting systems. This makes comparisons difficult, as differing methods must be taken into account. It also neglects many systems that could have interesting dynamical interactions within the system that have not been found yet.
5. It is possible that the solutions from existing modeling techniques are not dynamically stable and are therefore unrealistic.

When using a photodynamical model, these shortcomings are almost entirely avoidable. Since the model is capable of generating synthetic lightcurves, they can be directly compared to the Kepler data and there is no need for the extraction of 'meta-data' such as transit times. This also alleviates the challenge of planets being too small to model. A small signal-to-noise planet will still be difficult to work with in the sense that its presence is difficult to see in the data for comparison, but from a strictly modeling standpoint there is no size limit on the planets that can be simulated. Photodynamical models are flexible with the data they can take in and can return the model lightcurve at arbitrary points in time. This means they can easily make use of both the long- and short-cadence data from Kepler which in turn allows them to benefit from the increased precision from the short-cadence data. Given the overall flexibility of using a photodynamical model, it is possible to analyze all of the Kepler multis and not just the handful that have been identified as being particularly interesting. Lastly, it is possible to limit the range of parameters that are given to the photodynamical model in order to enforce that only stable configurations of planets are analyzed.

As mentioned previously, in order to fully leverage the Kepler database, a self-consistent analysis of the planets needs to be performed. Additionally, remembering that it is the systems with multiple transiting planets are the most information rich and best positioned for answering questions about the distributions of planetary parameters. Therefore, the clear conclusion on how to obtain a self-consistent catalog for the most important subset of Kepler systems, is to use a photodynamical model.

1.8 Long-term Stability

While a photodynamical model is able to extract the most information out of a lightcurve, there is still no guarantee that parameters such as mass, inclination, or eccentricity can be constrained by

the data alone. One potential way to constrain these parameters further is by considering whether a particular set of planetary parameters results in a configuration that is stable on the Myr-Gyr timescale. The reason that such an analysis can be performed is because the Kepler systems are understood to have existed for millions to billions of years and it is statistically unlikely to observe them before a cataclysmic event. Therefore, it must be assumed that the current configurations that we find these systems in must be stable on timescale similar to their age.

The difficulty in using stability as a method for constraining planetary parameters is in determining whether or not a configuration is unstable. The most direct method is to use n-body integration which can accurately model the positions of planets on the required timescales. Unfortunately, this is computationally expensive. In order to provide meaningful constraints on parameters, it is likely that thousands of potential configurations need to be considered, and thus n-body integration proves to no longer be feasible. In order to remedy the situation, many analytical methods have been developed to provide an estimate on whether or not a planetary configuration is stable.

One of the simplest methods for classifying stability is using the distance between planets in units of mutual Hill radii. The equation for the mutual Hill radius for two planets, as given in MacDonald et al. (2016), is:

$$R_{mH} \equiv \left(\frac{\mu_i + \mu_{i+1}}{3} \right)^{1/3} \left(\frac{a_i + a_{i+1}}{2} \right) \quad (1.6)$$

Where μ is the planet-to-star ratio. The separation between planets in units of Hill radii, Δ (not to be confused with the distance from resonance Δ given in Equation 1.4), is then:

$$\Delta = \frac{a_{i+1} - a_i}{R_{mH}} \quad (1.7)$$

It was shown by Gladman (1993) that if two planets on circular orbits are separated by a distance $\Delta > 2\sqrt{3} \approx 3.46$, then the planets will never develop crossing orbits. The fact that orbits will never cross is then strong evidence for the system being stable as a whole. This argument becomes more




Figure 1.6 Orbital separation between planets in units of mutual Hill radii versus stability time obtained from Figure 2 of Lissauer & Gavino (2021). Ignoring the smaller oscillations initially, it is seen that there is a clear linear trend between orbital separation and the stability time. Given that this figure has a logarithmic y-axis, this trend is actually an exponential relationship. Thus, as orbital separation increases, the planetary system becomes exponentially more stable. The red lines in the figure represent the locations of mean-motion resonances which are the primary cause for the oscillations within the relationship between spacing and stability.

difficult as it is extended to systems of more planets with eccentric orbits. For example, Smith & Lissauer (2009) finds that in the case of multi-planet systems, a distance of 10 mutual Hill radii is needed between each planet. In general, what has been found is that as planet separation increases, so does stability. This is most clearly shown in Figure 1.6 which is taken from Figure 2 of Lissauer & Gavino (2021). The x-axis is the distance between planets in units of mutual Hill radii, and the y-axis is the calculated stability times based off of numerical simulations. Ignoring some of the small oscillations which are the result of resonances, there is a very well-defined linear trend between stability times and planetary separations which in reality ends up being an exponential trend after factoring in the logarithmic scale on the y-axis.

A second class of stability classification is based off of the overlap of mean-motion resonances (MMRs). As explained earlier, a mean-motion resonance is where there is an integer ratio between the orbital periods of two planets. Beginning with work done by Wisdom (1980) and extending to more recent research by those such as Hadden & Lithwick (2018) and Tamayo et al. (2021b), it has been found that once MMRs overlap, planetary orbits begin to be chaotic and within small time scales, relative to the lifetime of the system, instabilities can occur. Thus, determining stability turns into a problem of defining the regions of parameter space where MMRs begin to overlap and where chaos is introduced. Once an analytic equation for this region is defined, it can quickly be applied to determine whether a set of planetary parameters is expected to lead to chaotic orbits and therefore constrain the parameters.

Beyond these two classes of stability metrics, there are many other analytical expressions aimed at defining the boundary between stable and unstable orbital configurations. The challenge remains to address how these metrics compare with each other as they are derived in various ways and define stability differently. Additionally, it is necessary to assess to what level a stability analysis can constrain parameters such as planetary mass, eccentricity, or inclination. The investigation of these problems is the main focus of Chapter 3.

1.9 Formation and Evolution

Planet formation and evolution are areas of active research and debate within the field of exoplanetary science, as evidenced by the thousands of papers on various aspects of the problem. The challenge in solving these problems stem from the fact that it involves modeling particles, fluids, magnetic fields, collisions, thermodynamics, etc. on a scale from microns to Jupiter sized planets. It is further complicated by the fact that the full process occurs over millions of years and our only observational evidence is snapshots of it along the way by observing young exoplanetary systems. Given the

overall complexity in presenting a complete exploration into the current understanding of planet formation, as well as the limited role that such information plays in the rest of this work, only a brief overview will be given here.

Formation begins in the protoplanetary disk which consists of gas and dust that was not incorporated into the star. Through various potential means, dust grains can grow into planetesimals which range in size from 100 m to a few hundred km (Armitage 2024). Generating these planetesimals was one of the primary challenges in formation theories. It was fairly straightforward to condense particles together and build up “rocks” approximately a meter in size, but beyond this point, collisions would lead to as much breakdown as buildup. The current working theory is that planetesimals form via the streaming instability. Youdin & Goodman (2005) explains that the streaming instability generates particle clumps that are dense enough to collapse gravitationally into planetesimals. After forming planetesimals, the next stage is growth into either terrestrial planets or the cores of ice or gas giants. Here, there is another fork in the road as to whether the growth occurs via planetesimal accretion or pebble accretion. In the case of planetesimal accretion, planetesimals collide with each other and if these collisions occur at a low enough velocity, the result is growth. Pebble accretion requires that the protoplanetary gas still exists and that it is able to bring small particles “pebbles” within the planetesimal’s Hill sphere which allows it to capture and accrete the pebbles (Armitage 2024). While older models favored planetesimal accretion, recent work is suggesting that these two methods work in tandem where pebble accretion follows a period of planetesimal accretion. However, Armitage (2024) also explains that it is possible for pebble accretion to be the only growth mechanism if planetesimals are large enough. Lastly, once the core of a planet is massive enough, it can start accreting significant amounts of gas. If a planet can accrete an atmosphere or envelope on the order of the mass of the core ($M_{env} \sim M_{core}$) before the gas disk dissipates, then it will enter a runaway growth phase which can lead to the formation of a gas giant. On the other hand, if this

point is not reached before the disk dissipates then the result is more likely to be akin to an ice giant or sub-Neptune.

Planetary evolution is also a complex problem as it depends on determining multiple sources of torques that lead to planets either migrating in or out. Armitage states that the most well studied source involves resonant interactions between the disk and planet which allow for the exchange of angular momentum with the disk and migration of the planetary orbit. Within the realm of resonant interactions there are also Linblad resonances and co-orbital resonance. The first causes migration that is almost always inward, while the second can result in outward migration. Outside of resonant torques new research has found that thermal and pebble torques exist and can affect migration. Thermal torques are the result of variations in the disk due to a luminous planet. Pebble torques are a consequence of scattering the particles that are coupled to the disk's gas when they drift past a planet (Armitage 2024).

While the work presented in the subsequent chapters did not rely heavily on the concepts of planet formation and evolution, what it did was provide estimates and constraints on orbital parameters, masses, and planetary compositions. This information can be used as a check on planet formation and evolutionary models as their final results must align with the systems that we are observing. Perhaps most directly applicable is the effort presented in Chapter 4 which looks at using planetary masses, radii, and temperatures to construct theoretical compositional models. If for example, there were a large population of interior planets that had a statistically significant ice layer, this would be clear evidence of planet migration over *in situ* formation and it would be the work of migration modeling to determine what processes led to the planets being in their current locations.

1.10 High Performance Computing and Book Keeping

Where the rest of this chapter has been about introducing concepts needed to fully understand and contextualize the research that will be presented in the following chapters, the point of this section is to provide advice and guidance to anyone working on a similar project.

The first area of discussion is in high performance computing (HPC), but is also tied to effectively making use of available resources which includes personal research time. BYU's supercomputer is available to any student that has a clear and defined research need for its high computing power and that can obtain sponsorship from a faculty member. To those that meet these criteria, they are given access to nearly 35,000 CPUs and over 170 TBs of memory. While these resources are shared among all that need them, throughout my work, I found that I could readily obtain 1,000-2,000 CPUs and with a little bit of time this could easily extend up to 10,000 CPUs. With this in mind, whenever I was performing a particular analysis, there was always the question in mind of whether the task was best done in series or if there was a potential to parallelize it and make use of the supercomputer to speed up the computation time. The answer to this question usually was based off of the independence of the analysis and the time needed to perform each task.

For example, when modeling the 719 Kepler multis, each system was completely independent of each other and furthermore each walker used to run the DEMCMC was also essentially independent. Given this, the photodynamical analysis of the multis definitely met the criteria for moving the work from a traditional computer to the supercomputer. This need was further enhanced when considering the time required to analyze a single system could be in the range of days to weeks. On the other hand, when plotting combined results from this work, the question to use the supercomputer became less obvious. Needing combined values removed a significant amount of independence. While it would be possible to have a file that is accessed independently for each system, the process to do so became more difficult. Thus it was necessary to address how much time would be needed in order to obtain the plotted values. If all that was necessary was to read in a value from the 719 files and

it took on average a few seconds, this could easily be done in a standard computing environment. However, if there were intermediate steps that would require on the order of minutes per system, then a method for performing this work on the supercomputer was considered and developed.

Once the decision to run a particular analysis on the supercomputer was made, code needed to be tested to ensure that it was reliable in the supercomputing environment and computing resources needs had to be assessed. From my considerable work, I have developed these lessons/reminders:

1. Test any code on the supercomputer thoroughly. Ensure that it runs on a test case or sample of test cases before jumping to a large dataset/sample size. The issue here is that as the supercomputer is a shared resource, it employs a scheduler to determine when particular jobs are going to be started. The more jobs or resources that an individual user utilizes, the lower their priority is in the scheduler which in turn means that they will wait longer for their supercomputing jobs to begin. If a large number of jobs are submitted with faulty code, then priority is wasted.
2. Run codes in a virtual environment to guarantee that the right versions of software are installed. This is especially important on the supercomputer as updates can occur without warning and change how individual software packages are linked to each other.
3. Have a clear understanding of what computational and memory resources are needed. When submitting a job to the supercomputer, it is necessary to define how many CPU are needed as well as how much time and memory will be used. If a piece of code attempts to go beyond any of these predefined boundaries, then the job will be killed and progress will halt.
4. Regularly monitor the progress of analyses being performed on the supercomputer. Specifically if jobs are finishing significantly faster than would be expected, it is possible that this is an indicator of a bug in the code or jobs being killed.

5. Consider moving processes around if it will increase the likelihood of progress being made. For example, if a particularly resource heavy process can be moved to the end of the analysis (such as parsing output for future plots), then it may be possible to perform the computationally light work on one job which would likely be started earlier due to needing fewer resources. Then submit a second job with the increased resource requirements, but potentially a smaller computational time needed.
6. Utilize the experience of others whether it is advisors, other students, or the supercomputing staff to aid in solving particularly tricky problems.

The second area of discussion is in book keeping. It was said throughout this work that the primary problem that I would face is in book keeping. This was largely true as many of the underlying codes and routines already existed and my task was to make them run on a large dataset. What I have learned from this is that to the best of one's ability, it is necessary to try and predict what values or information you might want at a future time. If there is a particular value that you might want to calculate, make sure that all the necessary variables are readily accessible. In other cases, make sure that all of the values used to produce a particular output are saved should the need to reproduce results arise. What both of these cases essentially amount up to is having well-defined and formatted log files. In these files, try to consider all the relevant information that will be needed down the road and include it in the file. If done well, this can save significant time as values are not always needing to be recalculated or looked up elsewhere.

A particular lesson that I only fully learned towards the end of this result is in regards to the planetary names. An unfortunate feature of working with the Kepler dataset is that every planet has about three names. There is the KOI number, the KIC, and potentially a Kepler number and letter. While each number is unique in its respective system, a significant challenge occurs in the fact that KOI-36 is not the same as Kepler-36. Throughout computation and analysis, I found that it was best and easiest to maintain a consistent naming routine by only utilizing the KOI numbers, but if logs

had also included the other naming values up front, later on it would have been easier when results needed to be correlated across different catalogs and inputs. Our particular analysis tool had the additional quirk of listing all of the planets by period, which is not necessarily the order of KOI numbers within a system (our KOI-41.1 is not necessarily KOI-41.01 in the catalogs). However, it was not until the final project of my work that I began including the true KOI number in the file names and not just the PhoDyMM number. By having the KOI number more readily accessible, getting the results ready for publication was significantly easier as these are the primary numbers used to reference each planet across the exoplanetary community. Ultimately, the better you can predict you needs down the road, the better chances you have of saving yourself some frustration when compiling results.

1.11 Dissertation Outline

The remaining chapters will be given as follows. Chapter 2 will cover my work in applying a photodynamical model to all of the Kepler systems with multiple transiting planets. While this work was performed in the middle of my graduate research, the other projects are meant to be extensions of this primary analysis. Following in Chapter 3, I explore a method that can be used to further constrain planetary parameters on the basis of requiring long term stability. Naturally, the goal of this analysis is to be able to eventually constrain the parameters found by applying the photodynamical model to the Kepler multis. Finally, in Chapter 4, I demonstrate one potential use of the planetary parameters obtained in Chapter 2 which is to use them as inputs to generate interior models. These interior models can then be used to inform theories about planet formation and evolution as well as increase our understanding of the Kepler database as a whole.

Chapter 2

Photodynamically Modeling the Kepler Multis

We apply our photodynamical modeling tool PhoDyMM to the 719 Kepler systems of multiple transiting planets (multis). After over a million CPU days, we obtained converged posterior distributions for all the physical and orbital parameters for all the planets in 661 systems, where the remaining 58 systems required either additional computation or error corrections. We obtained mass measurements for 104 planets with an uncertainty of $<25\%$ and density measurements for 139 planets where the averaged 1σ error is $< 0.5g/cm^3$. Using these results, we were also able to construct the most complete mass-radius diagram for the Kepler multis and present it using draws from the posterior distributions in order to provide the most accurate visualization of the uncertainties in the mass and radius values.

2.1 Introduction

In its four-year prime mission, the Kepler Space Telescope (“Kepler”) obtained photometric measurements for nearly 200,000 stars. By looking for evidence of transiting planets, these data have

been used to discover thousands of exoplanets, which has led to a significant leap forward in our understanding of exoplanetary formation, evolution, dynamics, and architectures, as evidenced by the hundreds of papers published in the last several years.

Of particular interest are systems with multiple transiting exoplanets. These multi-transiting systems or “multis” are special because of the additional information which can be gained by analyzing the planet-planet interactions and the additional context they provide on exoplanetary architectures (e.g., Lissauer et al. 2011). Furthermore, false positive scenarios are much less likely, allowing candidates in multis to be more easily validated (Lissauer et al. 2013; Rowe et al. 2014).

With stellar parameters (now precisely measured for nearly all Kepler targets by Gaia), transit observations strongly constrain Periods (P), the time of the transit near a particular epoch (T_0), the planetary radius (R_p), and the on-the-sky inclination (i). When the stellar density is well known and/or when comparing multiple planets transiting the same star, then the transit duration also loosely constrains the eccentricity (e) and argument of peripase (ω) through the so-called “photoeccentric effect” (Ragozzine & Holman 2010).

Additional constraints can be placed through planet-planet dynamical interactions, typically in the form of Transit Timing Variations (TTVs) and sometimes Transit Duration Variations (TDVs). TTVs are seen in $\sim 20\%$ of planets (e.g., Kane et al. 2012) and place additional constraints on e and ω (e.g., Lithwick et al. 2012) and the mass (m_p) of the planet (possibly unknown) that perturbs the planet experiencing TTVs (Agol et al. 2005; Holman & Murray 2005). TDVs are seen in $\sim 1\%$ of planets and can place additional constraints on inclinations and differences in the longitude of the ascending node (Ω). Thus, planet-planet interactions allow for much greater insight into planetary physical (m, r) and orbital ($P, T_0, e, \omega, i, \Omega$) parameters. Most studies so far focus on using measured TTVs and TDVs to infer orbital and physical parameters of planets, but it is also possible to place constraints on these parameters from the lack of detected planet-planet interactions.

The ability to constrain parameters can in turn be taken further by using the masses and radii in order to determine densities which inform investigations about composition, formation, evolution, and habitability. Given this, it is clear why systems with multiple transiting planets have been said to be the most information rich out of all Kepler systems (Ragozzine & Holman 2010). Furthermore, detecting planet-planet gravitational interactions can help to confirm system membership, as it is impossible for these to arise from a merely blended source. While there are more Kepler single-transiting systems than multis and hundreds that show evidence for TTVs, they are statistically more likely to be false positives and it is not possible to measure densities or other dynamical parameters.

Many of the Kepler systems have been analyzed using a variety of tools and procedures, but there is currently no complete catalog of all of the physical and orbital parameters for all of the systems with multiple transiting planets. This is unfortunate because the size and homogeneity of the Kepler dataset puts it in the perfect position to look at inter-system trends and behaviors such as the true underlying distributions for the masses, radii, and orbital parameters of exoplanets. To fully leverage the power of the Kepler data, any analysis to create a complete catalog of physical and orbital parameters would need to be self-consistent and homogeneously derived in order to not be plagued with the need to account for biases between calculation methods. This is the purpose of this work: to create a complete, self-consistent catalog of the physical and orbital parameters of all of the Kepler systems with multiple transiting planets, which will allow for exoplanet population-level investigations. It will also broaden the scope of any specific analysis of these systems since it will be possible to perform such an analysis on all of the systems with multiple transiting planets. Furthermore, we will make significant progress on individual systems that either have not been studied in detail or were studied early in the Kepler mission and have not had an updated analysis. Our analysis may not be optimal for detailed study of an individual system with the very best data available, but is still a contribution to dozens of previous studies of Kepler multis.

It has repeatedly been shown that the best tool for analyzing planet-planet interactions in Kepler systems is a photodynamical model (e.g., Ragozzine & Holman 2010) since it is a near-optimal inference of physical and orbital parameters directly from observational data. A photodynamical model is a model that combines both an n-body integrator, which is thus able to determine a planet’s location at any time, and a model that is able to calculate stellar dimming due to planetary transits based off of the positional information. For this work, we make use of our photodynamical model PhoDyMM, the PhotoDynamical Multi-planet Model, which also includes Bayesian parameter inference methods (discussed below).

Photodynamical modeling in general and PhoDyMM in particular has multiple benefits over other methods for parameter inference. First, PhoDyMM fits directly to the Kepler lightcurves, unlike TTV methods which gather transit times and then fit a dynamical model to the transit time metadata. This allows for a more detailed and self-consistent analysis of the system (e.g., including information on transit durations). Second, PhoDyMM can make use of Kepler’s long and short cadence data. Kepler’s primary cadence was to collect a new measurement every 30 mins and TTVs have only been produced systematically with this long cadence data. However, Kepler also made short cadence observations of selected targets every 1 minute and many of the Kepler systems with multiple transiting planets (multis) have short cadence data which are known to \sim double the precision of TTVs for small planets. Third, individual transits do not need to meet a minimum signal threshold to be modeled, allowing for the study of smaller or lower SNR transits. This avoids many issues that arise when using TTVs such as providing meaningful uncertainties (e.g., Judkovsky et al. 2024; Leleu et al. 2021; MacDonald et al. 2016). Photodynamical models are not limited to certain systems, especially when paired with Bayesian parameter inference as in PhoDyMM. Fourth, the main downside perceived for photodynamical models is computational speed and ease of use, but PhoDyMM is not much slower (per likelihood evaluation) than regular lightcurve calculations. Photodynamical models in general continue to improve in speed and ability (e.g., Langford &

Agol 2024). These strengths have been demonstrated by previous analyses that used PhoDyMM to characterize complex planet-planet interactions between small planets (e.g. MacDonald et al. 2021; Mills & Fabrycky 2017; Mills et al. 2016).

A self-consistent homogeneous photodynamical analysis of all Kepler multis thus represents a near-optimal extraction of information from the best exoplanet dataset for measuring all physical and orbital parameters for exoplanets with a detailed understanding of detection biases. Such a catalog can be used to improve the study of individual systems, measure new masses and densities to understand interiors and compositions, better characterize architectures, and allow for demographic studies. That is, such a catalog would generally provide extensive insight into the formation, evolution, and habitability of planetary systems. We are thus pleased in this work to present the Kepler Multis Dynamical Catalog (KMDC), a complete catalog for nearly all of the Kepler multis with posterior distributions for all of the physical and orbital parameters for each of the planets.

The remainder of the paper will be organized as follows. Section 3.3 will present our methods for this work such as the basics of initializing and running PhoDyMM, our lightcurve preparation, and our criterion for stopping PhoDyMM. Following this, we present the KMDC in Section 3.4 along with our mass-radius distribution and calculations for the number of well-constrained planetary masses. In Section 3.5, we highlight our performance in modeling several well-known systems of particular interest and importance in the Kepler database. In addition, we will discuss our performance compared to other works of similar scale. Finally, Section 3.6 will give our conclusions and final thoughts on this project.

2.2 Methods

This section begins with an overview of our photodynamical modeling software PhoDyMM with a particular focus on explaining the basics of the Differential Evolution Markov Chain Monte Carlo

algorithm that it implements in order to generate Bayesian posterior distributions for the modeled parameters. We then move into an explanation of our data selection and preprocessing routines that left us with the sample of 1858 planets across 719 systems. Finally, this section ends with the description of the methods that allowed this work to proceed nearly autonomously. Of particular importance to this automation process are the metrics that we employed to determine whether our solution was accurate and the posterior distributions has converged.

2.2.1 PhoDyMM

A summary of PhoDyMM’s tools and procedure will be given here, but for a more complete introduction to PhoDyMM, readers are directed to Ragozzine et al. and the GitHub repository (<https://github.com/dragozzine/PhoDyMM>). PhoDyMM implements both an n-body integrator and a Differential Evolution Markov Chain Monte Carlo (DEMCMC) algorithm in order to perform Bayesian inference of the posteriors of the physical and orbital parameters of a system with an arbitrary number of planets. Beginning with a set of stellar and planetary parameters, PhoDyMM uses its n-body integrator to integrate a system’s planets forwards and backwards in time to find potential transits. Near transits, the integration is refined to find the time, impact parameter, and velocity at the center of the transit; the planetary positions are then assumed to vary linearly with these parameters (after considering light-travel time). Using these positions and the lightcurve generation method of Pál (2012), PhoDyMM then produces synthetic lightcurves that can be compared directly to the Kepler lightcurves. We use uncertainties from the detrended Kepler lightcurves (discussed below) and (reasonably) assume all measurements are independent. With this noise model, our Bayesian log-likelihood \mathcal{L} is equivalent to $-\frac{1}{2}\chi^2$.

An initial guess for all parameters for all N planets and the star is required to begin the analysis. For parameters less related to dynamical interactions (Period P , time of transit “0” T_0 in Barycentric Kepler Julian Date - 67 = BJD - 2454900, inclination to the plane of the sky i , and planet to

star radius ratio $\frac{R_p}{R_*}$), initial values are taken from the best-fit parameters in Lissauer et al. (2024). Eccentricities (e) and arguments of periaapse (ω) are modeled as $\sqrt{e} \cos \omega$ and $\sqrt{e} \sin \omega$ which are both initialized to zero. The longitude of the ascending node (Ω) for all planets is set equal to zero and held fixed with tests showing this did not affect inference of other parameters even in the presence of Transit Duration Variations. Initial guesses for planet masses are taken from the most recent (as of 2023) mass estimate from NASA’s Exoplanet Archive Planetary Systems Table if available; if not, an initial guess of $M \approx R^{2.06}$ is used (as in Lissauer et al. 2011). The stellar mass (M_*) and radius (R_*) were also taken from the uniform GAIA parameters in Lissauer et al. (2024) from Berger et al. (2020). The initial guess for the limb darkening parameters (c_1, c_2) was also based on the stellar parameters calculated by LDTK (Parviainen & Aigrain 2015) discussed in more detail below. PhoDyMM’s DEMCMC uses an ensemble of multiple chains that each require an initial guess; this is provided by drawing from a very tight multi-variate Gaussian ball around the initial guesses (set by the PhoDyMM .in file) as is generally recommended.

DEMCMC samples from the Bayesian posterior distribution by sequential evolution of an ensemble of interdependent chains we call walkers that represent a position in parameter space. The DEMCMC sampling algorithm has walkers “take a step” (or proceed one link in the Markov Chain) by adding to its parameters, the difference between two other randomly selected walkers which is then multiplied by a scaling factor γ . The log posterior probability (\mathcal{L} modified by the priors discussed below) is calculated for the new walker positions and, as usual, this step is always accepted if the parameters are better and rejected with probability proportional to the difference in posterior probabilities. Rejected steps result in repeating the values of the parameters in the Markov Chain, as required for the DEMCMC to produce samples from the posterior distribution. This procedure continues until our stopping criteria discussed below which are designed to ensure that the best fit is reasonable and that we have drawn 1000 independent samples from the “burned-in” chains.

As mentioned, each proposed step is accepted or rejected according to the $\Delta\mathcal{L}$. Ter Braak has shown that the optimal percentage of accepted steps, f_{opt} , is 0.23 which, if maintained, can speed up overall performance. Therefore, if after a generation of walker steps, the acceptance fraction is greater than f_{opt} , then the step size scaling factor γ is increased (to better explore parameter space); otherwise gamma is decreased (to better sample locally good values). Unfortunately, it was often found that this process often led to an over-tuning of gamma which ultimately led to gamma, and thus the walker step size, going to 0. To remedy this, we set a range of ± 0.05 on f_{opt} where gamma would not be adjusted and also set a minimum gamma value (usually 0.003). This solution had the potential to allow the acceptance fraction to move well below optimal, but was seen as acceptable fix to ensure that walkers would continue exploration and not ‘freeze out’.

We sought to construct posteriors with an effective sample size (ESS) of 1000 so that the single-parameter marginal distributions would be well-sampled out to the 3σ tails. Given that each walker step is highly correlated with its previous location, new steps do not give an independent sample from the posterior. Instead, walkers need to travel a number of steps, called the autocorrelation length (ACL), before a new effectively independent sample is added to the posterior. To make the goal of an ESS of 1000 more computationally tractable, we initialized each system with 200 walkers, thereby requiring each walker to only need to travel 5 ACLs. Many walkers also help DEMCMC map the parameter space and thus choose better steps. The ability to have a large number of walkers became vital to this work because, while ACLs can be as low as a few hundred steps, our calculations showed ACLs $\sim 10^5$ presumably due to the high dimensionality of the parameter space ($6N + 4$ for multiplicity N).

2.2.2 Priors

PhoDyMM is built to have uniform priors by default which is preferred for this work because it does not bias the posteriors towards a preconceived distribution. If the true underlying distributions

for the physical and orbital parameters of the Kepler multis was known a priori then these would be the best priors, but given this is not known, uniform priors are an acceptable option. While not formal priors, limits were placed on both planetary densities and eccentricities. Densities were allowed to vary from $0.01\text{-}30\text{g}/\text{cm}^3$ and eccentricities were constrained to be between 0 and 1. The one exception to using uniform priors was for the quadratic stellar limb darkening parameters c_1 and c_2 . For these parameters, we used Gaussian priors with μ and σ calculated from the Python Limb Darkening Toolkit (LDTk) Parviainen & Aigrain. The values for $\log G$, T_{eff} , and metallicity which are used as inputs for LDTk were obtained from the Gaia catalog given in Berger et al.. To help broaden the limb darkening priors and allow for a larger exploration range, the errors returned by LDTk were multiplied by 10 before being passed to PhoDyMM.

2.2.3 Planet Selection

Out of the 8000 exoplanetary systems discovered by Kepler, only those containing multiple transiting planets were considered for this work. We used the planetary catalog given in Lissauer et al. (2024) to aid in filtering down the number of systems to be analyzed with PhoDyMM. Specifically, a system in the catalog was only added to our list of systems if there were multiple planets with valid non-negative periods and a high disposition score indicating they are believed to be true planets. Within the planetary catalog of Lissauer et al., each planet has four disposition scores which can be either “P” for positive planet candidate, “F” for false positive, “M” for Radial Velocity mass exceeds that of a planet, “R” for an unlikely large radius, “N” for not included in this catalog, and “S” for low SNR with the four scores coming from Lissauer et al., DR25 supplemental catalog, DR25 (Thompson et al. 2017), and DR24 (Coughlin et al. 2016). To be included in this work we required that the first value be either a “P”, “M”, or “R” or that the third score be a “P”. Planets that did not meet this criterion were removed from consideration and if this removal took a

system to less than two planets, then the system was disregarded as well. In the end, we were left with 1,824 planets distributed across 719 systems.

Graham & Ragozzine (in prep.) showed that there were at least 46 systems with TTVs where the TTVs could not be adequately predicted using the known planets in the system. In an attempt to better model these systems, they added an additional “hidden” planet and allowed PhoDyMM to solve for the best planetary parameters. After completing their analysis, Graham & Ragozzine found that 34 systems benefited from having the hidden planet added and as such we included these planets in our dataset, bringing our total number of planets up to 1858. To give this particular effort the best chance possible, we made use of the maximum likelihood posterior from Graham & Ragozzine in place of the initial conditions that would normally have been pulled from our starting catalog, but given these were not the result of an exhaustive search through parameter space, they did not guarantee that they held the correct parameters for the hidden planet. Still, it was worth adding these planets because excluding them would almost certainly result in a substandard fit to existing data.

2.2.4 Lightcurve Preparation

Lightcurves were taken from MAST and detrended as described in Ragozzine et al., in prep. In short, known transits were masked (including times from Holczer et al. 2016) and a polynomial was fit to the out-of-transit data with a width of 1440 minutes. This detrending length was chosen based on exploring various options. We also removed one Cotrending Basis Vector as we found through exploration of lightcurve scatter that this provided an optimal balance between detrending and overfitting. Regions near gaps are clipped. Even after detrending, there can be a few strong outliers. Though these do not affect the fit, we do a form of sigma-clipping where we first categorize datapoints by whether they are from out-of-transit data or within one of the known transiting planets. In each category, we calculate the median absolute deviation from the median and clip datapoints




Figure 2.1 Period-Radius plot for all the planets in the 719 systems modeled using PhoDyMM. These systems contain a total of 1858 planets with 34 of those being hidden planets that were added to their respective planetary system to better account for observed TTVs following the results of (Graham & Ragozzine in prep.). The clustering of planets with smaller periods is a known bias within transiting surveys as it is easier to detect these planets since they will transit more often and have a higher transit probability.

that are 15 times further away from the median than the equivalent Gaussian standard deviation. Inspection of the resulting data shows that extreme outliers are removed, but a few clear outlier points remain. Since these outlier points are relatively few (almost always less than 0.1% of all the data), we did not find that they affected our fits.

Kepler observed many stars using its 1-minute “Short Cadence” mode which provides more precise information on transit times especially for small planets. Through Quarter 9, it was possible to observe all systems with (then-known) multiple transiting planets with Short Cadence. After Quarter 9, Short Cadence target selection was potentially biased as the possibility of TTVs was included as a criterion for which targets were selected. To maintain the homogeneity of our results, we thus only include Short Cadence data (when it is available) through Quarter 9; afterwards, all systems use Long Cadence data even if Short Cadence is available. We note that this means that a photodynamical model that includes more Short Cadence data could reach a somewhat more precise mass measurement depending on the system. Even so, the general effect is that our inclusion of Short Cadence is a significant improvement over analyses that only use Long Cadence data (or TTVs derived from Long Cadence data).

2.2.5 PhoDyMM Input Generation

PhoDyMM requires two input files in order to run its analysis on a system. The first file is a planetary data input file (pldin) which contains the starting stellar and planetary parameters. For each planet the listed parameters are: period (p), epoch (T_0), eccentricity (e), inclination (i), longitude of ascending node (Ω), argument of periapsis (ω), mass (M_p), and radius (R_p). For the host star the specified parameters are: mass (M_{star}), radius (R_{star}), limb darkening coefficients (c_1, c_2), and the dilution fraction. Planetary parameters are obtained from the catalog published in Lissauer et al. which was created by taking the final Kepler data release catalog DR25supp, adding in KOIs from other sources and manually vetting the planet candidates, and then by using the TRANSITFIT5

software to obtain the transit models. This software performs an MCMC analysis to determine planet parameters, and as such the table given in Lissauer et al. gives both the median parameter values as well as the values from the maximum likely posterior. To ensure that the starting parameters for PhoDyMM are self-consistent we used those based of the posterior draw. Starting values for the orbital angles Ω and ω were not specified in the catalog and as such were initialized to 0. In the cases where M_p was not determined, we used the relationship $M_p/M_E = (R_p/R_E)^{2.06}$ which as explained in Lissauer et al. (2011) was obtained by fitting a power law to Earth and Saturn and provides good agreement to planets within that size range. The five stellar parameters were obtained using both the Gaia catalog from Berger et al. and the limb darkening code LDTk. From the catalog, the values for M_{star} and R_{star} were readily available and the additional parameters needed to calculate the limb darkening coefficients from LDTk were also used. In all cases, the dilution was set to 0.

The systems with hidden planets added to them also needed to be checked for valid radii. This is because Graham & Ragozzine focused primarily on dynamical interactions between planets, and there were several cases where the radius of the injected planet was fixed at 0. While starting our work with a radius of 0 would not be invalid, it would result in the planet being outside of the allowed density range. Therefore, in these cases, we used the published mass and applied the aforementioned relationship between mass and radius.

The second required input file contains the primary commands for how PhoDyMM should handle the system and includes information such as the number of steps and walkers to use, the number of planets, and what priors, if any, should be used, as well as parameters for controlling the behavior of the lightcurve model and n-body integrator. Of particular note for adjusting the n-body integrator is the integration error tolerance which was set to 1e-14 for all systems. In this file, it is also possible to specify which parameters are to be solved for and which are held fixed. Allowing for fixed parameters is useful, since it permits the user to explore simplified models such as those

where orbits are forced to be coplanar. Even if these simplifications are not fully justified from a physics standpoint, by employing them, it can greatly speed up the convergence time by reducing the dimensionality of the explored parameter space Ragozzine et al.. For this work, all planets had their Ω fixed at 0 and the jitter for all the stars was also fixed at 0.

An additional tool that was used, made accessible through PhoDyMM's second input file, is an option to disallow crossing orbits. The motivation for such a feature is that it is statistically unlikely for exoplanetary systems to be observed on the brink of a major upheaval, and once orbits are allowed to cross it takes very little time before planets either collide or there is a planet ejected from the system Lissauer & Gavino. Therefore, it is reasonable that if a walker attempts to take a step that would result in crossing planetary orbits, the step should be rejected. While it would be possible to enforce stricter stability requirements during the DEMCMC algorithm through the use of additional stability metrics such as those given in Gladman (1993); Hadden & Lithwick (2018); Petrovich (2015), it is possible that these may reject steps that are in reality stable and thus lead to inaccuracies in our results.

2.2.6 PhoDyMM Output

PhoDyMM raw output consists of several text files with three of them being of particular note. The first of these is the "demcmc_koiXXXX_cadence.out" file (where XXXX refers to a unique KOI index number) which is a log of the walker's current parameter values and is appended to every n steps as specified in the input file. From this file, the posterior distributions can be constructed factoring in the fact that because output is only saved every n steps, the posteriors are thinned by a factor of n. The second critical output file is the "mcmc_bestchisq_koiXXXX_cadence.aei" file. This file maintains a log of any step that improves the current χ^2 value which is important because these steps are not guaranteed to be contained in the demcmc out file. Lastly, the third important output file is the "gamma_koiXXXX_cadence.txt" file which tracks the acceptance fraction and

step size throughout the DEMCMC process. It is optimal if while exploring parameter space, the set of walkers have an acceptance fraction of 0.23 and a large enough step size that movement through parameter space occurs in a reasonable amount of steps. For these reasons, having the acceptance fraction and step size tracked is of use in answering questions about the success of a specific system's analysis.

Instead of performing a DEMCMC run using PhoDyMM, it is also possible to have PhoDyMM simply construct a model lightcurve given a set of initial parameters. By utilizing this process, we also are able to obtain as an output file, the model lightcurve that corresponds to set of parameters with the best χ^2 value. This file is saved as the "lc_koiXXXX_cadence.lcout" file and also contains the input lightcurve values to allow for easier comparison.

By utilizing some of our basic analysis codes, several additional output files and plots can easily be obtained. Specifically, by using a lightcurve analysis code, plots showing the original data lightcurve with marked transits, phase folded model versus data lightcurves, and O-C plots are constructed. There is also a DEMCMC analyze code which takes the "demcmc_koiXXXX_cadence.out" file and transforms it into a more readily useable format and generates plots showing the 2D parameter-parameter posteriors, 1D fully marginalized posteriors, and the movement of all the walkers for every parameter at each saved step which is helpful in quickly detecting any errors in how the walkers are exploring.

2.2.7 Transformation Matrix

In Tuchow et al. (2019) they point out that sampling algorithms such as DEMCMC face several challenges as they attempt to construct the posterior distributions for the variables of interest. Two challenges in particular are that parameters can vary in scale by several orders of magnitude and often parameters are highly correlated with each other. Both of these tend to slow down the convergence time because the sampler is not able to take reasonably sized steps. To alleviate these issues, we

follow the procedure set out in Tuchow et al. where they apply a linear coordinate transformation to rotate and scale the parameters space. Instead of steps being taken in the model parameter space p , steps are taken in the transformed space z where the two spaces are related by $z = \tilde{\Sigma}^{-1/2}(p - \tilde{p}_{means})$ where $\tilde{\Sigma}$ and \tilde{p}_{means} are estimates of the sample covariance matrix and the parameter means. Tuchow et al. found that by transforming their parameter space it was possible to obtain up to a 10x speed up in convergence time. For our work, we calculated the covariance matrix and means vector after completing our burn-in of 20,000 steps and once these were calculated all steps were taken in the transformed space. Unfortunately, in general, we did not see the same enhancement in convergence speed, but given the low computational cost of performing this transformation, we decided to apply it anyways given the inherent possibility of improving convergence times.

2.2.8 metaPhoDyMM

A challenge that this worked faced was how to apply the tools and procedures for a photodynamical model on 719 exoplanetary systems simultaneously with as little user input as possible. To aid in this endeavor, we designed a framework known as metaPhoDyMM. MetaPhoDyMM is a compilation of python scripts which can initialize, run, extend, and resume an analysis using PhoDyMM. Additionally, metaPhoDyMM is able to run several scripts which are able to determine whether a particular system has converged, failed, or needs to continue running. Each system proceeded in the same general flow through the metaPhoDyMM framework. Systems began with a small run of 25,000 steps with convergence or failure being evaluated after the first 5,000 steps and every subsequent 10,000 steps. Following this, the covariance and means were calculated allowing for future steps to be taken using the parameter space transformation. The status of a system continued to be evaluated every 10,000 steps until either a stopping condition of failure or convergence had been reached.

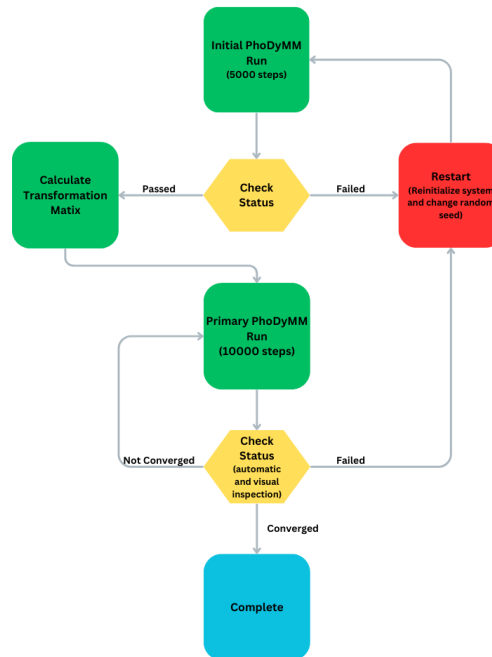


Figure 2.2 Basic flow diagram of how a system moves through the metaPhoDyMM framework. Each system begins with a small DEMCMC run of 25,000 steps before calculating the transformation matrix to de-correlate parameters. At this point the status is also checked to ensure that the system has not triggered a failure criterion. If a system does enter a failed state, it is restarted at the end of the supercomputer iteration. Systems that have not failed, proceed to run in cycles 10,000 steps where at the end of each cycle the questions of failure and convergence are once again addressed. Systems continued to run until either reaching a failed state or meeting the criteria for convergence (either automatically or by visual inspection).

Convergence Metrics

Convergence for any variety of MCMC is difficult to determine, and this especially holds true for DEMCMC. Often times, convergence is decided by a “by eye” analysis of walker and corner plots to evaluate if individual walkers have traversed the parameter space enough and if the posteriors look reasonably filled and smooth. A more robust method is to use the Gelman-Rubin statistic which compares the variance within a walker chain to the variance between walkers. Once this statistic reaches a value of 1 then convergence has been reached. Unfortunately, this statistic requires that walkers be independent of each other, which is not the case for DEMCMC. Given this, we employ a suite of our own metrics that, when all are met, provide convincing evidence the walkers have converged and a good solution has been found. In addition to looking at convergence specifically, we also made use of several metrics which were meant to give information about the quality of the best fit solution found by the DEMCMC analysis. This was done in order to give reasonable evidence that PhoDyMM was primarily exploring good regions of parameter space and that the resulting model built from the best parameters would indeed match the lightcurve data.

Before calculating any of these metrics, we remove the first 20,000 steps per walker if at least 25,000 steps have been taken, otherwise we remove the first 20% of steps. This is done because DEMCMC is known to have a warm-up/burn-in phase which generates samples that are dependent upon the initial parameter values. Final converged walkers/posteriors, will not depend on the starting parameters and therefore it is required to remove the burn-in from the data. The decision to set the burn-in at 20,000 steps was made after analyzing walker plots for a number of systems which indicated that in general the burn-in phase was on the order of 5-15,000 steps and so by removing 20,000 steps we added a slight buffer for systems that had a longer/slower burn-in.

- Out versus in residuals - this metric compares the mean of residuals in transit (as estimated by the model) to the mean of the residuals out of transit. Poor models are often characterized

by significant residual differences between data in and out of transit. We required the ratio of the out to in transit means to be greater than 0.75.

- Mean in residuals - To further investigate the quality of the model we calculate the mean of the in transit residuals between the data and our model lightcurves. To pass this metric, we require that the mean residuals be less than 0.2.
- P value between model and residuals - It is expected that if the model and the data are in good agreement, there should be no statistically significant correlation between the model and the residuals during transits. If there were a correlation, it would likely indicate that our model was failing to capture a fundamental aspect of the particular system. Ultimately, we required that the p-value between the model and the in transit residuals be greater than 0.01. While this is still a fairly small p-value which would nominally indicate there is statistical correlation between the residuals and the model, at this level you would still expect 1/100 trials to reach this level and have no true correlation. Since our work involved modeling thousands of planets, it became likely that several planets would have lower p-values while still having correlation between residuals and the model. Thus it was determined that the cutoff at 0.01 would still accomplish our purposes of catching significant issues in our analysis.
- Reduced χ^2 - We calculate the reduced χ^2 (χ_{red}^2) as a final check on the goodness for our model when the best found parameters are used. For this calculation we obtain the standard χ^2 between the input light curve and the synthetic one generated by PhoDyMM and divide it by the number of data points which is approximately equal to the degrees of freedom given the large number of data points. For this check we only require $\chi_{red}^2 < 5$ because the noise model is not always exact due to stellar variability.
- χ^2 slope - By the time the DEMCMC algorithm has converged, only minimal enhancements the best fit's likelihood should be made at each step. To require this to be the case, we

calculate the χ^2 slope for the best fit solution after removing the burn-in and if it is below our threshold of 0.001 then the system is deemed converged by this metric.

- One walker winning – Walkers are all expected to thoroughly explore the allowed parameter space and prioritize spending time in regions of high likelihood. However, it is possible that a single walker, might be able to find a region of parameter space significantly more favorable than the space being explored by the other walkers. In the case that this happens, we do not consider that convergence has been reached until a majority of the walkers have also explored this space. Formally speaking, we find the min χ^2 value for each walker $\chi_{min,i}^2$ and use these values to find the global min, χ_{min}^2 , and the median min value, χ_{median}^2 . Using these, we require that $\chi_{median}^2 - \chi_{min}^2 < 10$ thus indicating that there is not a walker that is maintaining a significantly better χ^2 value.
- Effective Sample Size (ESS) - We make use of public code published by Foreman-Mackey which provides a machine learning algorithm for predicting what the autocorrelation length is for a walker that has not traveled many autocorrelation lengths. This tool's ability to accurately predict the autocorrelation length is critical for this work because traditionally it is not possible to determine the ACL length for a walker until it has traveled 50 ACLs. Given that this level of exploration would be completely infeasible due to our average ACL, by employing this code we are able to continue our plan of having a large number of walkers travel a small number of ACLs. To get the ESS we took $(n_{steps} - n_{burn}) / ACL * n_{walkers}$. Once a system had an ESS of 1000 this metric was considered passed.

Failure Metrics

Our failed state for this work was a catch-all for systems that needed additional attention in order to continue progressing. The primary reason for ending up in this state was if walkers started to drift into regions of parameter space that would lead to highly implausible or impossible models for the

system. The rationale for this criterion is that once walkers make their way to unlikely parameter space, the likelihood surface can become incredibly flat, which would allow the walker to continue to explore worse and worse sets of model parameters. Once this occurs, walkers have been found to essentially become “lost” and it could require an infinite number of steps for the walker to return to the region of likely parameter space. We constructed a two stage test for whether or not this had occurred for T_0 , the impact parameter b , and R_p . While the other planetary parameters could still wander and lead to less reasonable values, by ensuring these parameters stayed within the specified ranges we required planets to remain transiting (or close to) in the synthetic lightcurves. As long as this was the case, PhoDyMM’s DEMCMC would be able to keep the walkers from getting completely lost. The first stage of the test was to see whether the best fit parameters had deviated considerably from the initial conditions. A particular system was flagged as potentially being in a failed state if any of the three following conditions were met for any of the planets:

- $\Delta T_0 > 0.5$
- $b > 1.5$
- $R_p > 25R_E$

Once flagged as potentially failing due to poor parameters in the best fit, the posteriors were checked to see if this was the case for more than 10% of the data. Given PhoDyMM’s probabilistic approach, it did not seem reasonable to fail a system based on a single step taken by one walker, which is what is contained in the best fit set of parameters. By also requiring at least 10% of the data to be in non-physical regions of parameter space, we allowed the walkers to move and explore, but failed the system when this exploration started to trend in an unfavorable direction that would compromise the final posteriors.

Another way that a system could be flagged as failed is if walkers stopped moving altogether. Given the aforementioned method for taking a step in the DEMCMC algorithm, it is possible that

walkers could eventually move to a region of parameter space where they can no longer find an allowed step to take. Once this has occurred for many walkers, no additional information is gained by running PhoDyMM longer and instead the system needs to be inspected. To check for this at the end of each computing cycle, we inspected the acceptance fraction of the last step, and if it was below 0.03, meaning only 6 out of the 200 walkers had an accepted step, then the system was decided to be in a failure state.

Finally, given that PhoDyMM was run on the BYU supercomputer that had time and memory limits, outages, and unexpected updates, it was also possible that a system could enter a failed state due to issues with the computational resources. For example, it was not wholly uncommon for a memory limit to be reached and as a result output could not be written to which then left PhoDyMM with fragmented files that were unusable for resuming PhoDyMM in the next metaPhoDyMM execution loop. These situations required output to be analyzed and cleaned up to remove any erroneous output so that the analysis could continue.

2.2.9 Restarting Systems

In the interest of completing this analysis for as many systems as possible, at the end of each computation cycle underperforming systems could potentially be determined to need restarting based off of several conditions. First, if the system was flagged as being in a failed state, then it was restarted from the current best fit solution. Second, if the system had one walker winning (OWW) by more than 15 χ^2 units, it was also restarted using the best fit solution. In the case of failed systems, it is possible that a best fit solution did not exist or was in an unacceptable region of parameter space and in these situations, the system was simply reinitialized. Since PhoDyMM allows the user to set the random number seed for reproducibility sake, it was also necessary to change this value whenever a system was reinitialized or else it would repeat the failure exactly as before. The third and last reason a system would be restarted is if metaPhoDyMM analysis codes repeated failed

during the computation cycle. These codes were robust when a system had acceptable output and PhoDyMM's analysis was working properly, when this was not the case, the analysis scripts also tended to fail. Therefore, repeated failures were seen as an indicator of an underlying problem requiring restarting.

2.2.10 Convergence by Visual Inspection

Our metrics for both catching failed systems and for determining convergence were primarily meant to allow for an automatic analysis of the 719 Kepler systems which was necessary because manually inspecting the output for every system after each cycle of computation would be physically impossible. However, after many of the easiest systems for PhoDyMM to analyze were completed, it became apparent that the metrics did not always give the most accurate representation of PhoDyMM's progress in constructing reliable posteriors. It was found that there were many systems where unmodeled features such as stellar variability or even simple outliers in the lightcurves could result in metrics not being passed. Given this, the option for a system to be deemed converged based on visual inspection became available. In these cases, output files were analyzed to determine if the model fit the data to reasonable standards and if the walkers had thoroughly explored parameter space leading to smooth, well-filled posterior distributions. This became especially important later on as it was found that the ACL code being used to calculate our ESS was not robust against outlier chains and the amount of data thinning which was necessary in order to keep this analysis computationally tractable. Systems that were passed based on this visual inspection are indicated in the data table.

2.2.11 Computational Resources

PhoDyMM was built to be parallelizable where individual walkers could run on their own CPU. This, along with the fact that each system is completely independent from each other, meant that

our photodynamical analysis of the Kepler systems could efficiently and effectively be performed on a supercomputer. We used the supercomputer located at Brigham Young University, which offers over 40,000 CPUs and 180 Tb of memory. Each system of the 719 systems were given 200 CPUs with a maximum wall clock time of three days (the maximum wall clock time allowed to be given for a single job). Given an average allocation of 7,000 CPUs this meant that all of the systems could be analyzed in approximately two months. This several week to 1-2 month computation cycle wherein all the systems had three days on the supercomputer was defined as a computing iteration. It was originally expected that each system would need approximately a week of computation time, which would require 2-3 iterations over the course of five months and would make use of over 1 million CPU days. This proved to be an underestimation as many systems were required to take millions of walker steps while other systems were repeatedly restarted at best fit solutions in order to ensure that the parameter space near the best solution was explored completely. In the end, this work required approximately 11 iterations over the course of a year and a combined total of $\sim 10^{11}$ walker-steps to be taken.

2.3 Results

Our results are presented in two parts. The first is our results related to the success of this analysis as a whole. For this, we look at the number of converged and failed systems, as well as the distributions of some of the metric values which were used in determining whether a system had reached convergence. Secondly, we present our Kepler Dynamical Catalog in three tables. These tables give the parameter values with their 1σ errors, the parameter values for the best fit posterior draw, and the values for 1000 randomly drawn posteriors. Additionally, we give some of the system information such as stellar parameters and convergence notes in Table 2.4. Using draws from our

posteriors as well as the best fit posterior, we are also able to give our version of the mass-radius distribution for the systems in our catalog.

2.3.1 Convergence Results

Of our original sample of 719 Kepler multis, 661 systems were determined to be converged either by automatically passing our convergence metrics or through visual inspection of the output files. Of the remaining 58 systems: 52 were classified as failed, as per our failure criteria, and 6 were simply unfinished/unconverged, but not in a failed state. As such, these systems will not be included in any of the following plots and tables, but the solutions found can be made available upon request. An in-depth analysis of what made these 58 systems difficult for PhoDyMM to solve went beyond the scope of this work. It is expected that several systems failed because initial parameters were not finely tuned enough, and if the system were to be restarted with more accurate planetary parameters, it would be possible for PhoDyMM to properly analyze the system.

Of more immediate interest than what caused the 58 systems to not reach convergence, was the question of whether these systems were statistically different from the 661 that converged without issue. To answer this, we employed a two sample KS test comparing the distributions for periods, radii, multiplicities, and level of TTVs between these two subsamples. The null hypothesis of a KS test is that the two samples are drawn from the same population, as the p-value obtained from this test decreases in value, the null hypothesis can be rejected at higher and higher confidence levels. The results of these KS tests are given in the first row of Table 2.1. As can be seen, there is enough difference in the distributions of periods, radii, and multiplicities between the failed and completed systems that it can be concluded that they were drawn from different populations. Looking specifically at multiplicities, Table 2.2 shows the number of systems with a given number of planets for both the failed and completed systems. Based off these numbers, it is evident that the KS test results are correct and that there is considerable difference between the subset of failed

	Period	Radii	Multiplicity	TTVs
Failed vs Complete	0.017	0.001	0.006	0.155
Complete vs All	0.999	0.997	0.999	0.999

Table 2.1 P-values resulting from two sample KS tests comparing the distributions of periods, radii, multiplicities, and strength of TTVs between the failed and completed as well as the completed systems to the entire dataset of 719 systems. The smaller the p-value, the greater the ability to reject the null hypothesis that these samples were drawn from the same population. Based off of the p-values comparing the failed and completed systems, there is strong statistical evidence that the periods, radii, and multiplicities were drawn from separate populations. This seems to indicate that the 58 failed systems are not simply a random subsample, but that there is some underlying principle that makes certain systems more likely to fail. Fortunately, when looking at the strength of the TTVs, there is little significant evidence that the failed and completed systems were drawn from different populations. This is important because it would be concerning if, systems with strong TTVs, and therefore better constraining capabilities, were statistically more likely to fail our photodynamical analysis. Of additional importance is the second row of the table which shows the p-values obtained from comparing the completed systems to the entire dataset. Here we see that the null hypothesis cannot be rejected at any statistically significant level and thereby there is no evidence suggesting that the completed systems represent a sample from a different population than the dataset as a whole. In turn it can be concluded that using the 661 completed systems instead of the 719 systems will not introduce additional biases in subsequent analyses.

systems and the completed ones. Primarily, the failed systems appear to have an over representation of the systems with 3-4 planets and an under representation of the 2 planet systems. Despite this and the implied differences in populations given by the tests on radii and periods, we are not concerned that our results are significantly biased. This belief is reinforced by the results of a second round of KS tests comparing our completed systems with the entire dataset of 719 systems which are shown in the second row of Table 2.1. For each of these KS tests, there is essentially no ability to reject the null hypothesis that these two subsamples originated from the same population. Thus, we are confident that the completed systems can be used as a single catalog without introducing any new biases into subsequent analysis.

	2	3	4	5	6	7	8
Failed Systems	20	22	8	2	0	0	0
Completed Systems	413	166	52	24	4	1	1

Table 2.2 Occurrence numbers based off of system multiplicities between the subset of failed and completed systems. Given that the number of failed systems is $\sim 1/10$ the number of completed systems, it would be expected that at each multiplicity an approximate factor of 10 would sufficiently link the two subsets. What is instead seen is an under representation of 2 planet systems and an over representation of 3-4 planet systems. While not a perfect match to expectations, the differences are not large enough to cause significant concern. The underlying cause of this behavior may receive some attention when the failed systems are explored further.

Additionally, to highlight some of the overall trends in our PhoDyMM analysis, we display in Figure 2.3 the distributions for four metrics which are the number of walker-steps per system, the out/in transit residuals, the reduced χ^2 between the model and the data, and the final step size walkers took for each system. Starting with the first subplot for the number of walker-steps, it is interesting to find that it is a fairly flat distribution. At the start of this work, it was expected that there would be more of an exponential drop-off of systems needing larger and larger numbers of steps, but this is not necessarily what was found. The plots of both the out/in transit residuals and the reduced χ^2 are essentially as expected. Given our convergence criteria, these values were constrained to be a very specific range indicating that we had achieved a good fit to the data. Points outside of this allowed range, which for the out/in residuals is anything below 0.75 and for the reduced χ^2 is anything above 5, have made it in our data due to visual inspections for convergence. Finally, looking at the step size distribution, it is also an unsurprising result that most of the systems trended towards the minimum allowed steps size of 0.03. This behavior has not been fully characterized and understood, but is likely a result of the difficulty in exploring the high-dimensional parameter space.

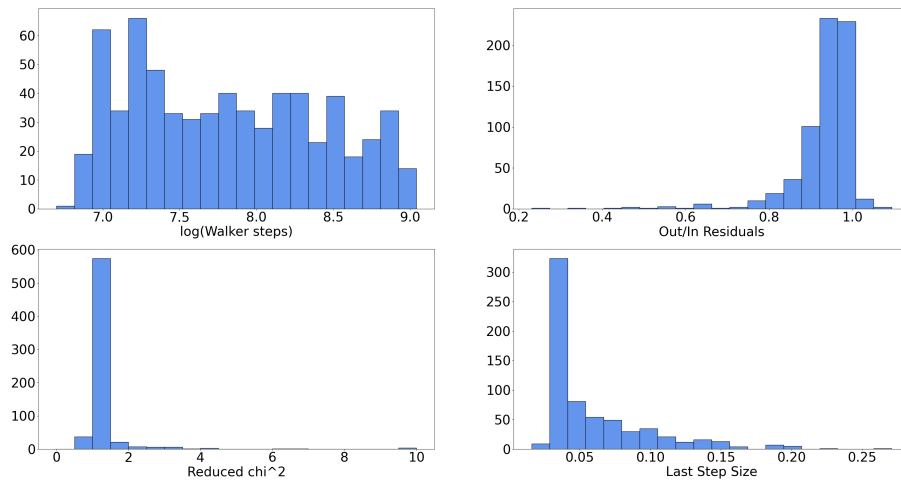


Figure 2.3 Final distributions for four metrics used in addressing convergence and analysis progress. The first plot shows the total number of walker-steps needed before the system reached convergence. Next is shown the distribution of the out/in transit residuals and the reduced χ^2 values between the model lightcurve and the Kepler data. The fourth metric is the last step size taken by the walkers. In the case of the out/in residuals and the reduced χ^2 these distributions are expected given our requirements for convergence. For step size, it is also not surprising that most systems moved towards the step size lower limit as this was a behavior observed throughout this work and is mostly attributed to the dimensionality of parameter space. Of most interest is the distribution for the number of walker-steps needed. From testing, it was predicted that each system could need up to 200 million walker steps, but it was also expected that there would be more of an exponential drop off of systems needing more and more steps which is not what is seen here.

2.3.2 Kepler Dynamical Catalog

Beginning with Table 2.3 we give our median values plus 1σ values for the period, epoch (T0), eccentricity (e), inclination (i), mass, radius, and density. As mentioned above, Ω was fixed at 0 and we do not include our values for ω here, since it is of slightly less interest in the solution for the system. This table has the primary purpose of showing the level which we were able to constrain physical and orbital parameters. It is clear that there are several planets where individual parameters were able to be tightly constrained and others where nearly the entire parameter space is within our error bars. The primary reason for this is that the majority of Kepler systems do not have detectable planet-planet interactions. While Kepler had incredible precision, it was not designed to measure masses and planet-planet interactions very quickly diminish when planets have orbits that are not within a factor of ~ 2 in period.

Unlike previous analyses that studied primarily systems where there was evidence of dynamical interactions, we studied all systems. Consequently, we found a range of cases from strong dynamical interactions that led to well-constrained parameters to weak dynamical interactions that were previously missed to no detectable dynamical interactions. In the most common case of no dynamical interactions, the system often explored the full prior range of densities ($0.01 - 30 \text{ g cm}^{-3}$) and eccentricities (up to crossing orbits).

It must be noted that the values in Table 2.3 are only the medians and confidence intervals for each parameter, they should not be used to initialize a model for a planet or system. The true output from PhoDyMM is posterior draws which are a self-consistent set of parameters that match the Kepler data and thus can be used to initialize a model (discussed below). If only the median values were used, it is possible that the starting parameters would actually lead to a bad fit to the data and start an analysis in an unlikely region of parameter space.

As a companion to our table of planetary parameters, Table 2.4 gives information and results for each system as a whole. Here we give the stellar mass and radius plus 1σ errors, the planetary

multiplicity for the system, whether or not the system had a hidden planet added (1 = yes, 0 = no), an indicator for whether the system was determined to be converged based off of a visual inspection (1=yes, 0=no) and the total number of walker steps that PhoDyMM used to analyze the system. For cases where a hidden planet was added to the system, the planetary multiplicity reflects this addition and will therefore not necessarily match other documentation that only lists the number of confirmed planets in the system. Also of note, the total number of walker steps is given by the number of steps taken by a single walker multiplied by the number of walkers used. Given that 200 walkers were used to analyze each system and that it was not uncommon for a single walker to take >1 million steps, a system with a total number of walker steps $> 10^8$ is also not unusual for this work.

Finally, Table 2.5 gives the same set of parameter values as shown in Table 2.3 (excluding densities), but in this case these parameters are those for the posterior draw that had the best log-likelihood value and thus represents the best fit to the starting lightcurve. No errors are included in these values, as they represent a single draw from the posterior distribution for each parameter and not the distribution itself. For any analysis which requires a single set of starting planetary parameters, these would be the values to start with. While these values allow one to construct a self-consistent model that out of all our posteriors matches the data the best, it should still be noted that the full value of this work is in the construction of the Bayesian posterior distribution for each of the modeled parameters. Due to this, it is more accurate and a better practice to use many posterior draws from our results and perform subsequent analysis in a statistical framework. To make this possible, as part of our results repository, we include a table with 1000 posterior draws per system. This table or any subsample is not included here as providing ~ 10 posterior draws for a single planet is not particularly useful.

One particular goal of this work was to improve the current understanding of the mass-radius distribution for the Kepler multistars. To this aim, we present Figure 2.4. Beginning with the first

KOI	Period (d)	T_0	e	i	M_E	R_E	Density(g/cm^3)
41.01	12.816 $^{+0.0}_{-0.0}$	786.454 $^{+0.0}_{-0.0}$	0.078 $^{+0.018}_{-0.036}$	89.501 $^{+0.353}_{-0.484}$	33.708 $^{+13.746}_{-11.322}$	2.295 $^{+0.045}_{-0.043}$	15.339 $^{+6.191}_{-5.063}$
41.02	6.888 $^{+0.0}_{-0.0}$	782.431 $^{+0.002}_{-0.002}$	0.023 $^{+0.032}_{-0.016}$	86.771 $^{+0.403}_{-0.231}$	2.185 $^{+3.268}_{-1.574}$	1.373 $^{+0.032}_{-0.034}$	4.641 $^{+7.18}_{-3.35}$
41.03	35.336 $^{+0.002}_{-0.002}$	758.309 $^{+0.003}_{-0.004}$	0.104 $^{+0.045}_{-0.058}$	88.501 $^{+0.099}_{-0.09}$	9.855 $^{+8.392}_{-7.016}$	1.627 $^{+0.046}_{-0.044}$	12.697 $^{+10.986}_{-9.107}$
46.01	3.488 $^{+0.0}_{-0.0}$	791.008 $^{+0.001}_{-0.0}$	0.05 $^{+0.046}_{-0.033}$	84.138 $^{+0.397}_{-0.39}$	263.67 $^{+124.633}_{-138.11}$	6.463 $^{+0.186}_{-0.184}$	5.354 $^{+2.713}_{-2.823}$
46.02	6.051 $^{+0.009}_{-0.011}$	789.011 $^{+0.03}_{-0.02}$	0.036 $^{+0.037}_{-0.025}$	85.317 $^{+0.403}_{-0.387}$	3.801 $^{+3.41}_{-2.514}$	1.396 $^{+0.079}_{-0.077}$	7.677 $^{+7.261}_{-5.094}$
70.01	10.854 $^{+0.0}_{-0.0}$	787.979 $^{+0.0}_{-0.0}$	0.018 $^{+0.021}_{-0.013}$	89.7 $^{+0.211}_{-0.272}$	15.915 $^{+12.09}_{-10.015}$	2.81 $^{+0.049}_{-0.045}$	3.956 $^{+3.004}_{-2.495}$
70.02	3.696 $^{+0.0}_{-0.0}$	791.941 $^{+0.0}_{-0.0}$	0.02 $^{+0.023}_{-0.013}$	88.306 $^{+0.461}_{-0.343}$	7.861 $^{+10.837}_{-5.735}$	1.725 $^{+0.034}_{-0.032}$	8.431 $^{+11.778}_{-6.162}$
70.03	77.612 $^{+0.001}_{-0.0}$	718.621 $^{+0.001}_{-0.001}$	0.165 $^{+0.093}_{-0.115}$	89.825 $^{+0.117}_{-0.11}$	12.972 $^{+25.5}_{-10.106}$	2.154 $^{+0.052}_{-0.041}$	7.12 $^{+13.94}_{-5.545}$
70.04	6.098 $^{+0.001}_{-0.001}$	788.558 $^{+0.002}_{-0.002}$	0.031 $^{+0.028}_{-0.021}$	88.466 $^{+0.35}_{-0.243}$	0.591 $^{+0.824}_{-0.43}$	0.782 $^{+0.023}_{-0.023}$	6.81 $^{+9.544}_{-4.966}$
70.05	19.579 $^{+0.002}_{-0.001}$	792.573 $^{+0.007}_{-0.005}$	0.029 $^{+0.02}_{-0.018}$	88.88 $^{+0.073}_{-0.097}$	1.224 $^{+0.81}_{-0.733}$	0.899 $^{+0.035}_{-0.031}$	9.208 $^{+6.178}_{-5.476}$
72.01	0.837 $^{+0.0}_{-0.0}$	788.167 $^{+0.0}_{-0.0}$	0.101 $^{+0.248}_{-0.076}$	83.704 $^{+0.962}_{-1.089}$	9.164 $^{+6.083}_{-6.187}$	1.491 $^{+0.032}_{-0.031}$	15.254 $^{+10.187}_{-10.313}$
72.02	45.294 $^{+0.0}_{-0.0}$	751.094 $^{+0.0}_{-0.0}$	0.22 $^{+0.261}_{-0.158}$	89.437 $^{+0.252}_{-0.12}$	37.407 $^{+25.849}_{-25.577}$	2.401 $^{+0.089}_{-0.078}$	14.904 $^{+10.331}_{-10.191}$
82.01	16.146 $^{+0.0}_{-0.0}$	778.165 $^{+0.0}_{-0.0}$	0.019 $^{+0.02}_{-0.014}$	89.294 $^{+0.086}_{-0.08}$	1.71 $^{+1.258}_{-0.95}$	2.285 $^{+0.04}_{-0.038}$	0.788 $^{+0.588}_{-0.438}$
82.02	10.312 $^{+0.0}_{-0.0}$	778.591 $^{+0.0}_{-0.0}$	0.027 $^{+0.029}_{-0.019}$	89.229 $^{+0.173}_{-0.135}$	0.914 $^{+0.425}_{-0.305}$	1.21 $^{+0.021}_{-0.021}$	2.836 $^{+1.345}_{-0.947}$
82.03	27.454 $^{+0.0}_{-0.001}$	764.37 $^{+0.001}_{-0.002}$	0.018 $^{+0.021}_{-0.013}$	89.269 $^{+0.043}_{-0.043}$	1.043 $^{+0.482}_{-0.489}$	0.938 $^{+0.021}_{-0.02}$	6.952 $^{+3.247}_{-3.238}$
82.04	7.071 $^{+0.0}_{-0.0}$	787.193 $^{+0.001}_{-0.001}$	0.017 $^{+0.029}_{-0.013}$	89.529 $^{+0.312}_{-0.341}$	0.471 $^{+0.393}_{-0.306}$	0.595 $^{+0.013}_{-0.012}$	12.348 $^{+10.309}_{-8.053}$
82.05	5.287 $^{+0.0}_{-0.0}$	787.878 $^{+0.002}_{-0.002}$	0.021 $^{+0.034}_{-0.017}$	89.515 $^{+0.332}_{-0.395}$	0.218 $^{+0.204}_{-0.152}$	0.466 $^{+0.012}_{-0.012}$	11.873 $^{+11.094}_{-8.264}$
82.06	22.409 $^{+0.001}_{-0.002}$	775.092 $^{+0.003}_{-0.004}$	0.02 $^{+0.024}_{-0.014}$	88.789 $^{+0.034}_{-0.037}$	0.082 $^{+0.111}_{-0.061}$	0.676 $^{+0.029}_{-0.028}$	1.461 $^{+2.054}_{-1.084}$
85.01	5.86 $^{+0.0}_{-0.0}$	785.813 $^{+0.0}_{-0.001}$	0.031 $^{+0.044}_{-0.021}$	87.26 $^{+0.379}_{-0.317}$	5.125 $^{+1.902}_{-1.634}$	2.691 $^{+0.069}_{-0.069}$	1.45 $^{+0.563}_{-0.469}$
85.02	2.155 $^{+0.0}_{-0.0}$	788.394 $^{+0.0}_{-0.001}$	0.221 $^{+0.115}_{-0.089}$	84.634 $^{+1.218}_{-0.63}$	9.986 $^{+5.343}_{-5.116}$	1.502 $^{+0.045}_{-0.046}$	16.284 $^{+8.921}_{-8.42}$
85.03	8.132 $^{+0.0}_{-0.0}$	786.54 $^{+0.0}_{-0.0}$	0.028 $^{+0.038}_{-0.019}$	87.435 $^{+0.221}_{-0.205}$	3.567 $^{+1.036}_{-0.975}$	1.616 $^{+0.041}_{-0.041}$	4.654 $^{+1.456}_{-1.312}$
89.01	84.687 $^{+0.001}_{-0.001}$	761.077 $^{+0.001}_{-0.001}$	0.067 $^{+0.021}_{-0.013}$	89.779 $^{+0.14}_{-0.122}$	85.627 $^{+38.604}_{-37.445}$	2.994 $^{+0.09}_{-0.076}$	17.623 $^{+7.986}_{-7.821}$
89.02	207.622 $^{+0.022}_{-0.025}$	638.013 $^{+0.025}_{-0.022}$	0.163 $^{+0.039}_{-0.022}$	89.523 $^{+0.016}_{-0.019}$	10.048 $^{+7.616}_{-6.141}$	3.934 $^{+0.117}_{-0.103}$	0.91 $^{+0.689}_{-0.56}$

Table 2.3 Median and 1σ errors for the periods, epochs, eccentricities, inclinations, masses, radii, and densities of all the planets in the systems that reached convergence. Given that the set of these parameters for an individual planet is not the result of a single posterior draw, the reader is strongly cautioned against using these values to initialize a model. This table value instead gives a quick overview of the expected parameter values for each planet.

KOI	M_*	R_*	N planets	Hidden planet	Hand passed	Walker steps
41	$1.062^{+0.09}_{-0.081}$	$1.548^{+0.028}_{-0.028}$	3	0	0	53000000
46	$1.122^{+0.047}_{-0.045}$	$1.716^{+0.042}_{-0.042}$	2	0	0	252610000
70	$0.924^{+0.04}_{-0.043}$	$0.89^{+0.014}_{-0.014}$	5	0	1	366278000
72	$0.918^{+0.066}_{-0.063}$	$1.074^{+0.021}_{-0.021}$	2	0	0	57678000
82	$0.75^{+0.032}_{-0.033}$	$0.747^{+0.01}_{-0.01}$	6	0	0	240702000
85	$1.235^{+0.058}_{-0.066}$	$1.458^{+0.031}_{-0.03}$	3	0	1	383410000
89	$1.632^{+0.074}_{-0.076}$	$1.565^{+0.039}_{-0.035}$	2	0	0	158234000

Table 2.4 Additional information relating to each Kepler system. Given are the stellar mass and radius values with their 1σ errors, the system multiplicity which includes hidden planets if one was added to the system, a key for whether or not a system had a planet added (0=No, 1=yes), an indicator of whether the system met convergence criteria based off of visual inspection (0=No, 1=yes), and lastly the total number of walker-steps used for analysis.

column of the figure, we show our mass-radius distribution using 100 posterior draws for each planet. We provide this data using two ranges for the plot axes to both highlight the catalog as a whole, and to zoom in on the range of smaller planets. Given that radii tend to be well-known/constrained, when plotting our posterior draws in mass-radius space, the results tend to be bands of points where the length of the band corresponds to the level of uncertainty in the mass measurement. This is what is primarily seen in Figure 2.4 with any difficulty in seeing the individual bands being caused by the overlaying of points from the hundreds of planet posteriors being simultaneously plotted. However, it is still clear that there are some denser regions of the plot where points are in a small cluster instead of a band. These clusters represent the planets that were able to have their masses constrained considerably and thus did not have a wide range of allowed mass values. In addition to plotting the posterior draws, we also show calculated mass-radius distributions for eight planetary compositional structures which include Earth-like planets, water worlds, planets with 1% H/He by mass, and planets with either a constant 1 g/cm^3 or 30 g/cm^3 density. These distribution

KOI	Period (d)	T_0	e	i	M_E	R_E
41.01	12.8158	786.4538	0.0861	89.8632	48.4749	2.2984
41.02	6.8888	782.4273	0.0149	86.6496	0.7853	1.377
41.03	35.3399	758.2988	0.1181	88.4173	1.2869	1.6461
46.01	3.4876	791.0072	0.0472	84.0243	306.603	6.4816
46.02	6.0527	789.0017	0.0348	85.2098	1.8499	1.4624
70.01	10.8542	787.9785	0.0357	89.5481	37.6911	2.7891
70.02	3.6962	791.9408	0.0147	88.2109	20.3502	1.7136
70.03	77.6132	718.6208	0.2645	89.8049	0.4935	2.1379
70.04	6.0979	788.5548	0.0432	88.4261	0.1545	0.7709
70.05	19.5818	792.5763	0.0418	88.9287	1.2549	0.8738
72.01	0.8375	788.1677	0.3376	84.0004	16.9312	1.4631
72.02	45.2944	751.0927	0.6908	89.1609	20.5533	2.4554
82.01	16.1461	778.1646	0.0438	89.2411	1.0184	2.3095
82.02	10.312	778.5902	0.0106	89.2268	0.5765	1.2156
82.03	27.4537	764.3702	0.047	89.2884	0.9611	0.9402
82.04	7.0713	787.1922	0.0258	89.89	0.9105	0.5916
82.05	5.2868	787.8776	0.032	89.9245	0.5245	0.4714
82.06	22.4096	775.0919	0.0675	88.8299	0.0043	0.6696
85.01	5.86	785.8128	0.0698	87.0468	6.4219	2.6738
85.02	2.1548	788.392	0.4845	83.9063	4.3911	1.485
85.03	8.1319	786.5393	0.0779	87.3501	4.0649	1.6052
89.01	84.6862	761.0776	0.0567	89.7069	120.7303	2.9925
89.02	207.6435	637.991	0.1403	89.5226	12.5293	3.9097

Table 2.5 Planet parameters as given in Table 2.3. In this case, the set of parameters are obtained from the posterior draw that had the highest likelihood evaluation and thus best fit to the data. Since the parameters come from a single posterior draw, they represent a self-consistent model and can therefore be used as initial parameter values for any additional analysis. Additionally this is the reason why the values have no error ranges associated with them.

curves are aimed at guiding visual inspection and can help to show that given a particular posterior point or cluster of points, what interior compositions are likely or unlikely. While the mass-radius distributions given in the first column are an accurate representation of the results of this work, due to planets that have a high uncertainty in their mass measurements and the large number of planets in our catalog, individual planets tend to get lost in the sea of data points. To alleviate this and highlight planets that could have their masses constrained considerably, the second column of Figure 2.4 shows only planets that have an uncertainty in their mass measurements of 25% or less where the uncertainty % is given by

$$\%Uncertainty = median(abs(mass - median) * 1.4826 / (median)) \quad (2.1)$$

With this constraint, it is now possible to see the 104 planets that have uncertainties less than 25% and what can further be seen is that they fill a wide range of the parameter space. Finally, the third column of Figure 2.4 shows the mass-radius distribution using only the masses and radii from our posterior with the maximum likelihood. Similar to Table 2.5, this is shown in order to display a quick and clear overview our findings at the cost of the statistical information that is made available when many posterior draws are plotted together. Unsurprisingly, there is no clear structure within the distribution of points that would indicate some overwhelmingly favorable composition. In order to gain the greatest amount of insight from this distribution, further work is needed to determine whether underlying trends exist within individual systems which in turn may yield information about formation and evolutionary mechanisms.

Using our equation for the mass uncertainty % given in Equation 2.1, Figure 2.5 shows our uncertainties for all the planets in our dataset of converged systems. We measure 104 planet masses to an uncertainty of 25% or less, which in turn means that these planets have a mass measurement that disallows 0 at 4σ and meets the criteria of a statistically significant mass measurement. While not affecting the planets that had meaningful constraints on their masses, it is expected that many of our higher uncertainty percent's actually represent a lower limit on the uncertainty as there were

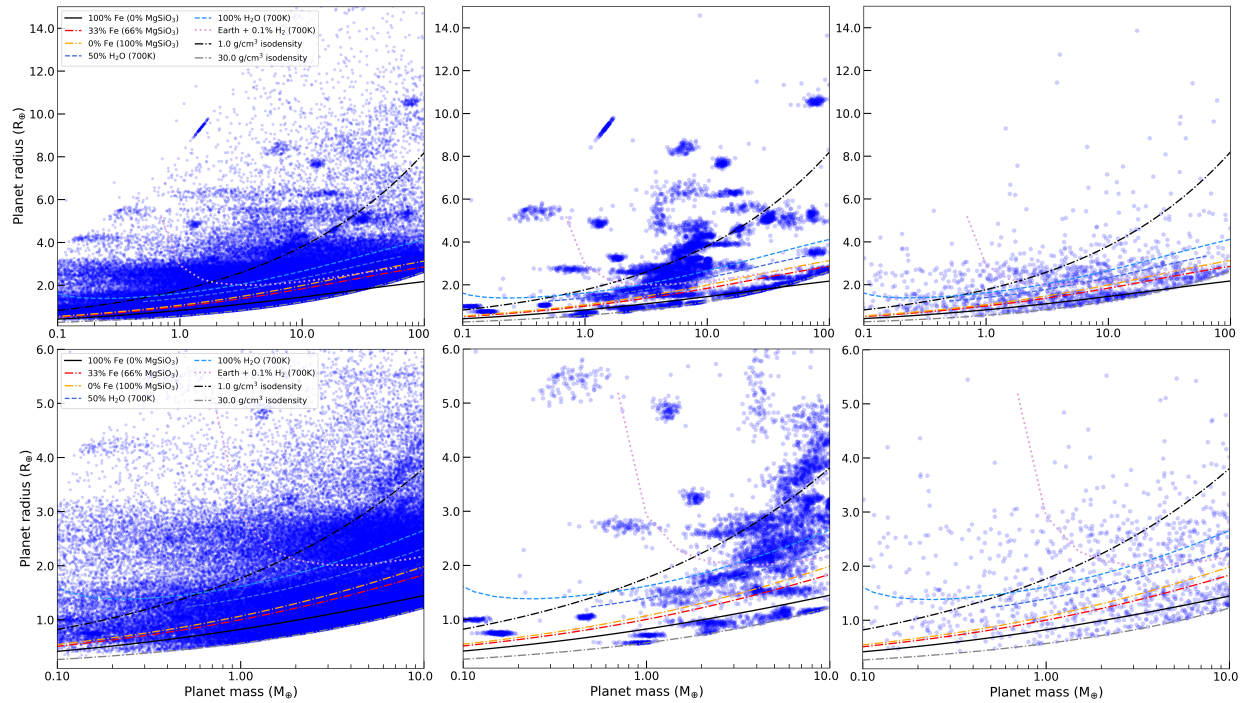


Figure 2.4 Column 1: Using 100 posterior draws for each planet, we construct the Mass-radius diagram for our Kepler Dynamical Catalog. Uncertainties in mass measurements result in bands of points with the length of the band directly correlating to the level of uncertainty in the mass measurement. Planets that have more certain masses are therefore seen as clusters of points. To begin looking into theoretical compositions for each planet, we also plot the mass-radius curves for eight theoretical compositions. While it is beyond the scope of this work to answer the question of what the interior composition is for each planet, based off of where the posterior draws are located in the diagram, it is at least possible to gain some insight into what elemental structures are likely present. Column 2: To better present the planets with well constrained masses, we give the same Mass-radius distribution as shown in column 1, but only for planets that have a mass uncertainty of 25% or less. It can now be seen that among these planets are several extremely low dense super-puffs as well as those that are likely to be very dense as they fall along the boundary of our density prior. Column 3: As a final interpretation of the Mass-radius distribution for our catalog, we plot the mass and radius values using only the best fit posteriors. This then represents the distribution of mass and radius values that would be used for any catalog level analysis where only a single posterior draw is used per system.

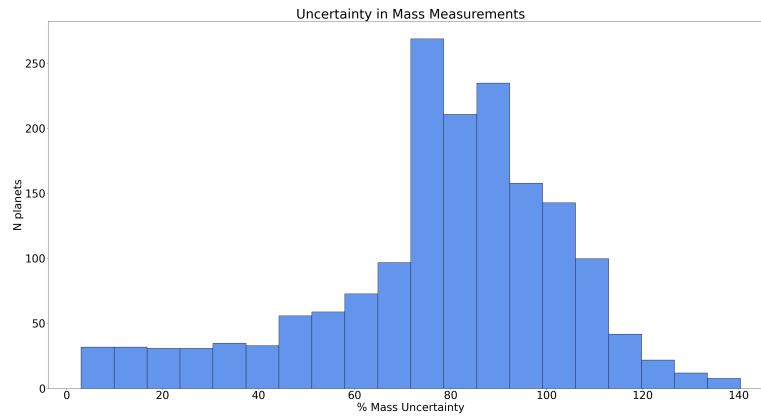


Figure 2.5 Obtaining constrained mass measurements was a primary goal for this work as masses are critical in calculating densities which thus gives information related to composition, formation, and evolution. Here we give the distribution of mass uncertainties for the planets in the 661 converged systems. We obtained 104 mass measurements that have an uncertainty of 25% or less meaning that a mass of 0 is disallowed at the 4σ level. The peak in uncertainties around 80% is not necessarily representative of the actual uncertainty in our mass measurement as it is possible that in some of these cases the exploration of possible mass values was limited by our density priors. Fortunately, this is not of much concern as it does not affect the planets with well constrained masses




Figure 2.6 Averaged 1σ errors on the density measurements for all of the planets in completed systems. As can be seen there are 139 planets where the averaged uncertainty is less than $0.5g/cm^3$ thereby indicating that the entire 1σ error range on the measurement is less than $1g/cm^3$. Providing well constrained density measurements is of particular interest as they can then be used to inform and improve planet interior models.

weak priors placed on the mass via our max density of $30g/cm^3$ and our min of $0.01g/cm^3$. We do not present the exact number of planets that had insignificant enhancements of their mass constraints due to our priors. To further highlight the amount of information gained by our photodynamical analysis, we show the averaged 1σ errors in the density measurements for each planet. This distribution, shown in Figure 2.6, is interesting and noteworthy for several reasons. The first is that it is bimodal, where the two peaks occur essentially in the first and last bins indicating that in general there is either enough information in the Kepler data to constrain densities well or almost no information and the 1σ error range spans nearly the entire density prior. Second, this distribution shows that for 139 planets, the average 1σ error limit is less than $0.5g/cm^3$ meaning that the entire error range only spans a difference of $1g/cm^3$. This is especially encouraging as densities are an important input value for planet interior models and the more precise we can constrain the planet's density the better the interior structure can be constrained.

2.4 Discussion

With our results given in the each of the data tables, it is necessary to address the level of success we obtained in the context of other works that have been performed. To accomplish this, we first look at our performance compared to other planetary parameter catalogs of a size similar to the Kepler Photodynamical Catalog that we have constructed. This is useful because we did not require the same level of attention given to each system as would have been the case if only a handful of systems were being studied. Therefore, a study that was only focused on a single system could reasonably obtain more precise results than those given in our catalog. By comparing our results to other larger catalogs, it is expected that this effect will be reduced considerably as these catalogs would also have had to balance precision with completeness. Following this, for a subset of systems of interest, we address the level of agreement between our results and those published previously.

Beyond just comparing the results, we also look at our performance in fitting a couple of additional systems that have been found to have interesting TTVs by (Kane et al. 2019). Our ability to correctly capture dynamical information, such as TTVs, is important, as TTVs are the primary method of breaking degeneracies in the data. However, since TTVs are not directly fitted in PhoDyMM it is possible that our best fit solution does not result in an exceptional TTV fit. Therefore, it is beneficial to see some of PhoDyMM's TTV models compared to the observational data. Lastly, we address the success level of our work in modeling systems with hidden planets added to them in an attempt to better predict their TTVs.

2.4.1 Comparison to Other Catalog Level Analyses

We compare our results with those presented by Hadden & Lithwick (2017), Judkovsky et al. (2024), and (Ofir et al. 2025) as these are three fairly recent studies that also analyzed a larger number of Kepler systems thus allowing us to look at the quality of our results on a broader scale. Hadden & Lithwick performed a TTV analysis of 55 Kepler multis which contained a total of 145 planets, Judkovsky et al. made use of their analytic method called AnalyticLC and applied it to 101 planets in 23 systems with the final result of obtaining significant mass measurements for 95 planets, and Ofir et al. used a simplified photodynamical model to fit 64 systems containing 218 planets which yielded 88 mass measurements. Our comparison is shown in 2.7 where our only selection criterion was that we had converged posteriors for the particular system which left us with a sample of 78 systems. Points without error bars represent upper limit mass estimates as given in their respective papers. Two readily observable trends are: first, with few exceptions, our results are in good agreement with those found previously; and second, in many cases our uncertainties are larger especially when considering the log-scaling for the y-axis. Still, there are a significant number of systems where we both obtain a mass measurement in good agreement with what has previously been reported and with tighter uncertainties on the measurement such as for KOI-250,

523, and 1576. Furthermore, we have mass measurements for all of the planets in each of these 78 systems, where in some cases, such as KOI-94 or KOI-1831, other studies did not calculate masses for every planet. Taking this one step further, is the fact that we have mass measurements (albeit sometimes with considerable uncertainties) for all of the planets in an additional 583 systems not shown in this figure. This once again highlights the point that even some of the largest analyses of the Kepler multis are still quite small compared to the Kepler Dynamical Catalog presented here.

2.4.2 Individual Systems of Note

While it would be ideal if we could compare our results to those published for every system, given the number of systems in our catalog, it is not possible. Instead, we present a condensed comparison between our results and recent publications for a selection of 12 systems. These systems were selected as being of particular interest due to having either well characterized TTVs, higher multiplicities, smaller signal-to-noise planets, or an overall popularity/familiarity in the exoplanet community. For each system, we compare our results to 1-2 of the most recent publications focusing on the agreement level between with mass measurements and when possible the measurements for eccentricity and inclination.

KOI-41/Kepler-100

KOI-41/Kepler-100 has three confirmed transiting planets and while it is not a particularly interesting system by way of dynamics, it was the primary system used for testing within our analysis and is therefore of note. Kepler-100 has been solved for by Marcy et al. (2014) and by Judkovsky et al. (2022). Marcy et al. obtains a mass measurement for Kepler-100b of $M_b = 7.34^{+3.2}_{-3.2} M_{\oplus}$ and upper limits on masses of $M_c = 4.8 M_{\oplus}$ and $M_d = 1.7 M_{\oplus}$ for planets c and d. Judkovsky et al. is able to obtain tighter constraints using their AnalyticLC results which gave masses of $M_b = 5.1^{+1.7}_{-1.8} M_{\oplus}$, $M_c = 14.6^{+2.6}_{-3} M_{\oplus}$, and $M_d = 1.1^{+0.5}_{-0.5} M_{\oplus}$. We found that PhoDyMM ultimately did not

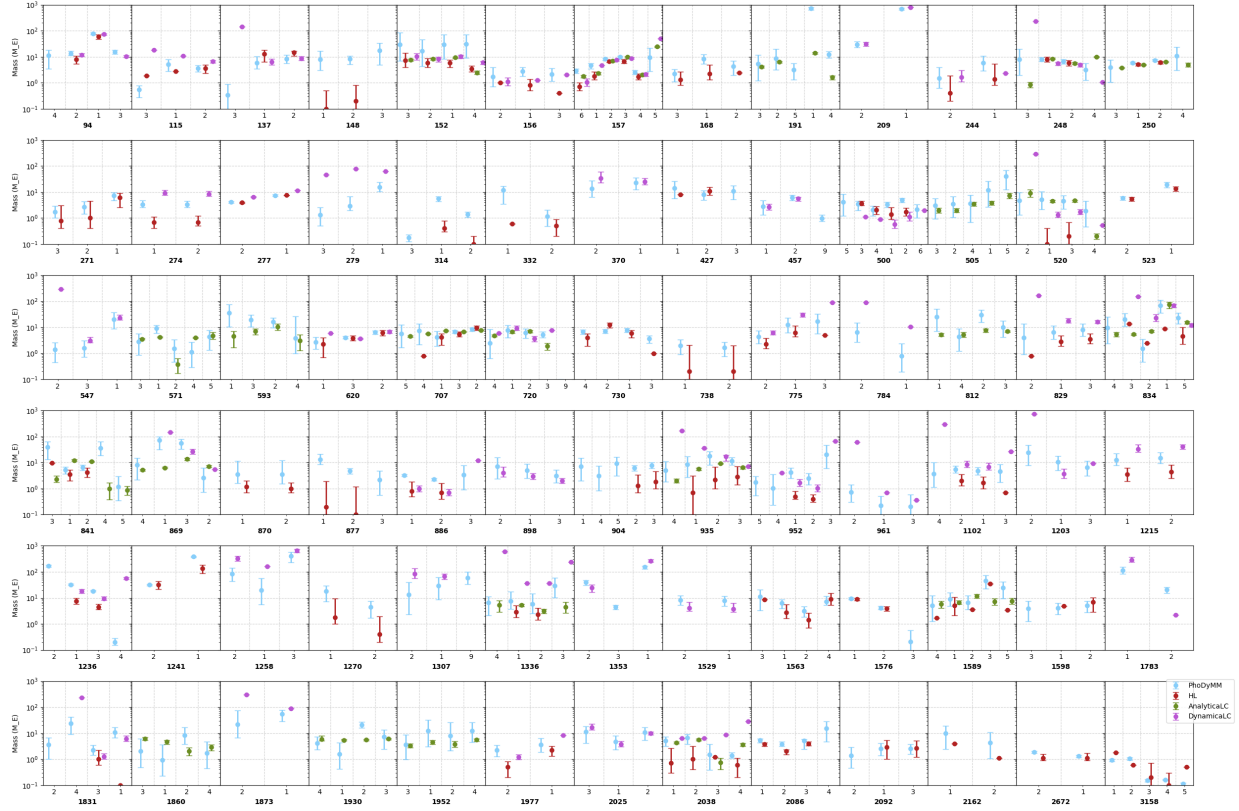


Figure 2.7 Mass values plus 1σ errors for 78 systems analyzed both in this work and in at least one of the studies performed by Hadden & Lithwick (2017), Judkovsky et al. (2024), or Ofir et al. (2025). It is beneficial to see how PhoDyMM performs compared to larger studies as it is well understood that we are not necessarily able to outperform a study focusing on only a couple of systems which can thus invest significant amounts of time in fine tuning inputs and procedures to obtain the very best results. What is seen here is that in nearly every case, PhoDyMM agrees well with previously published results. In general, but not always, our errors tend to be larger than those calculated by Hadden & Lithwick, Judkovsky et al., and Ofir et al.. Of the three other studies, we perform the best compared the results published by Hadden & Lithwick which is highlighted by inspection of the results for KOI-2092, 2672, and 3158. The good agreement combined with many cases where we provide tight constraints on the masses help to confirm the effectiveness of this work which is further magnified by the fact that we have parameter values for all of the planets in these systems and the remaining converged systems and not just for those that have a particularly interesting feature such as well defined TTVs.

recover significant mass measurements for this system, with our masses being $M_b = 2.18^{+3.27}_{-1.57}M_{\oplus}$, $M_c = 33.71^{+13.75}_{-11.32}M_{\oplus}$, and $M_d = 9.85^{+8.39}_{-7.16}M_{\oplus}$. While our masses are not fully inconsistent with Judkovsky et al. due to our large errors (except for planet c which differs by 2σ) it is more accurate to conclude that PhoDyMM did not provide meaningful constraints on the masses. These previous studies do not report limits on orbital parameters and thus no comparisons can be made to our results which are listed in Table 2.3.

KOI-70/Kepler-20

KOI-70/Kepler-20 is a tightly packed system of six planets where the outermost planet's orbit is interior to the location of Mercury's orbit within our own solar system. While the true multiplicity of the system is six, Kepler-20g has only been confirmed via RV analysis and has not been observed to transit. Due to this, it is not included in our results. Using RV data Gautier et al. (2012) obtain masses for Kepler-20b,c, and d of $M_b = 8.7^{+2.2}_{-2.2}M_{\oplus}$, $M_c = 16.1^{+3.5}_{-3.5}M_{\oplus}$, $M_d = 20.1M_{\oplus}$ where the mass for planet d is a 2σ upper limit on the mass. Using additional RV measurements, Buchhave et al. (2016) is able to obtain more precise measurements for the masses of planets b, c, and d with their results yielding masses of $M_b = 9.70^{+1.41}_{-1.44}M_{\oplus}$, $M_c = 12.75^{+2.17}_{-2.24}M_{\oplus}$, $M_d = 10.07^{+3.97}_{-3.70}M_{\oplus}$. Our photodynamical model gave masses of $M_b = 7.86^{+10.84}_{-5.73}M_{\oplus}$, $M_c = 15.91^{+12.09}_{-10.01}M_{\oplus}$, $M_d = 12.97^{+25.50}_{-10.11}M_{\oplus}$. Our errors are considerably larger than those from past studies; however, those results are obtained from RV measurements which are naturally going to be better at determining masses. It is still encouraging and worth mentioning that our median measurements fall within 1.5σ of these other published results. Buchhave et al. also obtain eccentricity measurements for planets b and c as well as inclination measurements for all the planets. For eccentricities, they found $e_b = 0.02^{+0.09}_{-0.03}$ and $e_c = 0.16^{+0.01}_{-0.09}$ whereas our results had $e_b = 0.02^{+0.02}_{-0.01}$ and $e_c = 0.018^{+0.02}_{-0.01}$. These values match quite well with some disagreement for planet c which is of smaller concern given Buchhave et al.'s error range is larger. In the case of inclinations Buchhave et al. reports

$i_b = 87.36_{-1.59}^{+0.22}$, $i_e = 87.62_{-0.13}^{+1.09}$, $i_c = 89.82_{-0.63}^{+0.04}$, $i_f = 88.79_{-0.07}^{+0.43}$, $i_d = 89.71_{-0.05}^{+0.17}$ and our results are $i_b = 88.31_{-0.34}^{+0.46}$, $i_e = 88.47_{-0.24}^{+0.35}$, $i_c = 89.70_{-0.27}^{+0.21}$, $i_f = 88.88_{-0.10}^{+0.07}$, $i_d = 89.82_{-0.11}^{+0.12}$ which show considerably good agreement.

KOI-137/Kepler-18

KOI-137/Kepler-18 is a three planet system first studied by Cochran et al. (2011) with follow-up studies done by Hadden & Lithwick (2017) and Weiss et al. (2024) and is noted to have planets near a 4:2:1 Laplace resonant chain as well as planets c and d showing significant TTVs. Using a combination of RV measurements and TTVs, Cochran et al. found masses for the three planets of $M_b = 6.9_{-3.4}^{+3.4}M_{\oplus}$, $M_c = 17.3_{-1.9}^{+1.9}M_{\oplus}$, $M_d = 16.4_{-1.4}^{+1.4}M_{\oplus}$. The solutions in Cochran et al. were based on Kepler photometry from quarters 0 to 6 assuming zero eccentricity; this was before the mass-eccentricity degeneracy was understood. Hadden & Lithwick used the full Kepler data and did not assume circular orbits and found masses of $M_c = 12.9_{-6.6}^{+5.6}M_{\oplus}$, $M_d = 14.9_{-4.2}^{+1.8}M_{\oplus}$ with no update in the mass measurement for planet b. Using additional RV measurements and without using TTVs, Weiss et al. found masses of $M_b = 12.9_{-4.1}^{+4.1}M_{\oplus}$, $M_c = 19.5_{-5.1}^{+5.1}M_{\oplus}$, $M_d = 24.0_{-6.5}^{+6.5}M_{\oplus}$. Our analysis in turn found a solution with significantly smaller masses where $M_b = 0.36_{-0.24}^{+0.57}M_{\oplus}$, $M_c = 5.95_{-2.49}^{+4.31}M_{\oplus}$, $M_d = 8.26_{-2.78}^{+3.38}M_{\oplus}$. These lower masses are consistent with Jontof-Hutter et al. (2022) (who added TESS data) and Ofir et al. and are only $\sim 1\text{-}\sigma$ off from the Hadden & Lithwick measurements. It is possible that the RV mass uncertainties are underestimated since there are not that many RV measurements and the planets are close to resonance which could lead to possible unaccounted for correlations between parameters. We confirmed that our fits look good, our stellar parameters are similar to previous studies, and we explored a range of parameter space. Kepler-18 is known to have Transit Duration Variations that are probably not accessible to our model because we force all longitudes to be 0. However, we explicitly tested using synthetic data based on this specific

system that this assumption would not bias the mass measurements. Our novel determination that the innermost $1.7 R_{\oplus}$ planet has a low mass strongly inconsistent with a rocky density is interesting.

KOI-157/Kepler-11

KOI-157/Kepler-11 is a six planet system of particular note because the five inner planets are near a series of first and second order resonances as explained by Hadden & Lithwick (2017). This system also displays chopping TTVs which allow the mass-eccentricity degeneracy to be broken for several of the planets. Of further interest is the fact that these planets have been shown to all have low densities. Lissauer et al. (2013) thoroughly analyzes this system and finds masses of $M_b = 1.9^{+1.4}_{-1.0}M_{\oplus}$, $M_c = 2.9^{+2.9}_{-1.6}M_{\oplus}$, $M_d = 7.3^{+0.8}_{-1.5}M_{\oplus}$, $M_e = 8.0^{+1.5}_{-2.1}M_{\oplus}$, $M_f = 2.0^{+0.8}_{-0.9}M_{\oplus}$, $M_g < 25M_{\oplus}$. Follow-up work done by Hadden & Lithwick published masses of $M_b = 0.7^{+0.3}_{-0.2}M_{\oplus}$, $M_c = 1.8^{+0.9}_{-0.5}M_{\oplus}$, $M_d = 6.8^{+0.7}_{-0.8}M_{\oplus}$, $M_e = 6.7^{+1.2}_{-1.0}M_{\oplus}$, $M_f = 1.7^{+0.5}_{-0.4}M_{\oplus}$ with no updated mass for planet g. We also found that masses could be constrained considerably for the first five planets with the last planet essentially having no constraint from the data. Our masses are $M_b = 2.91^{+0.55}_{-0.55}M_{\oplus}$, $M_c = 4.61^{+1.09}_{-1.11}M_{\oplus}$, $M_d = 8.14^{+0.74}_{-0.73}M_{\oplus}$, $M_e = 9.90^{+1.01}_{-0.98}M_{\oplus}$, $M_f = 2.53^{+0.47}_{-0.44}M_{\oplus}$, $M_g = 9.62^{+12.02}_{-6.88}M_{\oplus}$. Our results are generally consistent within $1-\sigma$. Since Lissauer et al. also gives eccentricity measurements it is valuable to compare their results to ours. Lissauer et al. report values of $e_b = 0.045^{+0.068}_{-0.042}$, $e_c = 0.026^{+0.063}_{-0.013}$, $e_d = 0.004^{+0.007}_{-0.002}$, $e_e = 0.012^{+0.006}_{-0.006}$, $e_f = 0.013^{+0.011}_{-0.009}$ and $e_g < 0.15$. We calculate eccentricities of $e_b = 0.015^{+0.016}_{-0.009}$, $e_c = 0.01^{+0.014}_{-0.007}$, $e_d = 0.009^{+0.005}_{-0.005}$, $e_e = 0.012^{+0.003}_{-0.004}$, $e_f = 0.012^{+0.005}_{-0.005}$ and $e_g = 0.026^{+0.037}_{-0.02}$ which are very similar to those found by Lissauer et al..

KOI-274/Kepler-128

KOI-274/Kepler-128 is a two planet system highlighted by the analysis done by Leleu et al. (2023) using the RIVERS approach. The reasoning for analyzing this system even after a TTV analysis

had been performed was to support their finding that in some cases the extraction method for TTVs led to their amplitudes being underestimated which in turn led to masses and densities being underestimated as well. Leleu et al. derives masses of $M_b = 3.79_{-0.66}^{+0.76}M_{\oplus}$ and $M_c = 3.38_{-0.59}^{+0.67}M_{\oplus}$ for the two planets. The results from PhoDyMM are $M_b = 3.34_{-0.79}^{+1.43}M_{\oplus}$ and $M_c = 3.26_{-0.73}^{+1.27}M_{\oplus}$ which is perfectly in agreement with Leleu et al.'s results albeit that our upper errors are approximately 2x. While the goal would be to have as tight of measurements as possible, this is also balanced by the fact that 719 systems were analyzed simultaneously and there was almost no tuning of individual systems to enhance results.

KOI-277/Kepler-36

KOI-277/Kepler-36 is a two planet system of interest as it was originally studied using an earlier version of our PhoDyMM codes and has thus been used as a test case while developing PhoDyMM further. From a dynamics standpoint it is also interesting system as the planets are near a 7:6 MMR. Using their DEMCMC modeling routine, Carter et al. (2012) find masses of $M_b = 4.45_{-0.27}^{+0.33}M_{\oplus}$ and $M_c = 8.08_{-0.46}^{+0.60}M_{\oplus}$. After reaching convergence with PhoDyMM we calculate the masses for the system at $M_b = 4.08_{-0.26}^{+0.27}M_{\oplus}$ and $M_c = 7.49_{-0.47}^{+0.50}M_{\oplus}$. These two solutions line up quite well with our analysis favoring slightly smaller masses and obtaining marginally smaller errors on those measurements. The original work on Kepler-36 only made use of the long cadence Kepler data whereas we make use of both the long and short cadence data which is a reasonable explanation for the shift in mass values and how we are able to obtain improved constraints on the masses.

KOI-351/Kepler-90

KOI-351/Kepler-90 is one of the more exciting systems to have converged without issue using PhoDyMM as it is an eight planet system which is the highest multiplicity in the Kepler dataset. It was originally announced as a seven planet system by several independent groups including Cabrera

et al. (2014), Rowe et al. (2014), and Schmitt et al. (2014) with the eighth planet being confirmed by an independent photometric search by Shallue & Vanderburg (2018) Weiss et al. (2024). Weiss et al. is able to provide upper limit masses for the planets Kepler-b,c,i,d,e, and f as well as actual mass measurements for planets g and h. Their findings are as follows: $M_b = 4.3M_{\oplus}$, $M_c = 6.5M_{\oplus}$, $M_i = 4.5M_{\oplus}$, $M_d = 7.2M_{\oplus}$, $M_e = 6.6M_{\oplus}$, $M_f = 6.9M_{\oplus}$, $M_g = 23_{-25}^{+25}M_{\oplus}$, and $M_h = 236_{-33}^{+33}M_{\oplus}$. Our calculations for all of the eight masses are: $M_b = 0.72_{-0.54}^{+1.41}M_{\oplus}$, $M_c = 1.73_{-1.25}^{+3.18}M_{\oplus}$, $M_i = 2.17_{-1.64}^{+2.78}M_{\oplus}$, $M_d = 5.52_{-3.18}^{+4.64}M_{\oplus}$, $M_e = 6.49_{-2.98}^{+3.76}M_{\oplus}$, $M_f = 3.20_{-1.92}^{+1.94}M_{\oplus}$, $M_g = 13.40_{-1.16}^{+1.01}M_{\oplus}$, $M_h = 252.43_{-15.57}^{+13.27}M_{\oplus}$. While our error ranges for the six inner planets are larger, we do manage to fall within the upper limits provided by Weiss et al.. Additionally, we obtained significantly tighter mass constraints for the outer two planets that are consistent with those found previously. To our knowledge this is the first photodynamical analysis of Kepler-90 and as such find it meaningful to also present our phase folded lightcurves and our model TTVs versus those estimated by the DR25 pipeline. From our phasefolded lightcurve in Figure 2.8 it can be seen that we obtain a quality fit for all eight planets. The 14.4 day planet Kepler-90i has the poorest model fit out of the set, which is understandable given that it was the most difficult planet to find in the first place. We achieve an exceptionally good fit for the outer two planets Kepler-90g and h which again makes logical sense as these are the largest planets in the system. Looking at the comparison between our model TTVs and those from Rowe & Thompson (2015) shown in Figure 2.9 it is a bit of a mixed result. There is passable agreement for planets d and e with most of our model points lying within the errors of the data points. For planets b and c, the errors on the TTV measurements are so large that it is difficult to say whether or not we obtain a good model for the TTVs. Our models for planets g and h are subpar, however there are not many data points for comparison. Lastly, for planets i and f, there is essentially no data to compare to.




Figure 2.8 Phasefolded model versus data plots for the 8 planets in KOI-351. To date our work is the first full photodynamical model of the 8 planet system which also is the highest multiplicity system in the Kepler dataset. Here the gray points represent the individual Kepler measurements which are then binned together and plotted in black. Our unbinned model points are plotted in magenta and the red points are the binned model. Below each phasefolded curve are the residuals between the model and data. As can be seen our model performs well at fitting the transits for all 8 planets with the most difficulty coming in fitting the small transits from the 14.45 day planet. Our residual plots reaffirm this agreement as the individual values tend to be small and the plots do not show significant structure which would indicate a systematic error in fitting the transits. It is also worth noting just how exact PhoDyMM is able to model the transits for the outer two planets with the models having a near perfect match to the Kepler lightcurves.

KOI-377/Kepler-9

KOI-377/Kepler-9 has three transiting planets, two of them are Saturn-sized and near a 2:1 orbital resonance, these were also the first planets that were confirmed using TTV analysis Holman et al. (2010). The third planet is a short period super-Earth that is decoupled from the other two planets in the system Weiss et al. (2024). There have been several published results for the masses of the three planets using both TTVs and RV measurements which were originally inconsistent. Later investigation found that the discrepancy can be attributed to poorly reduced RV measurements in Holman et al. as later RV observations from Borsato et al. (2019) are fully consistent with the TTVs. They arrived at masses of $M_b = 43.4^{+1.6}_{-2.0} M_{\oplus}$, $M_c = 29.9^{+1.1}_{-1.3} M_{\oplus}$ these results also align with those found by Hadden & Lithwick. Of the 719 systems that we analyzed using PhoDyMM there were 52 that entered a failed state as set by our failure metrics and this was one of them. It is likely that the extensive and precise TTVs were related to the failure as our initial guess or

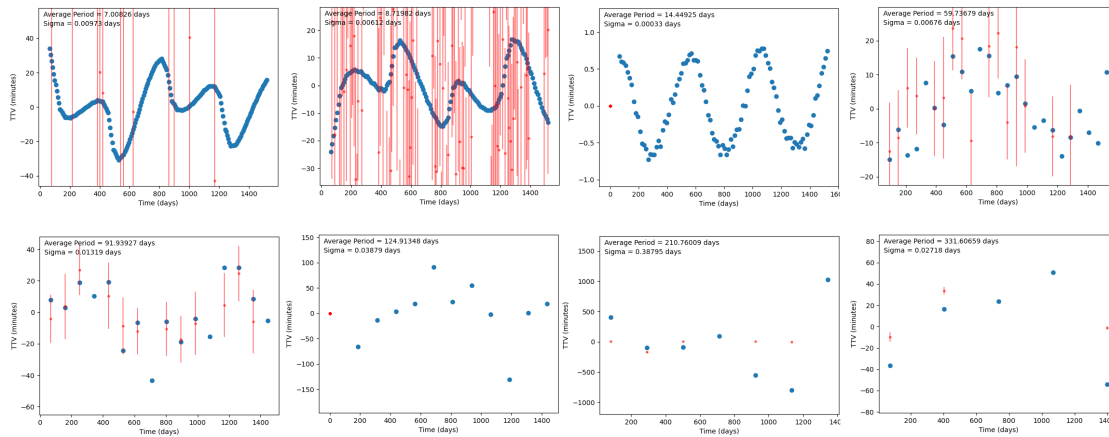


Figure 2.9 TTV plots for each planet in KOI-351. The red points with their errors are the observed TTVs obtained from the DR25 pipeline where the blue points are those generated by PhoDyMM using the best fit posterior draw. Very little is gained by looking at the TTVs for the first two planets as the errors on the observed TTVs span the full range of our model TTVs. There is also no TTV data for the third and sixth planets. For the 59.74 and 91.94 day planets (fourth and fifth), we see a good level of agreement between our model and the observational TTVs. While the agreement is aided by the larger error bars, in many many cases our model points fall very near the center of the error range. Finally for the outer two planets with periods of 210.76 and 331.6 days we find an overall disagreement especially for the 210.76 day planet. Given how well our model matches in the phasefolded plot, it is unlikely that our TTV model is as off as these results claim, therefore, it may be necessary to consider that the TTV measurements for these two planets are poorly estimated.

circular orbits was actually quite far from the expected solution. Our results for this system did not reach a final acceptable value, but our intermediate and incomplete estimates have masses of $M_b = 57.45^{+21.48}_{-8.63} M_{\oplus}$, $M_c = 39.79^{+48.93}_{-6.13} M_{\oplus}$ for the Saturn-sized planets and $M_d = 0.04^{+1.30}_{-0.04} M_{\oplus}$ which does maintain the mass ratio calculated by Borsato et al. and indicates that PhoDyMM is on its way to a plausible solution.

KOI-500/Kepler-80

KOI-500/Kepler-80 holds six super-Earth planets where at least four have been shown to exist in a 4:6:9:12 resonant chain, which made it the first exoplanetary system with a confirmed three-body resonant chain of longer than three planets. Some effort has been put into determining whether planet g is also included in the chain, but MacDonald et al. (2021) favors the solution with only the four middle planets in a resonant chain. Additionally, when MacDonald et al. solved for the planetary parameters for this system, they used an earlier version of PhoDyMM which gives us the ability to directly compare the quality of results between these two analyses. Beginning with masses, MacDonald et al. reports $M_d = 5.95^{+0.65}_{-0.6} M_{\oplus}$, $M_e = 2.97^{+0.76}_{-0.65} M_{\oplus}$, $M_b = 3.50^{+0.63}_{-0.57} M_{\oplus}$, $M_c = 3.49^{+0.63}_{-0.57} M_{\oplus}$, and $M_g = 0.065^{+0.044}_{-0.038} M_{\oplus}$ with no mass reported for planet f. Our work found masses of: $M_f = 4.08^{+4.27}_{-2.89} M_{\oplus}$, $M_d = 3.42^{+1.14}_{-1.42} M_{\oplus}$, $M_e = 2.02^{+0.89}_{-0.82} M_{\oplus}$, $M_b = 3.27^{+0.85}_{-0.71} M_{\oplus}$, $M_c = 4.88^{+1.13}_{-0.99} M_{\oplus}$, and $M_g = 2.12^{+1.26}_{-1.11} M_{\oplus}$. Our results are largely consistent with those found by MacDonald et al. with the biggest differences in planets d and c. We also have larger errors on each measurement, but they are still small enough to constitute a significant mass measurement in most cases. There is also a considerable disagreement between the mass measurements for planet g, but seeing as MacDonald et al. reports a radius of $R = 1.05^{+0.22}_{-0.24} R_{\oplus}$ our measurement on the mass seems to be much more reasonable. For eccentricities, MacDonald et al. found $e_f = 0.186^{+0.083}_{-0.049}$, $e_d = 0.0041^{+0.0037}_{-0.0028}$, $e_e = 0.0035^{+0.0032}_{-0.0024}$, $e_b = 0.0049^{+0.0036}_{-0.0032}$, $e_c = 0.0079^{+0.004}_{-0.0037}$, and $e_g = 0.1303^{+0.0034}_{-0.0037}$. The eccentricities found in

this work were $e_f = 0.419_{-0.069}^{+0.109}$, $e_d = 0.0056_{-0.0039}^{+0.0075}$, $e_e = 0.0084_{-0.0049}^{+0.0072}$, $e_b = 0.0059_{-0.0038}^{+0.005}$, $e_c = 0.0053_{-0.0035}^{+0.0051}$, and $e_g = 0.0152_{-0.0093}^{+0.0157}$. Thus, as a whole we found slightly larger eccentricities, but given how small each planet's eccentricity is, these results agree well with each other. Finally for inclinations, MacDonald et al. calculates $i_f = 85.99_{-0.52}^{+0.48}$, $i_d = 89.24_{-0.37}^{+0.46}$, $i_e = 88.59_{-0.16}^{+0.15}$, $i_b = 88.989_{-0.085}^{+0.09}$, $i_c = 88.744_{-0.046}^{+0.049}$, and $i_g = 88.26_{-0.07}^{+0.15}$. The inclinations calculated in this work were $i_f = 84.02_{-0.33}^{+0.27}$, $i_d = 88.18_{-0.31}^{+0.42}$, $i_e = 88.16_{-0.18}^{+0.22}$, $i_b = 88.53_{-0.12}^{+0.14}$, $i_c = 88.48_{-0.076}^{+0.087}$, and $i_g = 88.47_{-0.065}^{+0.083}$ which again shows good agreement with previous results.

KOI-730/Kepler-223

KOI-730/Kepler-223 is a four planet system where adjacent pairs of planets are in or near a first-order MMR Hadden & Lithwick (2017). This system was thoroughly analyzed by Mills et al. (2016) who performed their work using DEMCMC which is the same algorithm that PhoDyMM implements. Mills et al. calculates the following masses for each of the planets: $M_b = 7.4_{-1.1}^{+1.3}M_{\oplus}$, $M_c = 5.1_{-1.1}^{+1.7}M_{\oplus}$, $M_d = 8.0_{-1.3}^{+1.5}M_{\oplus}$, $M_e = 4.8_{-1.2}^{+1.4}M_{\oplus}$. Our work obtained masses of $M_b = 6.68_{-1.12}^{+1.28}M_{\oplus}$, $M_c = 7.07_{-1.17}^{+1.3}M_{\oplus}$, $M_d = 7.89_{-1.29}^{+1.29}M_{\oplus}$, $M_e = 3.67_{-1.15}^{+1.17}M_{\oplus}$ which fortunately agrees with Mills et al. at the 1σ level. For eccentricities, we also find good agreement with Mills et al.'s results as they find $e_b = 0.078_{-0.017}^{+0.015}$, $e_c = 0.15_{-0.051}^{+0.019}$, $e_d = 0.037_{-0.017}^{+0.018}$, $e_e = 0.051_{-0.019}^{+0.019}$ compared to our $e_b = 0.065_{-0.027}^{+0.027}$, $e_c = 0.039_{-0.026}^{+0.024}$, $e_d = 0.035_{-0.019}^{+0.018}$, $e_e = 0.034_{-0.017}^{+0.017}$. While we calculate a lower eccentricity for planet c, the rest of our eccentricities match up.

KOI-3158/Kepler-444

KOI-3158/Kepler-444 is a system of five mars sized planets orbiting within 0.1 AU from their star, which also has a wide-separation binary companion Weiss et al. (2024). This system was studied by Mills & Fabrycky (2017) who once again made use of a DEMCMC algorithm to obtain posterior distributions for the physical and orbital parameters of the planets. They only

report final masses for planets d and e, which are $M_d = 0.036_{-0.02}^{+0.065} M_{\oplus}$, and $M_e = 0.034_{-0.019}^{+0.059} M_{\oplus}$. From our application of PhoDyMM we are able to recover mass measurements for all five planets which are $M_b = 0.91_{-0.111}^{+0.087} M_{\oplus}$, $M_c = 1.04_{-0.095}^{+0.105} M_{\oplus}$, $M_d = 0.15_{-0.015}^{+0.016} M_{\oplus}$, $M_e = 0.16_{-0.009}^{+0.010} M_{\oplus}$, $M_f = 0.11_{-0.019}^{+0.013} M_{\oplus}$. For planets d and e, our errors are significantly smaller than those given by Mills & Fabrycky unfortunately, our median mass measurements are $1 - 2\sigma$ off from theirs. However, our median values are more consistent with those found by Hadden & Lithwick who derived masses of $M_d = 0.2_{-0.1}^{+0.5} M_{\oplus}$, and $M_e = 0.1_{-0.1}^{+0.2} M_{\oplus}$. Thus, it is possible that the masses for planets d and e are closer to the $M = 0.1 M_{\oplus}$ mark.

KOI-4772/Kepler-1705

KOI-4772/Kepler-1705 is a system with two confirmed planets and an unconfirmed candidate. It has primarily been studied by Leleu et al. (2021) and unfortunately, there is a disagreement between their initial conditions and ours. Leleu et al. have planets b and c at 9.03 days and 11.28 days whereas our inputs from Lissauer et al. have the periods at 9.02 days and 5.15 days along with a candidate planet at 3.38 days. Given our difference in initial conditions (likely due to a period alias), it is difficult to give a 1:1 comparison of results, but an attempt will still be made in the interest of addressing PhoDyMM's success or lack thereof when analyzing this system. What makes this system particularly challenging is that individual transits have low signal-to-noise and there are considerable TTVs displayed in the data. By making use of their RIVERS method, which is trained to fit planet parameters in these types of systems, Leleu et al. derive masses of $M_b = 4.47_{-0.43}^{+0.48} M_{\oplus}$ and $M_c = 5.42_{-0.57}^{+0.61} M_{\oplus}$. Our analysis with PhoDyMM yielded masses of $M_b = 0.11_{-0.061}^{+0.091} M_{\oplus}$, $M_c = 0.69_{-0.242}^{+0.432} M_{\oplus}$, and $M_d = 20.96_{-5.99}^{+5.53} M_{\oplus}$ where planet d is the 3.38 day planet. As expected, these results do not match and will likely require additional analysis in order to obtain more correct values.

2.4.3 Systems with TTVs

While TTV fitting is a useful tool for analyzing the subset of Kepler multis that display TTVs, it was not something that was directly accounted for in our process of modeling the multis using PhoDyMM. As a result of this, it is not guaranteed that our best solution for a given system is also going to match the measured TTVs. However, given that PhoDyMM fits the lightcurves directly using an n-body integrator, it is expected that a good fit to the data will also correctly model TTVs when they are present within a system. It is thus interesting and meaningful to present a couple of systems where there are particularly interesting TTVs which are modeled well by PhoDyMM. To answer the question of whether a system has “interesting” TTVs, we use the TTV score given in Kane et al. (2019) which gives KOIs a TTV score from 1-9 based off several factors including the TTV signal strength and cleanliness. We present results for KOI-168 and KOI-248 which both received 9’s for their overall TTV interestingness.

KOI-168/Kepler-23

KOI-168/Kepler-23 is a three planet system where Kepler-23b,c and d have periods of 7.11, 10.74, and 15.27 days respectively. Figure 2.10 shows our model against the lightcurve data and it can be seen that there is good agreement between the binned Kepler data points in black and the binned model points in red, while there is some scatter in the black points, the residuals (shown below the lightcurves) show no clear structure and are small which further indicates a quality model fit to the data. As mentioned above, this is what PhoDyMM is directly working on accomplishing through its analysis tools. It is then encouraging to look at the TTV plots for this system shown in Figure 2.11 which shows our model TTVs in blue and the published TTVs in red along with the errors for these points for planets b and c (which are the only planets with precise TTVs in the system). Our model points accurately match the phase, amplitude, and period of the TTVs with almost every model point well within the error bars for the data values for both planets. Given that when TTVs are




Figure 2.10 Phasefolded lightcurves for the three planet system KOI-168/Kepler-23 with residuals between the model and data shown below. As with other phasefolded plots, the gray points are unbinned Kepler photometric measurements which are then binned together to get the black points and the magenta and red points are the unbinned and binned PhoDyMM model points. It is clear that our model provides a good fit to the data as we accurately match the depth and shape of each planet’s transits. Where there is a slightly larger scatter in the base of the transit, our model tends to hit the middle of the distribution which in turn implies the model is neither over nor underestimating planetary parameters. Given this systems has TTVs, the fact that these transits are not blurred also confirms that are model and n-body integrator are accurately representing this system.

present it is possible to constrain physical parameters, we find a masses of $M_b = 2.2527^{+1.0277}_{-1.1508} M_E$ and $M_c = 8.2690^{+4.2306}_{-2.9647} M_E$. While not the tightest constraints, it is worth noting that these are not exceptionally large planets ($R = 1.6329^{+0.0519}_{-0.0845} R_E$ and $R = 3.0302^{+0.0696}_{-0.0649} R_E$) and the masses are still constrained at about the 2σ level.

KOI-248/Kepler-49

KOI-248/Kepler-49 is a four planet system with planets d,b,c, and e having periods of 2.57, 7.20, 10.91, and 18.59 days. Our phase folded lightcurves are shown in Figure 2.12. Planet e has some slight scatter in the non-transiting region of the data, but otherwise there is good agreement between




Figure 2.11 TTV plot for Kepler-23b and c which have periods of 7.11 and 10.74 days. Only plots for these two planets are shown as the outermost planet in the system did not have particularly interesting TTVs as determined by Kane et al. (2019). What is seen here is that despite not modeling TTVs directly, PhoDyMM’s model for Kepler-23, shown in blue, predicts TTVs that are a near perfect match to the measured TTVs published by Rowe & Thompson (2015), shown in red with errors. For the case of both planets, our model TTVs match the amplitude, period, and phase of Rowe & Thompson’s TTV values and in many cases our points fall almost exactly at the center of the error range for each point.

the Kepler data and the lightcurve models. Looking at Figure 2.13 we show the model versus data for the TTVs of planets b and c and there is a high degree of agreement between the two as seen by the matching phase, amplitude and periods for both planets. We are able to place significant constraints on the masses for both planets with planet b having a mass of $M_b = 8.0607^{+1.5038}_{-1.5270} M_E$ and planet c having a mass of $M_c = 6.5733^{+1.3422}_{-1.3315} M_E$ combining these with their well-known radii of $R = 2.5396^{+0.0765}_{-0.0579} R_E$ and $R = 2.3194^{+0.0703}_{-0.0567} R_E$ we also obtained meaningful constraints on their densities. Our calculated densities are $\rho = 2.6893^{+0.5587}_{-0.5545} g/cm^3$ and $\rho = 2.8856^{+0.6514}_{-0.6327} g/cm^3$ respectively.

2.4.4 Hidden Planet Systems

Of the 34 systems that had a hidden planet added to them, 12 systems were flagged as failed both through automatic and visual inspection, and it is expected that this high failure rate is linked

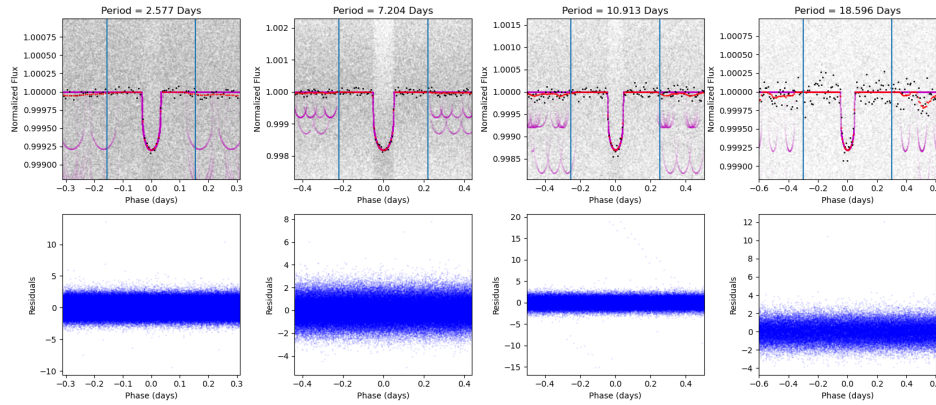


Figure 2.12 Phasefolded lightcurves and residuals between our model and the Kepler data for the four planets in KOI-248/Kepler-49. Inspection of the phasefolded lightcurves shows a high level of agreement between the data and our model as our model lightcurves accurately match the depth and shape of the transits for each planet. The 10.91 day planet Kepler-49c shows some outliers likely due to poor detrending of a single transit.




Figure 2.13 TTV plots for Kepler-49b and c. Out of the four planets these were the only two that had particularly interesting TTVs and are therefore the only ones presented here. Again it is seen that PhoDyMM’s TTV model in blue provides a good match to the observed TTV data in red. It is possible that our model is slightly underestimating the amplitude of the TTVs as it is seen that at the times of min and max transit time variation, our model tends to be on the upper and lower end of the error range respectively. Still given the fact that all of the measured TTV points have considerable errors, the model obtained using PhoDyMM is more than satisfactory especially when factoring in that we obtain statistically significant mass constraints of $M_b = 8.0607^{+1.5038}_{-1.5270} M_E$ and $M_c = 6.5733^{+1.3422}_{-1.3315} M_E$ for these planets.

to the increased uncertainty in planetary parameters that exists when a planet's presence is not being measured directly through transits. There were 13 systems that upon visual inspection, were determined to meet the criteria for convergence. Then lastly, 4 systems were classified as converged after passing our automatic checks. Given both our convergence metrics, and the poorly constrained parameters for the hidden planets, visual inspection was extremely important when determining whether the system had converged or not. This was the case because without a clear transit signal in the data, many parameters such as period, T_0 , inclination, and radius could float to arbitrary values without having any impact on the model fit to the data. This in turn would result in the system being flagged as failing, since we disallowed planets becoming non-transiting, having too large a radius, and significant changes in transit times. Ultimately, when inspecting one of these systems, if the model fits and posteriors for the known planets met our requirements for convergence, the system was passed even if this was not the case for the hidden planet. It is still worth reiteration, that there were the 4 systems that passed our automatic convergence metrics, which indicates that in these cases there was enough of a signal for PhoDyMM to lock onto. The list of systems that had a hidden planet added and the final status of the system is given in Table 2.6.

KOI-757/Kepler-229

KOI-757/Kepler-229 is nominally a three planet system with planets located at 6.25, 16.07, and 41.2 days. For this system, we add a planet at 64.45 days to better fit observed TTVs. Our results are given in Figures 2.14 and 2.15. Beginning with the phase folded lightcurve plot, we show the raw lightcurve points in grey, binned points in black, and our model points in red. For the first three planets, which are the known planets in the system, it is clear that we obtain a good fit to the data, which is further emphasized by the low residuals in the lower line of subplots. The fit for the outer planet is not remarkable, but this is also the injected hidden planet that currently has not been verified through transits. Only the 41.2 day planet, Kepler-229d, has observed TTVs greater

System Status	KOI Numbers
Converged	457, 720, 2199, 3184
Converged (Visual)	448, 456, 481, 598, 734, 757, 759, 1307, 1573, 1581, 1613, 2073, 3057
Failed	564, 638, 872, 1261, 1601, 1751, 1781, 2061, 2650, 2714, 3068, 3374

Table 2.6 Final statuses for the 34 systems that had a hidden planet added to them. Each system either converged automatically, was determined to be converged after a visual inspection, or was found to be in a failed state based off our defined failure criteria. In general, it was found that the additional planet, which did not necessarily have a transit signal, increased both the difficulty of obtaining a good model and the likelihood of failure. In all cases, the values of T_0 , inclination, and radius for the hidden planets were allowed to float beyond the limits set for known planets as these parameters were not well defined for the hidden planets. Still it was promising that through automatic and visual checking, 22 of these systems converged, making the efforts to improve the models for these system not wasted.

than 1-2 minutes and is thus the only one for which we display results for. Figure 2.15a shows our model TTV points in blue versus calculated TTVs from DR25 in red along with the errors in the calculations. With the exception of the first data point, our model matches the data well with almost every model point at least being within the errors of the calculated TTV values. This is contrasted with Figure 2.15b which shows our model TTVs when the hidden planet is not included. As can be seen without the hidden planet, the planet d is found to have essentially no TTVs and thus does not match the TTV data at all.

2.5 Conclusion

The Kepler database is currently the largest homogeneously constructed catalog of confirmed exoplanets and candidates. This also makes it the ideal dataset to perform population level statistics and analyses such as addressing the question of what the true distribution of a given parameter under a set of constraints or specific conditions. However, in order to perform such a work, there would

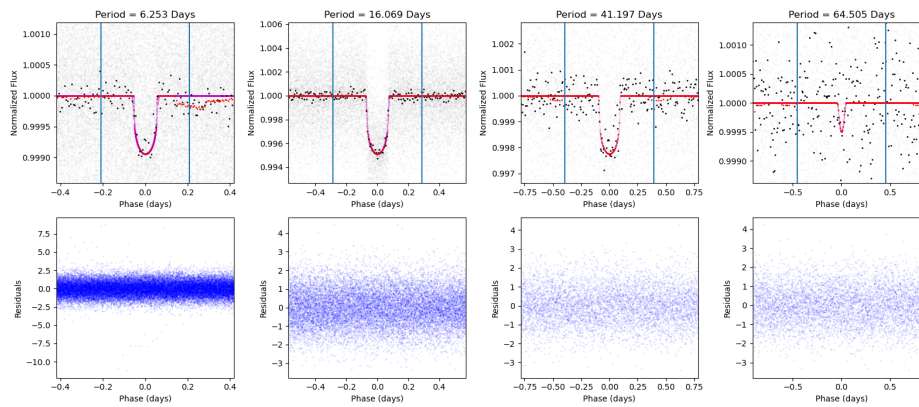


Figure 2.14 Phasefolded plot for KOI-757 along with the residuals between our PhoDyMM model and the Kepler data. For this system, we injected an additional planet with a period of 64.5 days as suggested by the results of (Graham & Ragozzine in prep.). Looking at the model fits, we find a good level of agreement between our model and the data for the three inner planets which is to be expected since these are the three known planets in the system. The phasefolded lightcurve for the hidden planet does not show much as PhoDyMM was trying to find a signal of a planet that is not necessarily there. However, it is worth noting that PhoDyMM did lock onto some sort of signal in the data as the parameter values for the hidden planet did not float to as extreme values as has been found in several of the other systems with hidden planets added to them.




Figure 2.15 TTV plots for KOI-757.02 with and without the hidden planet at 64.5 days. The goal of adding hidden planets to a select group of systems was to better model observed TTVs. In the left panel of this figure, we show our model TTVs overplotted with the observed TTVs when the hidden planet is added. It can be seen that our model provides a quality fit to the data as we accurately model the shape and amplitude of the TTVs with only ~ 5 points lying outside of the error range. This is contrasted with the right panel which shows our predicted TTVs in the case that the hidden planet is not included. Here it is found that KOI-757.02 has no TTVs and is in considerable disagreement with the TTVs that have been calculated previously for this planet. Therefore, it is clear that at least in the case of this system, adding the hidden planet provides a significant improvement in our ability to model its TTVs.

need to be a catalog of physical and orbital parameters for all the planets that was also derived in a homogeneous manner so as to avoid biases introduced by differing analysis methods. Since such a catalog did not exist, the goal of this work was to construct one.

Using our photodynamical modeling tool PhoDyMM, we modeled 719 Kepler systems that have multiple transiting planets. We made use of a photodynamical model since it has been shown to be the state of the art in modeling arbitrary exoplanetary systems, it can make use of both the long and short cadence Kepler data, and does not require individual transit times to be calculated. Our efforts were focused on just the systems with multiple transiting planets as these are the most information rich as they often contain additional dynamical information via planet-planet interactions which in turn allow physical and orbital parameters to be constrained and planetary membership confirmed.

Of the original 719 systems, we obtained converged posterior distributions for the physical and orbital parameters for 661 of them. From these results, we were able to construct the largest and most complete mass-radius distribution to date. Given that in order to constrain masses additional planet interactions such as TTVs must be present, our mass-radius plot shows considerable uncertainties for individual planets which is evident in the bands of points which represent individual draws from the mass posteriors. Still, there were many systems which did contain TTVs and thus allowed for considerable constraints to be placed on individual masses. We found that there were 104 planets where we detected masses above the 4σ level. While not all new mass measurements, this does further reinforce the fact that when the information is available in the Kepler lightcurve, PhoDyMM is able to constrain parameters to meaningful levels.

After broadly comparing our results to other studies of similar scope, as well as on a system by system level of a handful of particularly interesting and well-known systems, we found a significant level of agreement to what has previously been reported in terms of masses, inclinations, and eccentricities. There were both cases where PhoDyMM performed better and worse than previous studies, but at least some of regions of PhoDyMM's larger uncertainties can be attributed

to comparing a photodynamical model to RV measurements which can, in a few cases, measure masses more precisely. In addition to differences in techniques, it must be stated again that in the effort to create the most homogeneously derived catalog, almost no work was done to fine tune individual systems. If more fine tuning had been undertaken, it is possible that either additional systems could have reached convergence or those with some constraints could have had tighter ones placed. However, given the scope and time that this work already required, it would have been unfeasible to also adjust system individually. Fortunately, now that this initial catalog is completed, it would be possible in the future to look at individual systems that are believed to be able to have their parameters constrained further if some additional care is taken with their initial parameters or with how PhoDyMM handles movement throughout parameter space.

While it was not the goal of this work to find the individual best starting parameters for each system, there were several systems where it has been shown that the known planets do not allow for a model that accurately predicts the observed TTVs for that system. To give PhoDyMM the best chance possible at modeling these systems, we added an additional “hidden” planets to 34 of these systems and allowed PhoDyMM to work out the planetary parameters that best fit the data. Adding planets that have not been detected and thus have barely any signal in the data, increased the level of difficulty in modeling the system and was likely a significant factor in 12 of these systems outright failing our analysis. For the remaining 22 systems with hidden planets, we found results sufficient for meeting our convergence criteria even if in many cases the hidden planet was significantly unconstrained. As shown in Figure 2.15, adding the hidden planet was beneficial as it allowed observed TTVs to be modeled almost exactly whereas without the hidden planet there are no predicted TTVs. Therefore, we conclude that this extra endeavor was worth the effort.

We present the Kepler Dynamical Catalog for all to use and expect that it should be of considerable help in modeling both systems on an individual level as we now have complete sets of all

the physical and orbital parameters for all of the confirmed planets in each system as well on the population level of the Kepler database.

2.6 Future Work

With the completion of the Kepler Multis Dynamical Catalog, our immediate attention now turns to using the converged posterior distributions to deepen our understanding of many questions related to the dynamics, composition, and architectures of exoplanetary systems. Leveraging the known biases of the Kepler planets and our near-homogeneous analysis allows these questions to be directly connected to the true underlying distributions for a variety of parameters from mass-radius relations to eccentricities and from the true frequency of resonances to comparative planetology. In particular, we note that Hierarchical Bayesian Modeling could alleviate our use of wide priors and allow for improved inference of physical and orbital parameters by exploring the assumption (quantitatively in the hypermodel) that most Kepler planets are similar in nature.

In addition to this, we are planning on using our planetary parameters to enhance previous works that sought to fit interior composition models to the Kepler systems. By making use of our catalog which has posterior distributions for all of the physical and orbital parameters in the 661 converged systems, we expect to be able to further previous studies in two important ways. The first is through completion, earlier work was limited to the handful of systems that had mass estimates for the planets. Our catalog has mass estimates for all of the planets, even if the masses are not well constrained, which will instantly allow us to perform this work on a much larger scale than has previously been done. Second, we can use draws from our posteriors in order to model the interiors in a Bayesian framework. Beyond just simply modeling the interior compositions for these planets, we also plan to look at the composition trends that exist within the system such as gas mass fractions as a function of semi-major axis.

Finally, it is our goal that we can eventually revisit the 58 system that failed to converge throughout this work. It is expected that with a modest amount of fine tuning, it will be possible to obtain satisfactory results for these systems which in turn will allow us to add them to our dynamical catalog.

2.7 Acknowledgments

This material is based upon work supported by the National Science Foundation under Award No. 2143195. In addition, we would like to recognize the contributions of Burke Boyer and Steven Blodgett who were undergraduate researchers at the time this work was performed and assisted in generating some of the data processing and analysis methods. Finally, we acknowledge the Brigham Young University Office of Research Computing, without the help of the computing staff and the supercomputing resources that are provided through this department, this work would have never been completed.

Chapter 3

Enforcing Stability

Exoplanetary systems discovered through the transit method alone have well known planetary radii and periods, but many of the additional parameters such as masses, eccentricities, and mutual inclinations are poorly known. We employ a suite of recently published stability criteria in order to test whether dynamical stability can provide meaningful constraints on these parameters by requiring systems to be long term stable. We found that these metrics tended to agree well for systems that were far away from the stable/unstable boundary, but discrepancies occurred as the boundary was approached. This disagreement illustrated that the ability to constrain parameters depends on the particular application of a stability metric. Using the requirement that planets must be separated by at least 8 mutual Hill radii to be long term stable, we found that approximately 25 planets without strong TTVs could have their densities constrained to be less than 10 g cm^{-3} . For an additional 10 planets, stability constraints improved mass measurements even after an initial TTV analysis was performed.

3.1 Introduction

The success of the Kepler Space Telescope (hereafter “Kepler”) has made the transiting planet method the most successful for identifying candidate exoplanetary systems to date. In its four year Prime mission, Kepler identified several thousand candidates with ~ 2000 being validated as true planets using the transit method (e.g., Lissauer et al. 2014b; Rowe et al. 2014). However, planetary identification and validation is only the first of many steps towards understanding the characteristics of these systems. The next natural step is to determine the physical and orbital parameters for each planet in order to begin to address questions such as planetary composition, formation, evolution, and habitability. While exoplanetary systems discovered through the transit method alone have well known planetary radii and periods, many of the additional parameters such as masses, eccentricities, and mutual inclinations are poorly known. These unknown parameters are easiest to determine in systems with multiple planets which have the additional benefit of teaching us valuable insights about planetary system architecture.

Most exoplanetary systems are systems of multiple transiting planets discovered by Kepler that are too faint to reasonably follow up with additional techniques. This will remain the case for the foreseeable future; TESS discoveries can be followed up well (though mostly heterogeneously), but they will be quantitatively fewer than Kepler. Fortunately, Kepler observed most systems for four years continuously, allowing in $\sim 10\%$ of cases for dynamical interactions to be detectable, typically as Transit Timing Variations (TTVs) (e.g. Agol et al. 2005; Holczer et al. 2016; Holman & Murray 2005; Judkovsky et al. 2022; Kane et al. 2019). TTVs can effectively constrain masses and eccentricities – though often with degeneracies (e.g. Hadden & Lithwick 2017; Jontof-Hutter et al. 2021; Lithwick et al. 2012). True three-dimensional mutual inclinations for individual systems are much more challenging to constrain, though they can be revealed through Transit Duration Variations (e.g., Millholland & Winn 2021; Miralda-Escudé 2002; Payne et al. 2010; Shahaf et al.

2021). Measuring masses, eccentricities, and mutual inclinations for Kepler planets continues, especially for smaller planets and/or weaker signals.

The homogeneous nature of Kepler discoveries combined with the deeply studied completeness and reliability rates (e.g., Thompson et al. 2017) also mean that they provide the best hope for understanding the true underlying mass-radius distribution of exoplanets. As already shown in early studies of Neil & Rogers (2020) and Schlaufman & Halpern (2021), this has major implications for understanding the underlying composition distribution of exoplanets. Density constraints on a large ensemble of systems are a powerful way to address questions about planet formation processes such as *in situ* formation vs. migration from beyond the ice line or the evolution of exoplanetary atmospheres. It is also crucial to better determination of exoplanet habitability, such as the frequency of “born rocky” planets in the habitable zone of Sun-like stars (Neil & Rogers 2020).

In particular, methods that can *homogeneously* constrain the physical and orbital parameters of exoplanets have the potential to improve the information content that goes into analyses of the true underlying distribution of exoplanetary properties. One such technique that has not yet been used on the entire Kepler population is to require planetary parameters that are dynamically stable. As explained in Tamayo et al. (2020), due to the tight spacing of Kepler systems, a small variation in planetary parameters can sometimes result in instabilities after a time as short as 0.1% of the system’s age. Since it is highly improbable to discover a system on the verge of instability, planetary and orbital parameters that would create a system that is unstable on such short timescales can be rejected.

To address the question of whether or not a system is stable on Myr-Gyr timescales may require the use of computationally expensive N-body integration. A more tractable approach is to use analytic stability metrics, several of which have been developed with the goal of immediately estimating whether a system is stable. Many of these metrics are aimed at determining regions of parameter space where mean-motion resonances (MMRs) overlap which, as shown in Wisdom

(1980), causes planetary orbits to be chaotic. It has been found by Hadden & Lithwick (2018) and others that shortly after MMR overlap occurs, planetary systems tend to destabilize. Thus these metrics are able to estimate parameter values that will result in long-term stability.

There are many dynamical stability metrics that have been proposed over the years with varying levels of accuracy as tested on Kepler-like planetary systems with a wide range of varying initial parameters. Investigating many of these metrics is the first goal of this work. We performed numerical simulations on all of the Kepler systems with 3+ planets. For each realization of a system, a suite of stability metrics were calculated in order to allow direct comparison between these metrics on the real Kepler data. Of particular interest are the regions of parameter space where metrics do not agree as well as the transitions between stability and instability for the various metrics.

Our second goal is to better understand how these metrics can be used to constrain physical and orbital parameters by downweighting or eliminating unstable values. Although this was already a consideration in the first dynamical analysis of Kepler multis (Lissauer et al. 2011), the first extensive study was done by Deck et al. (2013). Recently, Tamayo et al. (2021a) performed a case study of the 3 planet system Kepler-23. Using their Stability of Planetary Orbital Configurations Klassifier (SPOCK), they find that stability can place upper limits on masses and eccentricities which are comparable to or tighter than other current state of the art methods. Building on these studies, this work will continue looking at the Kepler systems with the goal of identifying systems where densities, eccentricities, and/or mutual inclinations can be constrained to physically significant values such as those that would rule out particular compositions.

The remainder of this paper will be organized in the following manner. In Section 3.2 the stability metrics used for the subsequent analysis are summarized. Following this, the methods for initializing the system simulations are given in Section 3.3. The simulation results and corresponding analysis are given in Sections 3.4 and 3.5. Finally, conclusions about this work and discussion on initial questions about constraining physical and orbital parameters are presented in Section 3.6.

3.2 Stability Metrics

3.2.1 What is the Goal of a Stability Metric?

In practice, dynamical stability is not a well-defined binary determinant, but a probability distribution related to how long it will take for a system to reach a configuration that is somehow considered to be implausible. For example, Hill stability is defined as a configuration where two planets are disallowed to ever have close encounters by conservation of energy and angular momentum, though the orbits may be extremely chaotic or one planet may escape to infinity (e.g., Marchal & Bozis 1982). Generally, systems that are not Hill stable are prone to collisions on short timescales, thus it seems implausible that a system will be in a configuration that is not Hill stable. However, systems can be Hill stable while still evolving to unexpected configurations like an arbitrarily high eccentricity, leading to the collision of a planet and star. It is also important to point out that chaos and stability are related, but not identical concepts. Generally speaking, instability originates when there is dynamical chaos (e.g., from MMR overlap), but chaos itself is not a complete predictor of stability as in so-called “stable chaos.”

After an instability, a new planetary system arises (with one fewer planet), which itself is likely to be stable for a longer time than the previous system. As a result, a heuristic for the stability timescale of a given system is just the system’s age: a 10 Myr system is expected to have a stability timescale of approximately 10 Myr while a typical Kepler system is from a field star and has an approximate stability timescale of a few GYr.

How can we take such a fuzzy concept like “stability” and use it for practical exoplanetary analyses? We can make this idea more precise as we apply it to exoplanetary systems by appealing to the Copernican principle: for a given exoplanetary system, it is unlikely that we are observing it at a special time.

Let us consider the case where a Bayesian dynamical analysis of an exoplanetary system returns a list of samples from the posterior probability distribution that includes all physical and orbital parameters for all the known planets. That is, we have a large list of exoplanetary system parameters that are consistent with the observed data. Since the observed data span an extremely short time relative to the age of the actual exoplanetary system, it is quite plausible to have a set of parameters that match the data, but would be unrealistic for a real exoplanetary system because it would go completely unstable on timescales between ~ 10 years and the $\sim 10^9$ year age of the system. Our goal will be to identify which posterior draws (equivalent sets of orbital and physical parameters) can be ruled out in this way, thus better constraining our knowledge of these parameters as in Deck et al. (2013).

Consider some combination X of the given parameters that produces a quantitative output, similar to a derived parameter. Suppose that X is calculated every year from the age of the youngest system (~ 30 Myr) to the age of the oldest system (~ 10 GYr) for a large ensemble of *actual* planetary systems. This would give the true distribution of X values that would be expected from a random sampling of systems and times, e.g., from Kepler observations. If a posterior sample had a value of X that is extremely rare, this would provide evidence that this sample should be downweighted (by the probability of observing a value of X that extreme).

Of course, in practice, we can't know the exact distribution of X for actual planetary systems sampled at all times. But for well-chosen values of X , we can estimate a probability distribution based on the physical laws of orbital dynamics. For example, let X be the minimum of the difference of the semi-major axes between any two planets (not in the 1:1 resonance) measured in units of mutual Hill radii (defined below). Dynamically speaking, there is significant evidence that in the ensemble of real planetary systems sampled regularly in time, $X < 1$ is extremely rare. Whenever X goes below 1, the system rapidly rearranges into a new configuration (via merger, ejection, or some other kind of destruction) where $X > 1$ (Gladman 1993; Hussain & Tamayo 2020; Obertas

et al. 2017). Hence the probability of observing $X < 1$ for an arbitrary system at a random time is very low, and any posteriors corresponding to that configuration can be significantly downweighted. Note how we need the definition for the metric (how to calculate X) *and* a threshold value that is considered improbable.

By thinking about the probability distribution of a given metric over all exoplanetary systems at all times, we can avoid issues of the precise definition of stability. We can also recognize that in our sample of 227 Kepler systems (with a typical stability timescales of approximately a GYr) that it is not entirely improbable that one system is within ~ 1 Myr of an instability. Finally, we see how a particular threshold is not a binary discriminator between stable and unstable systems, but chosen to have some very low probability of occurring in a planetary system.

In the following subsections, we list all the metrics we consider for our stability analysis. As previously mentioned, this work focuses specifically on Kepler systems with 3+ transiting planets. This is because there is a fundamental difference between 2 planet systems and those with a higher multiplicity when considering their orbital stability. As shown in Gladman (1993), a pair of planets with initially circular coplanar orbits will never develop crossing orbits if the orbital separation is greater than $2\sqrt{3}$ mutual Hill radii. Therefore when considering the stability of these systems, calculating the dynamical separation in units of mutual Hill radii is typically sufficient to determine whether the system is long-term stable.

It is also necessary to note that none of these metrics are designed for an arbitrary number of planets. Most metrics only consider (adjacent) pairs or triplets of planets. In order to generalize these metrics to higher multiplicities, calculations were performed for each adjacent pair or triple within a system. The rationale for this is based on the dynamical heuristic that, except for unusual cases, adding additional planets to a system will only lower the probability of it being long-term stable. Therefore, by only considering pairs or triples of planets, the metric values will simply represent an upper limit on the systems stability and will not result in an unstable system being

erroneously classified as stable. This is also the recommended methodology for using these metrics on higher multiplicity systems (e.g. Tamayo et al. 2021b).

3.2.2 Hill Stability

Following the reasoning of Gladman, the first stability metric calculated is the planetary separation in units of mutual Hill radii R_{mH} which is defined as

$$R_{mH} \equiv \left(\frac{\mu_i + \mu_{i+1}}{3}\right)^{1/3} \left(\frac{a_i + a_{i+1}}{2}\right) \quad (3.1)$$

with μ being the planet-to-star ratio. The separation in units of R_{mH} is then given by

$$\Delta_{circ} = \frac{a_{i+1} - a_i}{R_{mH}} \quad (3.2)$$

Gladman determined that a pair of adjacent planets on circular orbits is Hill stable if $\Delta_{circ} > 2\sqrt{3} \approx 3.46$. However, we are also interested in eccentric orbits. It is nominally plausible to examine the energy and angular momentum constraints used to derive Hill stability (from, e.g., Marchal & Bozis (1982)) in any configuration. Here, we use a commonly-used approximation for circular orbits which measures Hill separation when the inner planet is at apoapse and the outer planet is at periapse. For this metric, the separation in units of R_{mH} is

$$\Delta_{ecc} = \frac{a_{i+1}(1 - e_{i+1}) - a_i(1 + e_i)}{R_{mH}}. \quad (3.3)$$

3.2.3 SPOCK

Recently, Tamayo et al. (2020) have developed the Stability of Planetary Orbital Configurations Klassifier (SPOCK) which is a machine learning model capable of estimating the long-term stability of compact multiplanet architectures based on short integrations. Given that SPOCK is the most

state-of-the-art stability metric implemented, its output is anticipated to provide the best indication of the long-term stability for a particular system.

SPOCK Stability Probability

In Tamayo et al., they explain that SPOCK begins by integrating a system for 10^4 orbits of the innermost planet. Following this integration a set of 10 parameters is calculated, which the machine learning model then uses to extrapolate the stability out to 10^9 orbits and return the overall probability of stability. Tamayo et al. found that SPOCK correctly predicted 94% of systems in their test set and since SPOCK only integrates 10^4 orbits, it is 10^5 times faster than direct N-body integration.

Deep Regressor

Another version of SPOCK is the Deep Regressor Model. As explained in Cranmer et al. (2021), the deep regressor is a Bayesian neural network model that is capable of determining the time before a system becomes unstable (as opposed to the stability probability after 10^9 innermost orbits) and provides confidence intervals on its predictions.

However, we found that while a single evaluation for the stability probability required ~ 1 CPU second, a single evaluation using the Deep Regressor Model required ~ 30 CPU seconds. Therefore, this model was only implemented for a smaller set of simulations which will be explained in further detail in Section 3.3.

Analytical MMR Overlap

Additionally, we implement the semi-analytical criterion for the presence of chaos given in Tamayo et al. (2021b). As explained in their paper, this metric determines the region of chaotic parameter space based on the overlap of MMRs including secular evolution to maximize eccentricities. We access this metric through the SPOCK interface which returns a probability that the specified system

is chaotic or normal. Although chaos and stability are not equivalent, this is still used as a point of comparison.

3.2.4 Petrovich

In Petrovich (2015), they present a stability criterion which was obtained by running a large number of long-term numerical integrations and the use of a Support Vector Machine algorithm to search for an empirical boundary that best separates stable systems from unstable ones. The condition that they find for stability is that

$$a_{out}(1 - e_{out})/[a_{in}(1 + e_{in})] > 2.4[\max(\mu_{in}, \mu_{out})]^{1/3}(a_{out}/a_{in})^{1/2} + 1.15 \quad (3.4)$$

with μ being the planet-to-star ratio and the subscripts in or out referring to the inner or outer planet of the planetary pair being considered.

For the purposes of this work and in the interest of being able to compare this criteria to other stability metrics, we slightly modify the form of Eq. 3.4 and define our Petrovitch metric as

$$P_{metric} = \frac{a_{out}(1 - e_{out})/[a_{in}(1 + e_{in})]}{2.4[\max(\mu_{in}, \mu_{out})]^{1/3}(a_{out}/a_{in})^{1/2} + 1.15}. \quad (3.5)$$

With the metric defined this way, the Petrovich stability criterion occurs at $P_{metric} = 1$ with a pair of planets is nominally stable if $P_{metric} > 1$ and unstable if $P_{metric} < 1$.

3.2.5 Hadden and Lithwick

Finally, we also implement the stability criterion given in Hadden & Lithwick. This metric specifically identifies the critical eccentricity above which chaos is triggered. The work done by Hadden & Lithwick builds on the pioneering work of Wisdom where it is shown that chaos occurs where mean motion resonances overlap. Hadden & Lithwick improve the metric originally found

by Wisdom (1980) by extending it to eccentric planets and examining resonances of all orders. It is extended by the analytical SPOCK metric discussed above. The formula given for the critical eccentricity e_{crit} is

$$e_{crit} \approx \frac{e_{cross}}{\sqrt{2}} \exp[-2.2(\mu_1 + \mu_2)^{1/3} \left(\frac{a_2}{a_2 - a_1}\right)^{4/3}] \quad (3.6)$$

with e_{cross} defined as

$$e_{cross} = \frac{a_2 - a_1}{a_1} \quad (3.7)$$

and μ again being the planet-to-star ratio. To turn this criterion into a value that could be compared to our other metrics, we took the ratio of an individual planets eccentricity to the critical eccentricity of the planet pair being considered. Therefore, our Hadden and Lithwick metric becomes

$$HL_{metric} = \frac{e_i}{e_{crit}} \quad (3.8)$$

where e_i is the higher eccentricity of the two planets. For this metric a pair of planets is not chaotic if $HL_{metric} < 1$ and is chaotic if $HL_{metric} > 1$.

3.2.6 Additional Metrics

In addition to the aforementioned metrics, we also implemented the stability criterion from Wisdom as well as the planetary system survival time given in Petit et al. (2020). However, in our main results and discussion, we will not be using these values.

Our reasoning for not including analysis of the metric from Wisdom is twofold. First, the metric was designed for circular orbits specifically and the majority of our simulations begin with planets on orbits with non-zero eccentricity. Second, the Hadden and Lithwick metric is essentially a

generalization of the Wisdom metric and considers additional MMRs which therefore makes it a preferred metric over the one given in Wisdom.

As for the survival time from Petit et al., it also has the condition that orbits are assumed to be circular. Furthermore, the goal of this metric was to understand long-term weak instability invoked by three-body resonances in configurations that otherwise appeared stable. As a result, the survival time is also only physically significant for extremely compact systems. This is most evident by inspection of Fig. 9 of Petit et al. where for planets with masses of $10^{-5}M_{\odot}$ the survival time tends to infinity once the period ratio between planets is greater than 1.2. For our dataset, it is uncommon for planets to be this tightly packed, and therefore the predicted survival time was not meaningful for this work.

3.3 Methods

To best accomplish the goals of this work – to understand how various stability metrics compare to one another and to what degree physical and orbital parameters can be constrained by enforcing stability – we initialized four runs of simulations for Kepler multi-transiting systems. By doing this, we were able to constrain certain parameters within a simulation run and thereby lower the dimensionality of the problem to better investigate how the remaining parameters affected the stability of the system.

3.3.1 Run Initialization

Each simulation run began with initial parameters for the 227 Kepler systems that were analyzed. These 227 systems represented all of the Kepler systems with 3+ transiting planets from the upcoming catalog of Rowe et al. (2014) after filtering out exoplanet candidates that are unlikely to be actual planets. Specifically, of the four disposition scores given in the catalog, which are

based on: Rowe et al., DR25supp, DR25, and DR24, we require that three of the four scores be 'P' indicating a planet/candidate. Additionally, two of the three P's had to come from the first two scores. There were four additional systems that had either 7 or 8 transiting planets that were not simulated for simplicity leaving just the 227 systems with 3-6 planets.

The starting parameters for each system consisted of the mass and radius of the host star, the radii of the known planets, and their periods. The specific values for these parameters were obtained using the exoplanetary data found in Rowe et al. which is similar to previously published results (e.g., Kepler Data Release 25, ThompsonetalRef), but with special care taken for multi-transiting systems. For each of the 10^4 realizations of each system, these initial parameters were held as constants. The remaining physical and orbital parameters for each realization were drawn from various distributions depending on the run and will be explained in the following subsections.

Unconstrained Parameters

Parameters for the first run were chosen from distributions that are believed to best model the actual parameter distributions for the Kepler systems and to allow for sufficient coverage of parameter space. Since we were primarily interested in the ability to constrain densities, a density value was drawn from a uniform distribution between 0 and 10 g/cm^3 . Then, using the predetermined radius, the mass of the planet was initialized. The eccentricities and inclinations were drawn from log-normal distributions where μ and σ depend on the system multiplicity with specific values given in He et al. (2020). These values are consistent with the underlying Kepler distribution, though it is important to note that this work already accounted somewhat for dynamical stability.

For the case of eccentricities, especially for low multiplicity systems, there was a very small chance that the draw from the log-normal distribution would return a value greater than 1. To forbid unbound orbits, we set a liberal upper limit of 0.9 on the eccentricity. The remaining orbital angles

were drawn from uniform $(0,2\pi)$ distributions; work by Lissauer & Gavino (2021) suggests that random angles do not significantly affect stability timescales.

e=i=0

The second simulation run was aimed at reducing the overall dimensionality of the stability problem as much as physically possible. For this run, masses and orbital angles were initialized in the same manner as in the previous run. However, eccentricities and inclinations were set to 0 thereby creating circular and coplanar systems. By initializing the systems in this way, the only remaining floating parameter that could affect the system's stability was the masses of the planets.

e=i=small

Realizing that circular, coplanar orbits are not realistic starting values for the Kepler planets, we then proceeded to simulate the systems by giving each planet a small eccentricity and inclination. Eccentricities were set to 0.01, inclinations were set to 1° , and all other parameters were initialized as before in earlier runs. This run therefore allowed for most of the benefits of the e=i=0 which are that eccentricity and inclination would not be the major causes of instability, but did not require a perfectly idealized system simulation.

m=i=small

In the previous runs, mass was the varied parameter. However, we also desired to have a simulation run where eccentricity was the varied parameter. To accomplish this, we began by setting the inclinations to 1° and the orbital angles were drawn from uniform $(0,2\pi)$ distributions, as before. To make stability rely as little on masses as possible, it was necessary to set masses to physically realistic values without simply setting them to 0 as this would make all of the realizations trivially stable. We chose to use the mass-radius relationship given in Lissauer et al. (2011) which is

$M/M_E = (R/R_E)^{2.06}$. As explained in Lissauer et al., this power-law was obtained by fitting to Earth and Saturn and is reasonable for planets within the Earth to sub-Neptune range. We therefore used the established radii for the planets and used the power-law to calculate the masses, which were then held fixed for all realizations of the system. We initialized the system with eccentricities drawn randomly from 0 to e_{cross} , the eccentricity that would lead to a crossing orbit assuming all other eccentricities are 0.

Mutual inclinations

To explore the effect that significant ($> 1^\circ$) mutual inclinations have on the stability of planet pairs, an additional set of runs was completed. As the effects of mutual inclination are not included in most stability metrics (mutual inclination is assumed to be small), a new metric was required. Our metric assumes that the angular momentum deficit (AMD, Laskar (1997)) contained within the inclination of a planet (measured from the Laplace plane of the planet pair) can be ‘redistributed’ into that planet’s eccentricity AMD. The new (higher) eccentricities of the planets’ are then input into the Hadden & Lithwick metric.

We evaluated all Kepler planet pairs using this metric, with inclinations drawn from a $\cos(\text{inclination})$ distribution with the masses set by the aforementioned mass-radius relationship, eccentricities fixed at 0.01 and additional angles drawn from uniform $(0, 2\pi)$ distributions. We limit our analysis to prograde orbits

Jontof-Hutter et al. (2021) posteriors

Finally, to better quantify how stability analysis can aid additional methods to constrain planetary parameters, we applied our suite of stability metrics to a set of posterior distributions from Jontof-Hutter et al. (2021). These posteriors were created by performing transit time variation (TTV) analysis on a set of 48 Kepler Systems with multiple transiting planets. Each system’s posterior file

contained a set of 10^4 posteriors whereas each posterior contained five parameters for each planet which are: $\frac{M_p}{M_*} \frac{M_\odot}{M_E}$, $P(\text{days})$, $e \sin \omega$, $e \cos \omega$ and T_0 . Since the posteriors did not contain all of the parameters needed to simulate the system, the remaining parameters were set in a similar manner to the other simulation runs. Specifically, the stellar parameters and the planetary radii were read from the data table given in Rowe et al. and its later updates, the inclinations were set to 1° , and the orbital angles were drawn from uniform $(0, 2\pi)$ distributions.

3.3.2 Metric Calculation and System Simulation

After physical and orbital parameters were determined for all the planets in a system, all the aforementioned metrics were calculated directly. The only exception to this was for a smaller run that made use of the Deep Regressor from the SPOCK package. Due to the additional computation time for the Deep Regressor, when it was implemented for the $i=e$ =small case, only 1000 realizations were initialized. Again as mentioned, many of these metrics were specifically designed for either pairs or triplets of planets. To account for this, systems were broken up into pairs or triplets and the suite of metrics were calculated on the planetary subsystem. For each realization, the returned metric value was the one corresponding to the most unstable pair or triplet of planets. This value then represented the upper limit on the stability for the system for the particular metric because considering the entire system as a whole would generally only lower the probability that the system would be stable. Tamayo et al. consider adjacent triplets using SPOCK and find that the lowest SPOCK probability is a valid estimate for the stability of the entire system.

All of the computations were executed on the supercomputer at Brigham Young University, with each system analyzed in parallel. For a standard run with 10^4 realizations, all of the 3 planet systems could be simulated in a few hours. However, the time for systems with higher multiplicities was essentially $t_n = (n - 2) * t_3$ due to the fact that each additional planet required an additional calculation for all the metrics. In its entirety this analysis used approximately 70 CPU-days.

3.3.3 N-body Integrations

It is understood that the stability metrics implemented for this analysis are an approximation of the effects of the true underlying dynamical interactions within each system. Given this, it was also informative to select a small subset of systems and perform the full n-body integration using the software package REBOUND. Out of the 227 Kepler systems initially simulated, seven were selected to be integrated for 10^6 years. These seven were selected after analyzing the results from the simulations and were specifically chosen due to interesting characteristics that were present when searching for density constraints based on the SPOCK stability values. To better aid in the comparison between the SPOCK and n-body integration results, two of these systems were rerun with an integration time of 10^9 inner orbits which matches the stability period that SPOCK predicts. Once the systems were selected, 1000 realizations of the $e=i$ =small simulations were integrated using initial parameters obtained from the already generated output files. This gave the added benefit that the n-body integrations were performed on the exact configurations of each system that were analyzed by the other metrics such as SPOCK.

3.4 Results

3.4.1 Metric Comparisons

In general, the stability metrics agreed well with one another. One example of this agreement is shown in Fig. 3.1 which shows the Petrovich versus the Hadden-Lithwick metric for KOI-0116 for our first run with unconstrained parameters. There is a nearly linear relationship between the two metrics. While it is unsurprising that the metrics agree with one another, it is unexpected that there would be such a distinct linear relationship given the major differences in the functional forms of the metrics. However, in addition to the linear trend, it can also be seen that the Hadden-Lithwick metric predicts a stable planetary configuration before the Petrovich metric. As mentioned, the

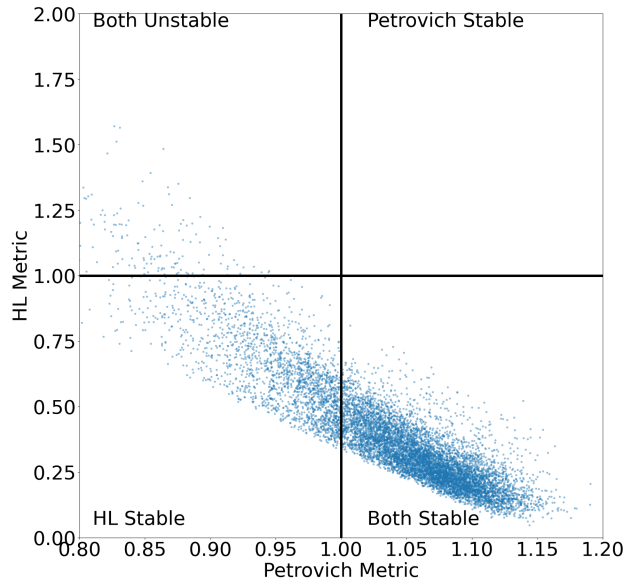


Figure 3.1 Hadden-Lithwick metric values plotted against the Petrovich metric for KOI-0116. The two metrics show a nearly linear correlation with each other. At the boundary of stability the level of agreement declines as seen by a stable Hadden-Lithwick value with $HL_{metric} < 1$, but an unstable Petrovich value with $P_{metric} \approx 0.8$.

Hadden-Lithwick metric is stable whenever the metric value is less than 1 and the Petrovich is stable whenever the metric value is greater than 1. As can be seen in the figure, many systems that are stable by Hadden-Lithwick still have a Petrovich value of $\approx 0.8 - 0.9$ indicating the system is not quite stable by our Petrovich metric. This result helps to show that the boundary between stable and unstable is challenging to define and within this region it is possible for metrics to disagree.

We can also compare the SPOCK values with the Eccentric Hill Metric and the Petrovich metric for 100 realizations of each system using data from our unconstrained parameters as shown in Figs. 3.2 and 3.3. The trendlines show the binned results based on the system's planetary multiplicity. Two trends are made clear. The first is that for a given Hill separation or Petrovich value, the corresponding SPOCK value tends to decrease with increasing planet multiplicity. This is not too surprising and merely serves to reinforce the understanding that additional planets do not increase the probability that a system is long-term stable. The second trend is how SPOCK values transition




Figure 3.2 Minimum planetary separation in Hill radii plotted with the corresponding SPOCK probability of stability for the system. For each system, 10^4 realizations were initialized with the metric values for the first 100 being shown in the figure. The data were then binned by planet multiplicity and by planetary separations and plotted trendlines show the median values for each bin. As can be seen, the probability of stability increases as planetary separations increase with a significant jump in stability as separations go from $\Delta_{ecc} \simeq 4$ to $\Delta_{ecc} \simeq 8$. It is also clear that as planetary multiplicity increases the SPOCK probability of stability decreases.

from generally unstable to generally stable. This is most evident for the SPOCK versus Petrovich data in Fig. 3.3 where the SPOCK values are essentially 0 until the Petrovich value is 0.9 and then the SPOCK values jump to between 0.8 and 1 by a Petrovich value of 1 (the stability threshold) indicating a very high probability of stability. While not as dramatic for the Eccentric Hill metric, it is clear that when the planetary spacing is less than ≈ 4 mutual Hill radii the systems have a very low probability of stability. The stability probability then increases quickly as the spacing increases and once the spacing is greater than ≈ 8 mutual Hill radii, most systems have a high SPOCK stability probability. The implications of this trend are further considered in Section 3.5.

It is also worth noting that SPOCK rarely assigns stability probabilities near 1, even for systems that appear quite stable from other metrics.

Fig. 3.4 shows the SPOCK Deep Regressor times versus the Eccentric Hill metric using data from our $e=i$ =small run.




Figure 3.3 SPOCK probability values plotted against the Petrovich metric values. Similar to Fig.3.2, 100 realizations are plotted with trendlines showing median values within each multiplicity and Petrovich value bin. A Petrovich value that is greater than 1 indicates a stable system and this agrees well with where the SPOCK values transition from low probabilities of stability to high probabilities. While for the Petrovich metric, the median values based on planetary multiplicity are less dispersed there is still evidence that stability decreases as the number of planets increases.

It is once again clear that systems with more planets tend to be less stable and therefore have a lower predicted time from the Deep Regressor metric. Interestingly, around Δ_{ecc} of 8, the stability timescales stabilize. The near-linear relationship between the logarithm of the stability timescale and the separation in mutual Hill radii matches previously known results (e.g. Obertas et al. 2017).

3.4.2 Unconstrained Parameters

Our analysis of the ability to constrain physical and orbital parameters began with looking at results for our run using unconstrained parameters. We generated scatter plots for planetary parameters versus the SPOCK metric in order to identify regions of instability within the parameter space. Once such plot is given in Fig. 3.5. Plotted is the density for the first planet versus the second planet for KOI-0041 with the coloring based on the SPOCK probability of stability. Upon inspection of this figure, it becomes clear that for the majority of the simulations KOI-0041 has a high probability

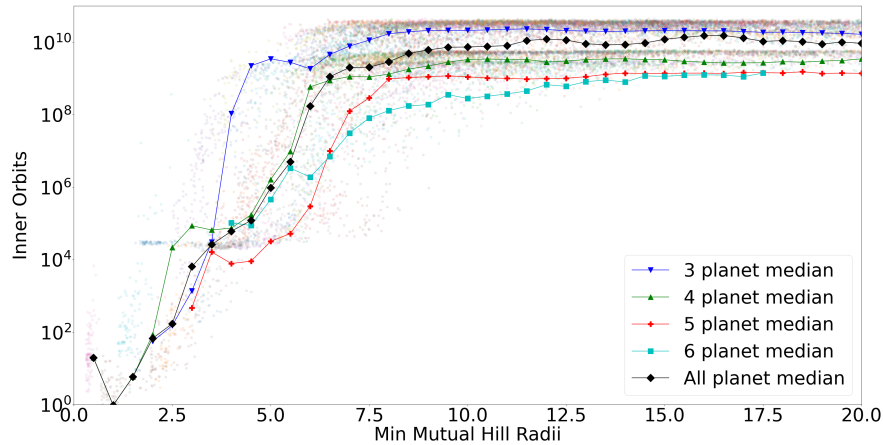


Figure 3.4 Stability time determined by SPOCK’s deep regressor model versus minimum planetary separations in units of Δ_{ecc} . Similar to other metric-metric plots, stability increases as separations increase and systems with fewer transiting planets are generally more stable as indicated by the increase in stability time.

of being long term stable. However, there still remain a number of realizations where the specific combination of parameters led to an overall configuration that SPOCK determined to have essentially no chance of being stable. What is also clear about these unstable configurations is that they appear to be randomly dispersed throughout the figure. This is likely a result of individual draws having particularly high eccentricities that were not correlated with densities in our setup. Thus, the high dimensionality of the parameter space makes it challenging to draw any specific meaningful results.

In the interest of being able to draw meaningful conclusions about the ability of constraining parameters using stability analysis, we moved on to more idealized system configurations, which reduced the number of floating parameters.

3.4.3 $e=i=0$

Our second run of simulations considered systems with planets initially on circular, coplanar orbits which thereby reduced the number of parameters by $2 * n_{Planets}$ and left masses as the primary floating parameter. Starting as we did for the previous run, we generated scatter plots to identify

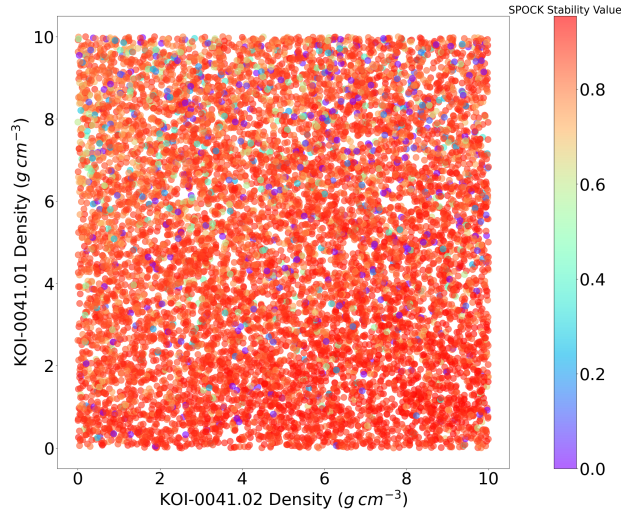


Figure 3.5 Scatter plot for the densities of the two innermost planets in KOI-0041. The color values are determined by the SPOCK probabilities with red indicating a higher probability of stability and blue indicating a lower probability. While it is clear that most realizations of this system have a high probability of stability, the unstable configurations appear to be randomly dispersed in such a way that there is no clear density cutoff to ensure long term stability.

particularly interesting systems where our stability metrics would allow the remaining parameters to be constrained. One such plot is given in Fig. 3.6 where the densities for the first and second planets are plotted against each other for KOI-0041 again with the colors determined by the SPOCK values but using a different range of colors. Unlike Fig. 3.5 this plot shows a clear region where the stability is slightly enhanced (SPOCK probability difference of 0.1) by having the innermost planet have a density below 5 g/cm^3 . This system would thus somewhat favor lower densities for both of these planets.

To further quantify the expected number of planets that could have their density constrained based on our stability analysis (assuming $e=i=0$), we looked at how SPOCK values varied simply as a function of density for each planet and plotted the results in Fig. 3.7. We emphasize a density limit of $\rho = 4000 \text{ kg/m}^3$ where it becomes possible to rule out certain compositions for typical Kepler planets. For the majority of planets, their density did not affect the SPOCK value for their




Figure 3.6 The same setup as Fig. 3.5 with a different colorbar range except data for this plot were generated from the run where eccentricities and inclinations were set to 0 which reduced the overall dimensionality of the problem. With this reduction, the range of SPOCK values is reduced such that the minimum stability value is 0.86 indicating that every iteration has a relatively high chance of being long term stable. Still, there is a clear increase in the stability probability when the density of the first planet is below $5g/cm^3$, for example. Therefore, in this system it is anticipated that density values could in fact be somewhat constrained by requiring the system to be long-term stable

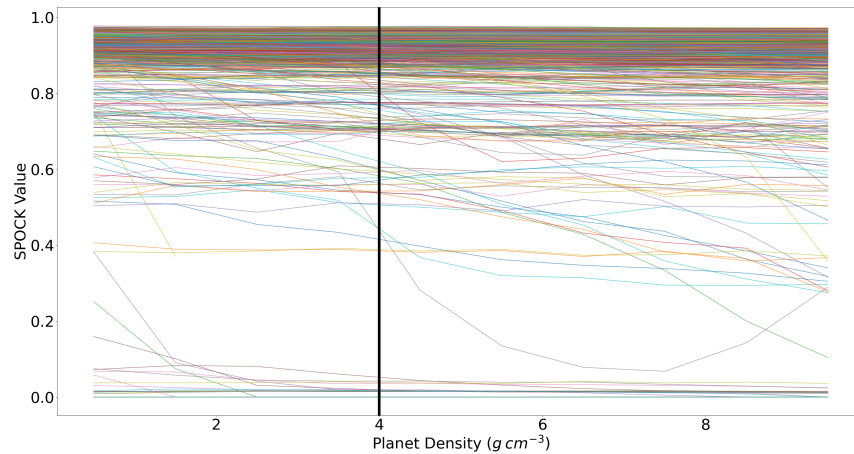


Figure 3.7 Median SPOCK values versus planet densities for the particular density bin in the case where eccentricities and inclinations are fixed to 0. If an individual planet’s density has no effect on the system’s stability, the result is a line at a constant SPOCK value. Some planets could have densities constrained by using SPOCK values as probability weights. We highlight $\rho = 4g/cm^3$ where many planets would have meaningful compositional constraints.

system and therefore their densities cannot be constrained using SPOCK. However, a fraction of planets have SPOCK values that are inversely correlated with density, allowing stability to constrain masses.

To assess the strength of SPOCK in constraining densities, we calculated the average value of SPOCK above and below $\rho = 4000kg/m^3$ which we call Δ SPOCK. Fig. 3.8 shows Δ SPOCK as a function of planetary radius. When SPOCK cannot place stability constraints, we expect Δ SPOCK to be 0. Most planets have no such constraints, but there are ≈ 20 planets that have a Δ SPOCK ≥ 0.1 .

One concern with the effectiveness of dynamical stability constraints is that systems that are dynamically unstable will also have significant TTVs, allowing for precise mass measurements independent of stability. To assess the validity of this concern, we use the TTV “interestingness” score from Kane et al. (2019) to color the points in Fig. 3.8: the strongest TTVs are plotted in blue, moderate TTVs in green, mild TTVs in red, and no TTVs in gray. This shows that only $\approx 5/20$

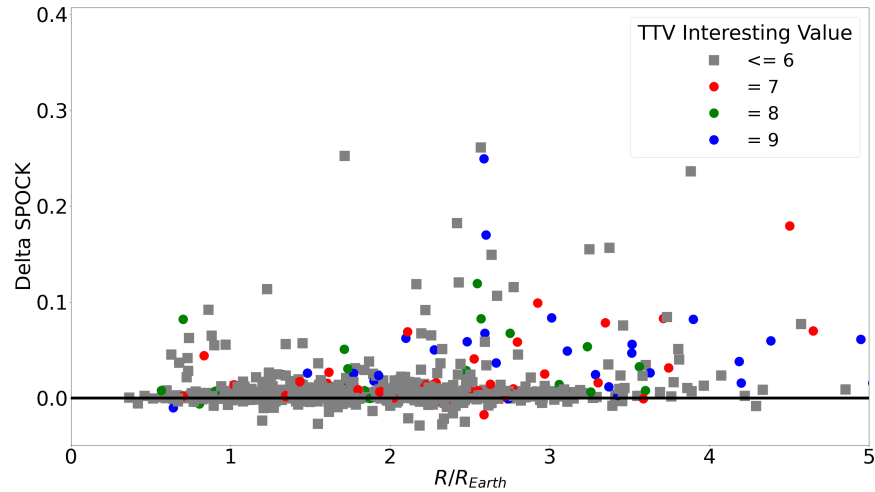


Figure 3.8 Δ SPOCK values versus planetary radii. The Δ SPOCK value is the difference in the average SPOCK value above and below a $4g/cm^3$ cutoff (see Figure 3.7). A positive Δ SPOCK value indicates that the system is more stable when the planet’s density is less than $4g/cm^3$. While there is a degree of scatter around zero, about 20 planets have a Δ SPOCK ≥ 1 . The plot points are coded by whether or not the planet was identified as having interesting TTVs by Kane et al.. Blue points indicate the strong TTVs, green points moderate TTVs, red points mild TTVs, and gray squares no TTVs. While there is some correlation between TTV strength and Δ SPOCK, there are several planets where stability constraints would provide new density information not available from TTVs.

systems with Δ SPOCK ≥ 1 have strong TTVs. This suggests that there are many planets where stability can provide novel mass constraints.

Supposing that our uniform distribution of mass from 0 to $10 g cm^{-3}$ is like a Bayesian prior and that SPOCK gives the likelihood (in the form of a probability), we can estimate the posterior distribution for each planet’s density. In practice, stability can only provide upper limits, so we consider the 95th percentile density with realizations weighted by the SPOCK probability of stability. The results of these calculations are shown in Fig. 3.9.

Based on early results, it was expected that about 20 planets would have meaningful density constraints. Unfortunately, this is not what was found. Instead, only about 8 planets were found to have upper limit densities that were below a value of $8g/cm^3$. Of these planets, half come from either KOI-284 or KOI-2248 which are known to have planets falsely identified as members of the




Figure 3.9 Constrained 95th percentile upper limit densities for the Kepler planets versus their radii. Since densities were drawn from a uniform (0,10) distribution, the nominal 95th percentile density would be $\rho = 9.5$. To determine upper limit densities using stability constraints, realizations were weighted by their SPOCK probability and then the 95th percentile density was calculated. As can be seen, only eight planets have a constrained density below $8g/cm^3$ thereby indicating that SPOCK’s probability of stability does not have much ability to constrain densities on its own. Four of the eight significantly constrained planets come from KOI-0284 and KOI-2248 which are known to have false planetary members which dramatically affect the systems stability when they are included.

system (Fabrycky et al. 2014). These systems include planets that are unreasonably close together probably due to blending of two planetary systems within the Kepler aperture. Thus, while they form an interesting point of comparison for our analysis, they do not yield meaningful density constraints.

SPOCK therefore only provides interesting upper limits for a few planets, which is less than expected from looking at Δ SPOCK.

3.4.4 N-body results

The challenges with SPOCK became more apparent once they were compared with n -body integrations. As mentioned earlier, the systems that were directly integrated were selected due to interesting features in parameter-parameter plots colored by SPOCK's probability of stability, similar to Figure 3.6. Examples of these features included extremely small stable regions where if true would give a strong parameter constraint, or a sudden drop in stability values where values went from approximately 0.9 indicating a high probability of stability to 0.7-0.8 which is still expected to be long-term stable, but to a less degree of certainty. However, none of these features were present when showing actual survival time. In almost every realization that was integrated, the system survived to the end of the integration time, despite cases where SPOCK found nearly all realizations of a system to be potentially unstable.

It is not immediately clear what the cause was for the disagreement between SPOCK and the n -body integration results. One consistent feature is a sudden change in SPOCK probability, but not associated with a large change in SPOCK values. If, for example, SPOCK found that low-density planets were stable and high-density planets were unstable, you could have Δ SPOCK values approaching 1. In comparison, the largest Δ SPOCK value was 0.25. For some reason, SPOCK did not show significant changes for any planet over the range of densities tested here. In the overall ensemble, SPOCK did distinguish strongly between unstable and stable systems (although probabilities >0.9 were rare). Keeping in mind that SPOCK is not 100% accurate and may identify features based on unknown internal representations, we focus the rest of our analysis on other metrics.

3.4.5 $e=i$ =small

Allowing for small eccentricities and inclinations did not show significant differences as an ensemble from the $e=i=0$ run. Since these are more realistic but still have the reduced dimensionality, we consider this run to be among our most meaningful.

We now turn to using the Hill Stability metric and specifically analyzed the number and strength of density constraints depending on the circular Hill stability (Δ_{circ}) threshold values of 3, 4, 5, 6, 7, or 8. Systems below the threshold were deemed completely unstable (weight of 0) and above the threshold were deemed completely stable (weight of 1). This is not as ideal as assigning a smooth probability distribution, but the effect of such a distribution can be estimated through considering multiple threshold values.

The resulting density “posterior” upper limits – calculated as above – are presented in Fig. 3.10. For the starting cutoff at 3, very few planets have meaningful density constraints. As the stability threshold increases and a greater dynamical separation is required, additional planets are found to have stability-constrained density upper limits. The challenge then becomes choosing the threshold that we believe has a very low probability of being seen in the ensemble of planetary systems. We chose to take a closer look for the case where the cutoff between stable and unstable is $\Delta = 8$ since there is some precedence for this being a reasonable value for multi-planet systems and a better indicator of long term stability than the minimum $\Delta = 2\sqrt{3}$ value.

Focusing in on a Hill stability threshold of $\Delta_{circ} > 8$, we again investigate which systems also have significant TTVs in Figure 3.11. Of the ≈ 50 planets with a density constraint below $4g/cm^3$, approximately half have strong TTVs thereby resulting in ≈ 25 planets that can have their density uniquely constrained by enforcing stability.

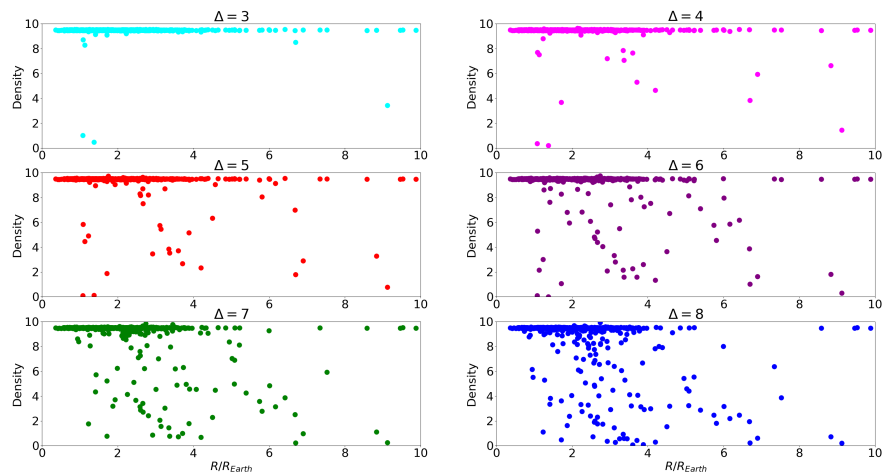


Figure 3.10 Constrained 95th percentile upper limit densities versus radii for the Kepler planets using Hill stability. For each subplot a cutoff value of either 3,4,5,6,7, or 8 was used to determine the minimum number of mutual Hill radii that must separate planets. After rejecting unstable systems, the 95th percentile upper limit densities were calculated. It is clear that as the stability criterion becomes stricter, the number of planets with meaningful density constraints increases dramatically. It is also valuable to note that while the nominal stability cutoff for Hill stability is $\Delta_{circ} = 2\sqrt{3}$, requiring a minimum separation of $\Delta_{circ} = 8$ is still conservative for long-term stability with 3+ planets.

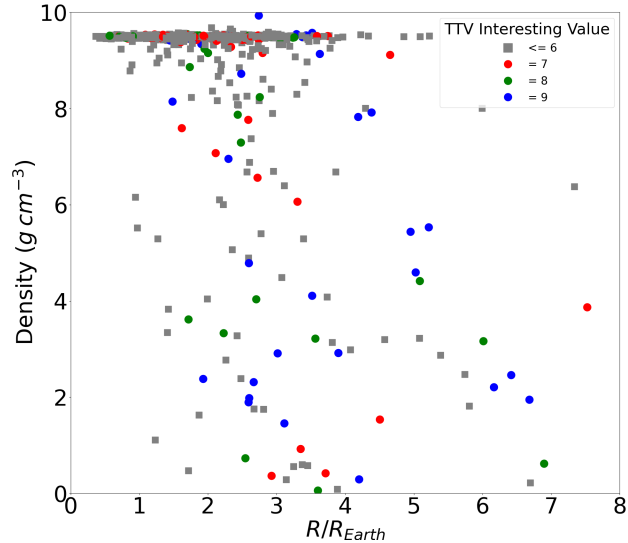


Figure 3.11 Constrained 95th percentile upper limit densities versus radii for the Kepler planets using Hill stability with a threshold of $\Delta_{circ} > 8$ from our $e=i$ =small runs. Points are colored based on TTV interestingness as in Figure 3.8. Of the ≈ 50 planets that have a density constraint below $4g/cm^3$ (meaningful for ruling out compositions), ≈ 25 do not have significant TTVs.

3.4.6 $m=i$ =small

For the run where eccentricity was allowed to float between 0 and e_{cross} we calculated the 95th percentile upper limits on eccentricity with the stability criterion set at $\Delta_{ecc} \geq 8$. In this case, we used Δ_{ecc} because we allowed for larger eccentricity values and thus wanted to account for the eccentricity as we searched for constraints. The results of these calculations are shown in Fig. 3.12. On the left, is shown the constrained eccentricity value versus the planet radius. However, since all the eccentricities were not drawn from the same uniform distribution, we consider on the right the ratio between the constrained and unconstrained 95th percentile values.

Several features from this plot are worth mentioning. First, it can be seen that the majority of eccentricities are constrained to be below values of 0.2. While this is not an exceptionally strong constraint on the Kepler systems as a whole, it still rules out 80% of the eccentricity parameter space and would allow for a tighter prior distribution for eccentricity. Second, we are able to constrain the




Figure 3.12 Constrained 95th upper limit eccentricities versus planetary radii using Hill stability and the criterion that $\Delta_{ecc} \geq 8$ (left). Ratio of constrained to unconstrained upper limit eccentricities (right). For this run, eccentricities were allowed to float up until crossing orbits. As can be seen by the left plot, enforcing stability results in eccentricities clustering below 0.2 which is a loose but non-negligible constraint. From the right plot, it is evident that majority of planets can have their eccentricities constrained to below half of the value for crossing orbits. Therefore, using $\Delta_{ecc} > 8$ reduces the parameter space for physically probable eccentricities could by about half.

eccentricity to below half of e_{cross} for about 60% of the planets. Finally, we will note that there are a number of cases that, no matter the eccentricity, the planetary system had no stable configurations in which case the constrained eccentricity and eccentricity ratio were both 0. While the exact reason for this requires additional analysis of the specific systems, the aforementioned systems with false planetary membership are the cause of eight planets having an eccentricity constrained to 0.

3.4.7 Mutual inclinations

For the run exploring possible mutual inclinations of the system, we find the value of mutual inclination where, if the AMD was redistributed to the eccentricity, the planet pair crosses e_{crit} from the HL metric.

As can be clearly seen in Figure 3.13, the resulting i_{max} is strongly correlated with the period ratio of the two planets. The closer the two planets, the lower i_{max} . Systems with more planets tend

to have lower i_{max} between planet pairs, but this is just due to tighter planet spacing in systems with more planets. The lowest upper-limit mutual inclination between a pair of planets is KOI-1665, where the two known planets with $\frac{P_{in}}{P_{out}} = 0.853$ have $i_{max} = 10.8^\circ$.

Of the 1191 planet pairs we analyzed, 785 can have constraints placed on their mutual inclination. The remaining 406 planet pairs can have any prograde mutual inclination and not violate the HL metric. The majority of the planets with no inclination constraints are planets with small period ratios ($\frac{P_{in}}{P_{out}}$), although some pairs with larger period ratios have no constraints, mostly due to small planetary masses. Clearly, our stability analysis shows that there are few interesting constraints on the mutual inclination of Kepler systems. This justifies the use of stability metrics that assume coplanar orbits, although some care must be taken when analyzing systems with the lowest i_{crit} .

3.4.8 Jontof-Hutter posteriors

Finally, we show the results from applying our stability analysis to the posteriors generated from the TTV analysis from Jontof-Hutter et al.. This analysis will quantitatively identify how stability can trim posteriors from a significant set of Kepler TTV systems.

We began by calculating the upper limit density using our Eccentric Hill metric for three different cutoff values and once again plotted the ratios of the constrained to unconstrained values because the densities were not drawn from a single distribution. The results are shown in 3.14. As seen before, as we increase the Hill stability threshold from 4 to 6 to 8, the number of planets that can be constrained increases. Overall, a large number of planets can be constrained to interesting values. For the highest threshold, several planets have upper limit masses decreased by a factor of ~ 2 from the TTV analysis. Thus, even a fraction of systems with TTVs can benefit from stability constraints.

The posteriors from Jontof-Hutter et al. did not constrain eccentricities as strongly as our other runs. Therefore, we also considered stability constraints on densities using the standard circular Hill metric Δ_{circ} . This modified analysis is shown in Fig. 3.15. Unsurprisingly, the circular Hill metric's




Figure 3.13 The maximum mutual inclination of a planet pair, as calculated based on AMD and HL metric, as a function of planet period ratio. Color shows the total number of planets in the system. Maximum mutual inclination is closely correlated with period ratio, as expected from theory. Number of planets in each systems does not seem to be closely correlated to the maximum mutual inclination.




Figure 3.14 Upper limit density constraints versus planetary radii (left) and ratios of constrained versus unconstrained densities (right) on the analysis of the posteriors from TTV fits given in Jontof-Hutter et al.. The constraint values were calculated using Hill stability with threshold values of $\Delta_{ecc} = 4, 6,$ and 8 from top to bottom. As the stability criterion becomes stricter, the constrained densities become smaller. For the highest threshold, several systems have upper-limit densities reduced by a factor of 2, showing that stability can provide meaningful information even on systems with known TTVs.

ability to constrain density is weaker especially for cutoff values less than 8. Still, using a cutoff of 8 yields ≈ 10 planets where the ratio between their constrained and unconstrained densities is ≤ 0.5 thus indicating a meaningful gain in information for these planets.

3.4.9 Constraint Comparisons

We have shown that different metrics can give different stability constraints.

Although we think it is reasonable to choose $\Delta_{circ} > 8$ as a metric and threshold to constrain stability, we provide our upper limit densities and eccentricities for a wide range of metrics and thresholds. These are tabulated in Table 3.1 for our $e=i$ =small runs and Table 3.2 for our analysis of the posteriors from Jontof-Hutter et al..

These tables also include a third version of Hill stability indicated by Δ_{test} . For this metric, R_{mh} is calculated with the neighboring planet having a mass set to zero. Then we find the maximum




Figure 3.15 Upper limit density constraints versus planetary radii (left) and ratios of constrained versus unconstrained densities (right) on the analysis of the posteriors from TTV fits given in Jontof-Hutter et al., Plots are the same as in Figure 3.14, but use the circular Hill stability metric $\Delta_{circ} = 4, 6,$ and 8 from top to bottom. There are significantly fewer planets that have density constraints based on stability. However, for the case where stability requires $\Delta_{circ} \geq 8$ there are ≈ 10 planets that have their 95th percentile densities constrained to less than half their unconstrained value.

density that the considered planet could have before $\Delta_{test} = 4, 6,$ or 8 . Given that this represents a much more theoretical upper limit for densities, it allows for extremely large densities when planet separations are greater. For this purpose, reported densities are capped at $30g/cm^3$ since this value is well outside our considered density range and far from being a meaningful constraint. Finally we also give the value for the 95th percentile upper limit on eccentricities which were found using the criterion that $\Delta_{ecc} \geq 8$ for a realization to be stable.

3.5 Discussion

Looking closer at Figs. 3.2 and 3.3 one question that deserves additional consideration is whether or not metrics such as the Petrovich or Hill stability metric are sufficient for stability analysis on their own. While SPOCK is a highly optimized code, the fact that it performs a short N-body integration makes it require a significantly greater amount of time than the any analytical metric.

KOI	P (days)	R/R_E	$\Delta_{circ} = 4$	$\Delta_{circ} = 6$	$\Delta_{circ} = 8$	$\Delta_{test} = 4$	$\Delta_{test} = 6$	$\Delta_{test} = 8$	SPOCK	e
41.02	6.8871	1.253	9.489	9.489	9.489	30.000	30.000	30.000	9.490	0.26
41.01	12.8159	2.237	9.484	9.484	9.484	30.000	30.000	30.000	9.480	0.17
41.03	35.3335	1.507	9.476	9.476	9.476	30.000	30.000	30.000	9.476	0.33
70.02	3.6961	1.718	9.466	9.466	9.466	30.000	30.000	30.000	9.469	0.18
70.04	6.0985	0.812	9.510	9.510	9.510	30.000	30.000	30.000	9.508	0.13
70.01	10.8541	2.813	9.511	9.511	9.511	30.000	30.000	24.789	9.496	0.13
70.05	19.5777	0.870	9.510	9.510	9.510	30.000	30.000	30.000	9.510	0.14
70.03	77.6113	2.500	9.515	9.515	9.515	30.000	30.000	30.000	9.515	0.44
82.05	5.2869	0.461	9.465	9.465	9.481	30.000	30.000	30.000	9.487	0.09
82.04	7.0714	0.580	9.499	9.499	9.500	30.000	30.000	30.000	9.506	0.09

Table 3.1 Density and eccentricity constraints calculated using both Hill stability and SPOCK for the Kepler planets from our $e=i$ =small run. Constraints using Hill stability were found by requiring planets to be separated by the minimum number of mutual Hill radii given in the column header. We also include a version of Hill stability denoted Δ_{test} which assumes the neighboring planet has zero mass with an upper limit of $30g/cm^3$ on reported densities. The eccentricity values were found using the requirement that $\Delta_{ecc} \geq 8$. The first 10 rows of data are reported here with the remaining data being provided in MRT format.

KOI	P (days)	R/R_E	$\Delta_{circ} = 4$	$\Delta_{circ} = 6$	$\Delta_{circ} = 8$	$\Delta_{ecc} = 4$	$\Delta_{ecc} = 6$	$\Delta_{ecc} = 8$	$\Delta_{test} = 4$	$\Delta_{test} = 6$	$\Delta_{test} = 8$	e
85.02	2.1549	1.441	30.000	30.000	30.000	30.000	30.000	30.000	30.000	30.000	30.000	0.19
85.01	5.8602	2.588	3.366	3.366	3.357	3.144	2.910	2.329	30.000	18.885	7.967	0.07
85.03	8.1328	1.525	12.080	12.080	12.031	11.845	11.449	10.764	30.000	30.000	30.000	0.06
115.03	3.4364	0.513	30.000	30.000	30.000	30.000	30.000	30.000	30.000	30.000	30.000	0.10
115.01	5.4119	2.480	2.615	2.615	2.339	2.370	2.191	1.572	30.000	10.102	4.262	0.03
115.02	7.1261	1.411	12.270	12.270	12.102	12.385	12.172	11.284	30.000	30.000	23.116	0.02
137.03	3.5047	1.567	26.918	26.918	26.918	26.918	26.918	26.922	30.000	30.000	30.000	0.16
137.01	7.6412	4.186	1.097	1.097	1.097	1.097	1.097	1.097	30.000	28.194	11.894	0.08
137.02	14.8596	5.023	0.636	0.636	0.636	0.636	0.636	0.636	30.000	16.322	6.886	0.09
152.03	13.4843	3.286	3.403	3.403	3.403	3.403	3.417	3.434	30.000	30.000	30.000	0.05

Table 3.2 Density and eccentricity constraints calculated using both Hill stability and SPOCK for the Kepler planets considered in Jontof-Hutter et al.. Constraints using Hill stability were found by requiring planets to be separated by the minimum number of mutual Hill radii given in the column header. We also include a version of Hill stability denoted Δ_{test} which assumes the neighboring planet has zero mass with an upper limit of $30g/cm^3$ on reported densities. The eccentricity values were found using the requirement that $\Delta_{ecc} \geq 8$. The first 10 rows of data are reported here with the remaining data being provided in MRT format.

Therefore, if an analytic metric could be a robust predictor of the probability of stability which is given by SPOCK, then SPOCK may not be strictly necessary. Given the almost step-like nature of the relationship between the Eccentric Hill metric and SPOCK shown in Fig. 3.2 and between the Petrovich metric and SPOCK shown in Fig. 3.3, it seems like these two metrics may have the ability to accurately predict the SPOCK probability of stability.

One of the most disappointing findings from this work was SPOCK's inability to provide meaningful constraints on planetary densities as shown in Fig. 3.9. This came as a greater surprise given the clear evidence that when the density for many planets was increased the probability for stability decreased as seen in Fig. 3.7. We therefore, decided to further analyze the reasons for this apparent contradiction. To do this we calculated a theoretical upper limit density which depends on the ratio of SPOCK values above and below a predefined density cutoff value which is similar to the procedure that created Fig. 3.8.

The results of this analysis are shown in Fig. 3.16. Each line represents a different density cutoff value and the ratio values represent the ratio of the SPOCK value above and below that particular cutoff. Therefore, the two extremes are 0 which means that beyond the cutoff the SPOCK probability of stability drops to 0 and 1 which means that the average SPOCK values above and below the cutoff are the same. What can easily be seen is that there is a near logarithmic growth in the theoretical upper limit density as the ratio between density regions approaches unity. This is the reason why a Δ SPOCK value of 0.1 to 0.2 between regions is unable to constrain the density much below the nominal 95th percentile value. If for a density cut at $4000\text{kg}/\text{m}^3$, such as the one used in Fig. 3.8, we assume an average SPOCK value of 0.9 below the cut and a SPOCK value of 0.7 above the cut we would have a Δ SPOCK value of 0.2 which is on the higher end of values found in this work. However, the SPOCK ratio would end up being 0.778 which would produce a theoretical upper limit density only slightly lower the nominal 95th percentile value for a [0,10] uniform distribution like the one used for most simulations. Given this result and with our




Figure 3.16 Theoretical upper limit densities as a function of the ratio of average SPOCK values above and below a particular density value. Each line shows how the density constraint is affected when the location of the break in SPOCK values is changed. What is found is that a very small SPOCK ratio is needed in order to constrain densities significantly below their nominal 95th percentile value. Specifically, if there was a discontinuity in SPOCK values at $\rho = 4$, implying that SPOCK favors a planet less dense than $4g/cm^3$, a ratio of ≈ 0.1 is necessary to constrain the density below $4g/cm^3$ which would essentially require a jump in SPOCK values from 1 to 0.1 at this boundary. This was typically not the case for SPOCK values which tended to transition from stable to slightly less stable (instead of completely unstable) which in turn limits its ability to constrain density values.

simulation setup, the only way for meaningful density constraints while using SPOCK to weight the realizations would be if at some density value SPOCK jumped from high probability to essentially 0 probability of stability. Unfortunately, as shown, this generally does not seem to be the case for SPOCK.

3.6 Conclusion

This work began with two primary objectives. First, given the publication of several metrics for evaluating the stability of exoplanetary systems, we sought to apply a suite of these metrics to the Kepler systems with 3+ transiting planets. After doing so, we would then be in a position to directly compare these metrics and provide recommendations as to which appears most accurate

and effective in determining the stability of these systems. We found that in cases far away from the stability thresholds, the metrics we employed agreed well with one another and with standard reasoning. Near the thresholds, different metrics made different predictions.

Our second objective was to determine the number of planets that could have parameters such as density and eccentricity constrained by enforcing stability. However, we found that the number of planets that are sensitive to stability analysis is directly linked to the metric applied and where the line between stable and unstable is drawn. For example, if probabilities from SPOCK are employed in calculating upper limit densities, then the resulting constraints tended to be quite weak. We therefore conclude that the particular metric and its application can make a significant difference in the results of enforcing stability for exoplanetary systems.

Given this finding, we chose to analyze our realizations of the Kepler systems using the Hill stability metric with the minimum number of mutual Hill radii between planets being 8 as this has been shown to be a reasonable criterion (e.g. Obertas et al. 2017; Pu & Wu 2015). Using this as our definition for stability, we found that ≈ 50 planets could have their density constrained to physically meaningful values. Since only about half of these planets show significant TTVs, stability provides new constraints on planetary densities. When TTVs are present, stability can still contribute. This was shown by applying stability cuts to a set of TTV posteriors from Jontof-Hutter et al.. An additional $\approx 10/50$ planets could have their densities constrained.

Between these two analyses, we conclude that by enforcing stability, the densities of ~ 35 Kepler planets can be constrained. While this is relatively a small number compared to the entire Kepler catalog of exoplanets, it would still greatly benefit the mass-radius diagram for exoplanets. These enhanced density estimates will also aid in determining compositions which, when combined with orbital information, can assist the work in modeling exoplanetary formation and evolution.

In addition to densities, we found that a significant fraction of eccentricities could be constrained (compared to crossing orbits) based on the requirement of long-term stability. By applying the

same criterion as for densities, we found that for most planets, their upper limit eccentricity was less than half of their crossing eccentricity e_{cross} which resulted in many planets being required to have eccentricities less than 0.2. A similar investigation into mutual inclinations did not provide particularly meaningful constraints.

We provide tables summarizing our constraints for all 227 Kepler 3+ planet systems and a variety of metrics and thresholds. Future investigations can use these results to better understand how stability constrains the orbital and physical parameters of Kepler exoplanetary systems.

3.7 Acknowledgments

We thank Antoine Petit and others for helpful discussions related to this paper. This material is based upon work supported by the National Science Foundation under Award No. 2143195.

Chapter 4

Interior Modeling

We make use of the recently published Kepler Multis Dynamical Catalog in order to perform a large scale modeling of interior compositions of 830 exoplanets. This was made possible through the application of the software tool ExoMDN which constructs posterior distributions for the mass and radius fractions of a four component interior model. We present the 1σ uncertainty ranges for the eight modeled parameters for all 830 Kepler systems along with an in-depth look at KOI-277, 351, and 3158 as these are particularly interesting/important systems within the Kepler dataset. Additionally we look into the Kepler Radius Gap in the context of the mass and radius fractions of the H/He layer and find that when the distinction between rocky and gaseous is set by the radius contribution of the H/He layer, we can accurately predict the location of the radius gap.

4.1 Introduction

The Kepler Space Telescope (“Kepler”) was the single most successful mission for detecting exoplanetary candidates. After observing a field of approximately 200,000 stars for four years, Kepler identified thousands of candidates which have been verified as true exoplanets. One of the primary benefits of using the Kepler dataset is that the large number of planets allow for inferences

to be drawn on a class of exoplanets and not just for an individual one or system. The fact that all of the Kepler planets were identified in a homogeneous manner means that analysis can be performed between systems or on the entire catalog of systems without introducing additional biases and thus making it a simpler task to work with the data.

One problem that has inhibited large scale analyses of the Kepler dataset is the lack of a source of self-consistently derived physical and orbital parameters for all the planets. Recently, the Kepler Multis Dynamical Catalog (KMDC) has made a significant step forward in remedying the situation. Jones et al. (in prep.) performed a photodynamical analysis on all 719 Kepler systems with multiple transiting planets which when completed yielded converged Bayesian posterior distributions for all the physical and orbital parameters for 661 of those systems. The reasoning for focusing solely on the systems of multiple transiting planets (“multis”) is that these are the most information rich systems due to planet-planet interactions allowing for parameters to be constrained considerably as explained by Ragozzine & Holman (2010). The fact that the KMDC was derived in a Bayesian framework to obtain posterior distributions means that it is simple to obtain a self-consistent set of model parameters, and the uncertainties in each parameter are very well characterized.

In recent years, considerable effort has been put into understanding the demographics of exoplanetary interiors. This problem is significant because it is directly related to planet formation and evolution. Clearly, a planet can only form out of the materials available at its orbital distance, which implies that if a planet has a considerable mass fraction of water or gas, for example, it must have formed where these elements and compounds were located in the proto-planetary disk. If the planet is currently found in a location where this would not be the case, ie a tight orbit around its star, then it is obvious evidence of planet migration. Specific endeavors include looking at the compositional structures of individual planets and systems performed by groups such as Egger et al. (2022), Brinkman et al. (2024), Carter et al. (2012), and many others. Complex interior modeling codes like the one given in Huang et al. (2022) allow the user to define the mass fractions

of model layers, choose various compositions and equations of state, and can then return the planet's theoretical radius. Updated mass-radius plotters have been published by Castro-González et al. (2024) and Parviainen et al. (2024) which also allow model mass-radius trend lines for various compositional structures to be over-plotted with the data based off of recent publications. The abundances of individual compounds in exoplanets have been examined by groups such as Rogers et al. (2025) who specifically analyzed the amount of water in super-earths. Additionally, some of the research for understanding the Kepler Radius Gap includes work done by Luque & Pallé (2022), Ho & Van Eylen (2023), and VanWyngarden & Cloutier (2024).

However, despite the results from these publications as well as those not listed, there has not been a significant effort to obtain theoretical compositional models for the Kepler dataset as a whole. Two of the current largest analyses include work by Armstrong et al. (2025) who presented results from the NCORES observing program and combined their results for 35 planets along with archival known planets to look at compositions around and below the radius gap. The second larger study is presented by Dainese & Albrecht (2025) who employs the interior modeling software ExoMDN to investigate the evidence for water worlds on a sample of 60 exoplanets that have well constrained masses.

Given, that the KMDC contains planetary parameters for nearly 1700 planets across 661 systems, it provides exactly what is needed to perform a true large scale analysis of the interior compositions of the Kepler multis. Similarly to Dainese & Albrecht (2025), we make use of the ExoMDN code as it can predict the mass fractions of a four component model consisting of an iron core, silicate mantle, water/ice layer, and a H/He atmosphere. ExoMDN makes use of a mixture density model which, as it is a trained machine learning algorithm, makes obtaining model results extremely fast and computationally inexpensive. Our methods for performing this analysis are given in Section 4.2 where we give a summary of ExoMDN, our planet selection and input generation process, and the required computational resources needed. After applying the constraints of ExoMDN's priors,

we were left with a sample of 830 planets. Our results, including an in-depth look at a couple of systems, are presented in Section 4.3. Following this, we present an additional discussion and analysis of how our results fit in the larger framework of exoplanetary theory in Section 4.4. Finally, our conclusions are given in Section 4.5 with some additional thoughts for future work and potential projections found in Section 4.6.

4.2 Methods

Our methods for this work were quite simple. We obtained posterior draws from the KMDC, used them as the initial values for an interior modeling code ExoMDN, and obtained posterior distributions for the interior compositions of the modeled planets. Specific details on how this was performed are in the following subsections, where we cover a summary of ExoMDN, our generation of input parameters and the number of planets in our sample, and the computational resources necessary to perform this work.

4.2.1 ExoMDN

We utilized the publicly available exoplanetary interior modeling software ExoMDN. For a complete explanation of ExoMDN, we direct the reader to Baumeister & Tosi (2024), but a short summary of their methods will be given here as it pertains to this work as well. ExoMDN is a machine learning software that implements a Mixture Density Network (MDN) which is a combination of a standard neural network and a mixture density model as introduced in Bishop (1994). This combination allows the MDN to extend the functionality of a neural network and be able to generate Bayesian posterior distributions for the fitted parameters.

The model was trained on a set of 5.6 million synthetic planets. The planets were initially constructed by sampling from Baumeister & Tosi (2024)'s priors on mass, equilibrium temperature,

and layer mass fractions. Each planet was defined as having four distinct layers that were an iron core, silicate mantle, water/ice layer, and a H/He atmosphere. The sampling for the masses of each layer was done uniformly with the constraint that the fractions had to add to 1. The total masses of the planets were chosen uniformly between 0.1 and $25 M_{\oplus}$ and the equilibrium temperatures were sampled from 100-1000K. With the total mass, equilibrium temperature, and mass layers defined, these values were passed to their private modeling code TATOOINE which calculates the radial profile of mass, pressure, and density using equations for mass conservation, hydrostatic equilibrium, and the equation of state Baumeister & Tosi.

ExoMDN was validated by using the TATOOINE software to access radius accuracy and by an independent inference approach which involved assessing how the ExoMDN posteriors compared to a standard Monte-Carlo sampling method. After both of these tests, it was found that ExoMDN performed well overall, but had a tendency to overestimate the radii by about 0.4%. The final ExoMDN model consisted of three hidden layers with 384 nodes per layer, takes in as inputs the planets mass, radius, and equilibrium temperature, and returns the log-ratio transformed values for the masses ($\ln w_{\text{mantle}}/w_{\text{core}}, \ln w_{\text{water}}/w_{\text{core}}, \ln w_{\text{gas}}/w_{\text{core}}$) and thicknesses ($\ln d_{\text{mantle}}/d_{\text{core}}, \ln d_{\text{water}}/d_{\text{core}}, \ln d_{\text{gas}}/d_{\text{core}}$) of each layer. These log-ratio parameters are then internally transformed to give the mass and radius fractions of each layer compared to the total mass and radius of the planet.

We also make use of Baumeister & Tosi's software for plotting the posterior distributions for the mass and radius fractions of each layer. Given the requirement that the total mass and radius equal 1 times the mass or radius of the planet, the full range of mass or radius fractions for each layer can be thought of as being contained within a 4 dimensional hyper-plane. Each subplot within the corner plot, generated by the modeling code, is then a 2D slice of this 4D space, where the remaining two dimensions have been marginalized over.

4.2.2 Inputs

ExoMDN assumes the initial parameters are exact and therefore does not attempt to model any uncertainty in those values. To account for uncertainties in the input parameters, ExoMDN must be run multiple times with the input parameters drawn from the error ranges for each of the parameters. Additionally, it is necessary to specify the number of draws from the generated posterior distribution. In the end if there are n sets of initial conditions that are sampled from the uncertainty range for the planet, and m draws per set of conditions, the result is an $n \times m$ sample of posterior draws which can be flattened and combined to yield the full posterior distribution for the compositional layers that accounts for the uncertainty in the input parameters.

Rather than just having error ranges for each parameter, we made use of posterior draws from the Kepler Multis Dynamical Catalog (KMDC) published in Jones et al.. These posterior draws were obtained after applying a full photodynamical model to all of the Kepler systems with multiple transiting planets and thus represent the most complete catalog of planetary parameters. Each posterior is a self-consistent set of parameters for the planets, which can be used to accurately model the given system. To provide a reasonable estimate and exploration of the uncertainties on the masses and radii within the KMDC we began with a random sample of 1000 posterior draws for each of the 1858 planets in the catalog. Given that ExoMDN is only useable in the region of parameter space that it was designed and trained on, it was then necessary to thin our sample of planets based off of these constraints. Specifically, we could only use a given posterior draw if the mass was between 0.1 and $25 M_{\oplus}$. In the interest of modeling as many planets as possible, and not wanting to skew results by throwing out a significant number of posteriors draws for a planet, we applied the following criterion: If at least 90% of the posterior draws (ie 900) fell within the allowed range, then the planet was kept in our sample and its interior was modeled using ExoMDN. After applying this criterion, we were left with a sample of 830 planets across 465 systems.

As noted, one of the required input parameters for ExoMDN which is not directly contained in the KMDC is the equilibrium temperature for the planet. To calculate this we used the equation:

$$T_{eq} = T_* \sqrt{\frac{R_*}{2a}} (1 - A_b)^{1/4} \quad (4.1)$$

Both the radius of the host star and the planet's semi major axis were contained or could be calculated directly within the posterior draw. To obtain stellar temperatures we used the Gaia catalog given in Berger et al. (2020) as it provides the most accurate and uniformly obtained stellar parameters. This was also the same catalog that was used to obtain required stellar parameters for generating the KMDC and thus we were able to maintain self-consistency by utilizing it here as well. Exoplanetary albedos are not well-defined and thus it was not possible to obtain a distinct value for each planet based off of its known properties. Given this, we initially gave every planet a reasonably moderate albedo of 0.5. As stated previously, ExoMDN also had an allowed range on temperatures which was 100-1000K. Unlike with masses where posteriors were removed which could in turn lead to the removal of a planet, we applied a simple adjustment to calculated equilibrium temperatures if they nominally fell outside of the allowed range. If the temperature was greater than 1000K it was set to 990K and if it was less than 100K it was set at 110K. The rationale for this was that masses are of greater importance within modeling the interiors, and adjusting them to fit within the allowed range would have a much more significant effect compared to a small adjustment in the equilibrium temperatures.

Although it was generally expected that the dominant parameters would be the planetary mass and radius, after some exploration of the models it was found that equilibrium temperatures played a meaningful part in determining the mass and radius fractions of the atmosphere and water/ice layers. This is due to the fact that with only a small amount of H/He by mass, it is possible to greatly increase the planet's radius especially if the planet has a significantly high equilibrium temperature, which is the case for many Kepler planets. To investigate this further, we decided to generate two

datasets, which will be presented in the following sections. The first follows the method already described where the equilibrium temperature was calculated using an albedo of 0.5. The second dataset was generated where all equilibrium temperatures were set to 150K and gives the results for the case where volatile layers are not significantly inflated due to stellar heating.

4.2.3 Computational Resources

As explained in Baumeister & Tosi because the model is already trained, it is extremely easy and computationally inexpensive to obtain a sample from the posterior distribution for all parameters of the interior model. For this work, we chose to obtain 100 posterior samples for each of the 1,000 input posterior draws to give a total sample size of 100,000 samples. Each set of 100 samples could be obtained in 1 second and therefore the entire posterior sample was obtained in under an hour for each planet. Although this is remarkably fast, given that we were obtaining interior models for 830 planets, we chose to run this analysis using the supercomputer located on the campus of Brigham Young University. Each of the 465 systems were modeled in parallel as it was easier to parse starting data within a system than completely individually for every planet. Since the posterior samples for an individual planet could be obtained in under an hour, the time needed for an entire system ranged from under an hour to a couple of hours for the higher multiplicity systems. Using the large number of CPUs available, the entire dataset was obtained in approximately 1 day.

4.3 Results

Due to the large number of individual planets that we modeled in this work, it is impossible to analyze and assess the results of the interior models on a planet by planet basis here. Instead, we will primarily present our findings on a catalog level and allow the reader to investigate specific systems further by making use of our publically accessible results. We do display in depth results

for KOI-277, 351, and 3158 as these systems are important either because they are modeled in Baumeister & Tosi, have well defined mass measurements from the KMDC, or in the case of KOI-351 are the system with the highest multiplicity in the Kepler database. Where applicable and informative, we will also present a side by side comparison between our hot and cold datasets.

4.3.1 Catalog Compositions - Hot

It is of particular interest to determine whether there are significant trends in the compositional structures of the Kepler multis as a whole. Our work in addressing this is shown in Figure 4.1. Each cornerplot displays two key sets of information. The first is all the combinations of 2D posterior distributions for layer mass fractions. Secondly, we give the 1D posterior distribution for each layer's mass fraction along with the 1σ errors. To generate these combined posterior plots, we only made use of the systems that had mass uncertainties of less than 25% which was done in order to limit the results to the systems that had statistically significant mass measurements in the first place and could thus in theory have more accurate interior models constructed by ExoMDN. When the mass uncertainty is not taken into account, the result is that there is an extremely dominant population of posteriors favoring a solid iron core with no other layers and then a smaller population of with variable compositions. As such, the median mass fractions have these extreme values (0 or 1) with very large uncertainties. With the constraint on mass uncertainty, we were left with a subsample of 65 planets to generate these results.

The first subplot of Figure 4.1 simply shows the combined posteriors for the layer mass fractions obtained from the 65 planets. One of the most prominent features here is the spike at 1.0 in the 1D posterior for the mass fraction of the core. For the models that generated this feature, the planet in question was consistent with a composition composed of 100% iron. While it is possible for a rocky planet to be quite dense, it is nearly impossible for a planet to exist that has this composition. The problem here that ExoMDN was trying to resolve was that there are several planets in the KMDC

that have well measured masses that when combined with their known radii values results in a planet with an unrealistically high density. This observed bias in the KMDC has yet to be addressed in its entirety, but is likely primarily a result of missing features (ie stellar variability or unseen planet) in the model for the system. With this in mind, we decided to remove posterior samples where the mass fraction of the core was greater than 0.9 and the results of this clipping of unrealistic compositions is shown in the second subplot of Figure 4.1. What is now more easily seen is that within the 1σ error range nearly any compositional structure made up of these four layers is allowed. Each 1D posterior distribution also displays a significant spike at 0 indicating that it is also often possible to describe the planet with a model with a maximum of three layers and potentially fewer if one of the layers is moved to even more extreme values. A potentially interesting and surprising result is how small the gas component tends to be for this subsample. It is generally accepted that the Kepler systems are either rocky planets with little to no atmosphere or rocky planets with a significant component of gas which then inflates their size to the sub-Neptune regime. Therefore, it would make more sense to see two features in the gas posterior, a spike at 0 and another spike at some slightly larger mass fraction. The second spike that is seen is at essentially 1.0 indicating a composition of 100% H/He which is more likely a result of a bias in the KMDC. Finally, in the third subplot, we show the posteriors weighted by planet occurrence rates. Using the values from Hsu et al. (2019), we assigned an occurrence rate to each of the planets in our subsample. Then each of the 100,000 posterior draws per planet were added into this combined result with a probability equal to the planet's occurrence rate. The difference that weighting the posteriors by occurrence rate is almost negligible except for the case of the atmospheric layer's mass fraction. Unweighted the median mass fraction and 1σ errors are $0.02^{+0.93}_{-0.02}$ whereas the weighted median and error range are $0.00^{+0.39}_{-0.00}$. This results is explainable when considering the context of the Kepler dataset. It is known that rocky planets are more common than gaseous ones so when taking this into account, the gas component is going to be significantly suppressed (Hsu et al. 2019).

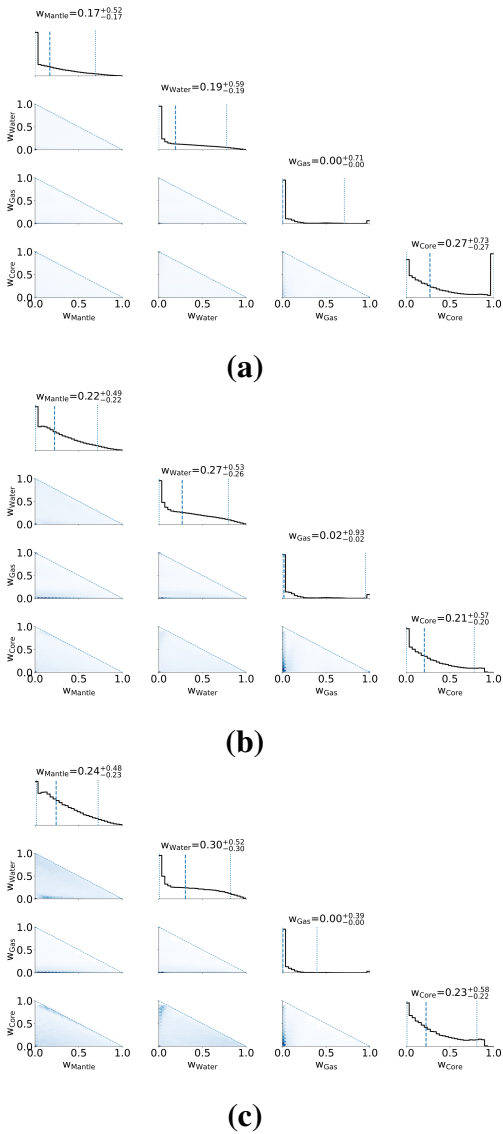


Figure 4.1 Combined corner plot giving the posterior distributions for the four compositional layers plotted against each other and then marginalized to show the one dimensional distribution using the hot dataset. These distributions were obtained by simply combining the 100,000 posterior draws for all of the planets with a mass uncertainty of 25% or less. This purpose in focusing on this subset of planets was so that the distributions would not be dominated by high mass and therefore compositional uncertainties. The first subplot gives the raw combined posteriors which contained a large number of samples that specified an unrealistically high core mass fraction. To focus in on more probable compositional models, subplot b shows the same data only where any posterior with a core mass fraction greater than 0.9 has been removed. Finally, subplot c has high core mass posteriors removed and each posterior has been weighted based off of the planet’s occurrence rate given by Hsu et al.. One interesting feature that can be seen is that every posterior has a large spike at a mass fraction of 0, indicating that it is possible to model the planet with 3 or less layers.

4.3.2 Catalog Compositions - Cold

As mentioned, it became clear that it was of interest to investigate the differences in model compositions when the planets were given reasonably accurate equilibrium temperatures and when they were given temperatures near the minimum of ExoMDN's prior. Figure 4.2 shows the exact same results as Figure 4.2 except that in this case the input equilibrium temperature for every planet was 150K. Between the first two subplots of each respective figure, there is essentially no difference in the median and 1σ error ranges except when it comes to the gas and core components of the planets. Where in the hot dataset, the ranges for gas and core were $w_{Gas} = 0.00^{+0.71}_{-0.00}$ and $w_{Core} = 0.27^{+0.73}_{-0.27}$, for the cold dataset the ranges are $w_{Gas} = 0.07^{+0.91}_{-0.07}$ and $w_{Core} = 0.19^{+0.81}_{-0.19}$. First off, these results are consistent with one another at the 1σ level due to the large error range. However, by cooling the planets to subfreezing temperatures, ExoMDN allowed a statistically significant greater amount of gas to be present in the planet's atmospheres. Looking next at the second subplot of each figure, which removes the planets with non-physical compositions, the trend of enhanced gas for the cold dataset becomes more pronounced and it can be seen that there may in fact be a secondary small maximum in gas mass fraction around 0.35. The overall larger amount of gas contained within the cold dataset combined with the difference in the posterior distribution itself does seem to agree better with current expectations for the Kepler database. In the case of the third subplot, by weighting the posteriors based on the occurrence rate of the planets, there is again a suppression of the gas component as the 1σ error range is changed to $w_{Gas} = 0.07^{+0.70}_{-0.07}$, but this is a larger component than all three cases in the hot dataset.

4.3.3 Catalog Compositions - Summary

Looking at the results of this work as a whole regardless of which temperature condition is used, a couple of general trends stand out. First, as mentioned when looking at the hot dataset, but also seen in the cold, ExoMDN seems to prefer models with a lower complexity/dimensionality than

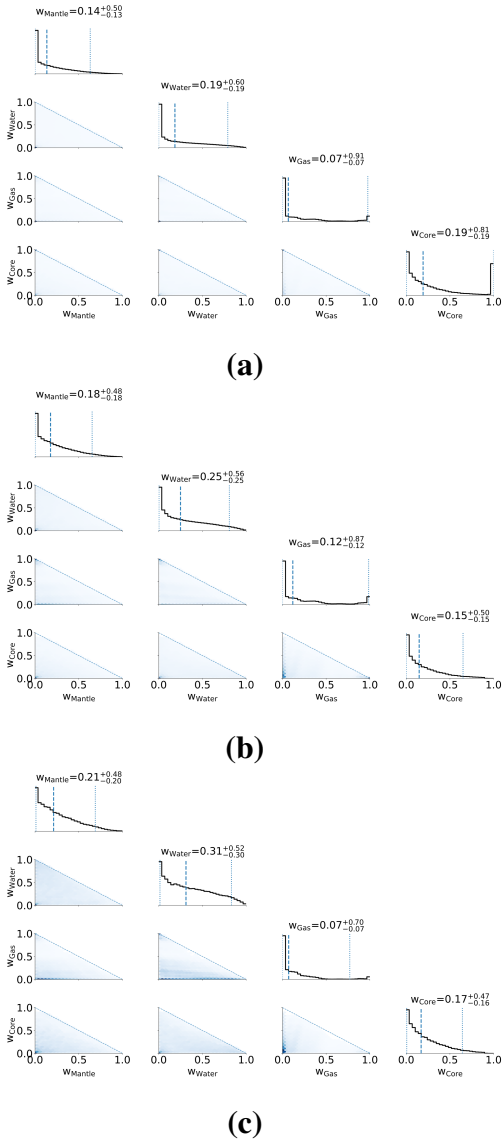


Figure 4.2 The exact same plots as in 4.1 except for the cold dataset. Once again, subplot a has all of the posterior draws from the planets that have a mass uncertainty of 25% or less, subplot b has the posteriors draws claiming a core mass fraction of 0.9 or greater removed, and in subplot c the posteriors are weighted by planet occurrence rates given in Hsu et al.. While generally the posteriors are quite similar. The main difference between the hot and cold datasets is the amount of gas that is predicted by ExoMDN. For the hot dataset the median gas mass fractions across the three subplots are 0.00, 0.02, and 0.00 whereas for the cold dataset the medians are 0.07, 0.12, and 0.07.

the four component model. While it is not immediately clear why this behavior exists, it does help to illustrate the inherent level of degeneracy in modeling the interiors of exoplanets based on just three parameters. Given, the wide range of physical properties between each of these components, it is possible to simulate planets of intermediate densities with either an extreme amount of gas and core or a moderate amount of mantle and water. Without additional constraints on the planet's structure itself, it is not an easy task to determine which model is correct. The second general trend is simply that these planets do not contain a large percentage of H/He by mass. This is perfectly understandable in the sense that it only takes a small percentage of lighter gases to inflate a planet's radius significantly. Therefore, planets of moderate radius values do not have extreme mass percentages of gas. It is also worth noting again that the maximum mass value for ExoMDN is $25\oplus$ and as a result of this all of the gas giant planets have been excluded from this analysis.

4.3.4 Individual Planet Analysis

As previously mentioned, it is not feasible to give an in depth analysis of the ExoMND results for all of the systems of planets that were modeled. Here we will present a closer look at three particularly noteworthy systems, which are: KOI-277/Kepler-36, KOI-351/Kepler-90, and KOI-3158/Kepler-444. Following these analyses, Tables 4.1 and 4.2 give the KOI number, period, mass, radius, equilibrium temperature, mass fractions ($w_{Core}, w_{Mantle}, w_{Water}, w_{Gas}$), and radius fractions ($d_{Core}, d_{Mantle}, d_{Water}, d_{Gas}$) for the planets in the first 10 systems analyzed. The tables in their entirety can be obtained from the supplemental materials to this paper in a machine readable format. Table 4.1 gives the parameter values for the hot dataset while Table 4.2 gives the parameter values for the cold dataset.

KOI-277/Kepler-36

KOI-277/Kepler-36 is an important system to look at for a couple of reasons. First, it has been studied extensively using photodynamical models like the one used to generate the KMDC which means that there is a considerable amount of information about this system in previously published papers. Second, this is a modeled system in Baumeister & Tosi which means that we can directly compare our results in order to validate our process for using ExoMDN. This system contains two planets orbiting in a 7:6 mean motion resonance (MMR) with periods of 13.869 and 16.217 days for Kepler-36 b and c respectively. The masses, radii, and equilibrium temperatures used were $M_b = 4.06^{+0.254}_{-0.256} M_{\oplus}$, $R_b = 1.54^{+0.045}_{-0.039} R_{\oplus}$, and $T_{eq-b} = 902.7\text{K}$ for planet b with $M_c = 7.45^{+0.464}_{-0.451} M_{\oplus}$, $R_c = 3.78^{+0.096}_{-0.084} R_{\oplus}$, and $T_{eq-c} = 856.9\text{K}$ for planet c. These values differed slightly from Baumeister & Tosi who used $M_b = 3.83^{+0.11}_{-0.11} M_{\oplus}$, $R_b = 1.5^{+0.061}_{-0.061} R_{\oplus}$, and $T_{eq-b} = 978\text{K}$ for planet b with $M_c = 7.13^{+0.18}_{-0.18} M_{\oplus}$, $R_c = 3.68^{+0.096}_{-0.096} R_{\oplus}$, and $T_{eq-c} = 928\text{K}$ for planet c. Looking now at the parameter values obtained for each layer which are plotted in Figure 4.3, we obtained for planet b: $d_{Core} = 0.66^{+0.07}_{-0.11}$, $d_{Mantle} = 0.12^{+0.2}_{-0.1}$, $d_{Water} = 0.11^{+0.12}_{-0.08}$, $d_{Gas} = 0.09^{+0.1}_{-0.05}$ and for planet c: $d_{Core} = 0.19^{+0.09}_{-0.11}$, $d_{Mantle} = 0.13^{+0.17}_{-0.11}$, $d_{Water} = 0.17^{+0.23}_{-0.15}$, $d_{Gas} = 0.49^{+0.11}_{-0.11}$ whereas Baumeister & Tosi found $d_{Core} = 0.68^{+0.09}_{-0.12}$, $d_{Mantle} = 0.10^{+0.19}_{-0.09}$, $d_{Water} = 0.10^{+0.12}_{-0.08}$, $d_{Gas} = 0.08^{+0.10}_{-0.04}$ for planet b, and $d_{Core} = 0.19^{+0.09}_{-0.11}$, $d_{Mantle} = 0.13^{+0.19}_{-0.12}$, $d_{Water} = 0.19^{+0.24}_{-0.17}$, $d_{Gas} = 0.46^{+0.12}_{-0.11}$ for planet c. Thus, despite our slightly differing input parameters, the resulting models are nearly identical. This seems to suggest that the inherent degeneracies and uncertainties within ExoMDN can dominate over the exact input parameters to a certain level. A detailed analysis of Kepler-36 was given by Carter et al. (2012) and as part of this work they performed Bayesian parameter inference to determine the limits on the compositional layers. From this, they find that Kepler-36b is likely to be rocky with a core mass fraction of $w_{Core} = 0.29^{+0.11}_{-0.10}$ a potential water layer with $w_{Water} < 0.23$ and an upper limit on the gas layer of $w_{Gas} < 0.01$. From our analysis we find $w_{Core} = 0.8^{+0.14}_{-0.27}$, $w_{Water} = 0.04^{+0.08}_{-0.03}$, and $w_{Gas} = 0.0^{+0.0}_{-0.0}$. These results are roughly in

agreement with each other in the sense that Kepler-36b is unlikely to have any H/He and could have a small water layer. Our values disagree more substantially for the core mass fraction, but given our larger uncertainty region this is not a major issue and the ultimate conclusion of what this planets physical properties are remains the same. Carter et al.'s main purpose in their compositional analysis of Kepler-36c was to determine the mass fraction of its gas layer given its low density. They found two potential solutions for the mass fraction which were $w_{Water} = 0.086_{-0.01}^{+0.01}$ when there is no substantial water layer and $w_{Water} = 0.016_{-0.004}^{+0.004}$ when there is a water layer equal in mass to the core. Our results found $w_{Water} = 0.07_{-0.04}^{+0.05}$ which fortunately are in good agreement with both scenarios. Given all of these results, it can be concluded that ExoMDN performed well when modeling Kepler-36's planets in that it was able to match the current best understanding of the planetary compositions.

KOI-351/Kepler-90

KOI-351/Kepler-90 is an exciting system to model as it has the highest multiplicity of any Kepler system with a total of eight confirmed planets. It has been studied by several independent groups including Cabrera et al. (2014), Rowe et al. (2014), Schmitt et al. (2014), and Weiss et al. (2024). As of yet there has not been an attempt to model the interior compositions of these planets as current methods are still primarily putting upper limit constraints on the masses. Our results for this system are shown in Figure 4.4. We only were able to model six of the eight planets as Kepler-90b and h had too many posterior draws outside of ExoMDN's mass prior. Going through our findings for the remaining planets, we found that Kepler-90g has a H/He component $w_{Gas} = 0.98_{-0.005}^{+0.01}$ essentially requiring the entire planet to be made up of gas. This result is caused by the fact that from the KMDC, Kepler-90g has a density of only $\rho = 0.16_{-0.013}^{+0.013}$ and then ExoMDN is forced to find a composition that can match such a low density. This problem can be solved significantly when looking at the mass measurements published in Weiss et al. which gives $M_g = 49_{29}^{29} M_{\oplus}$ which

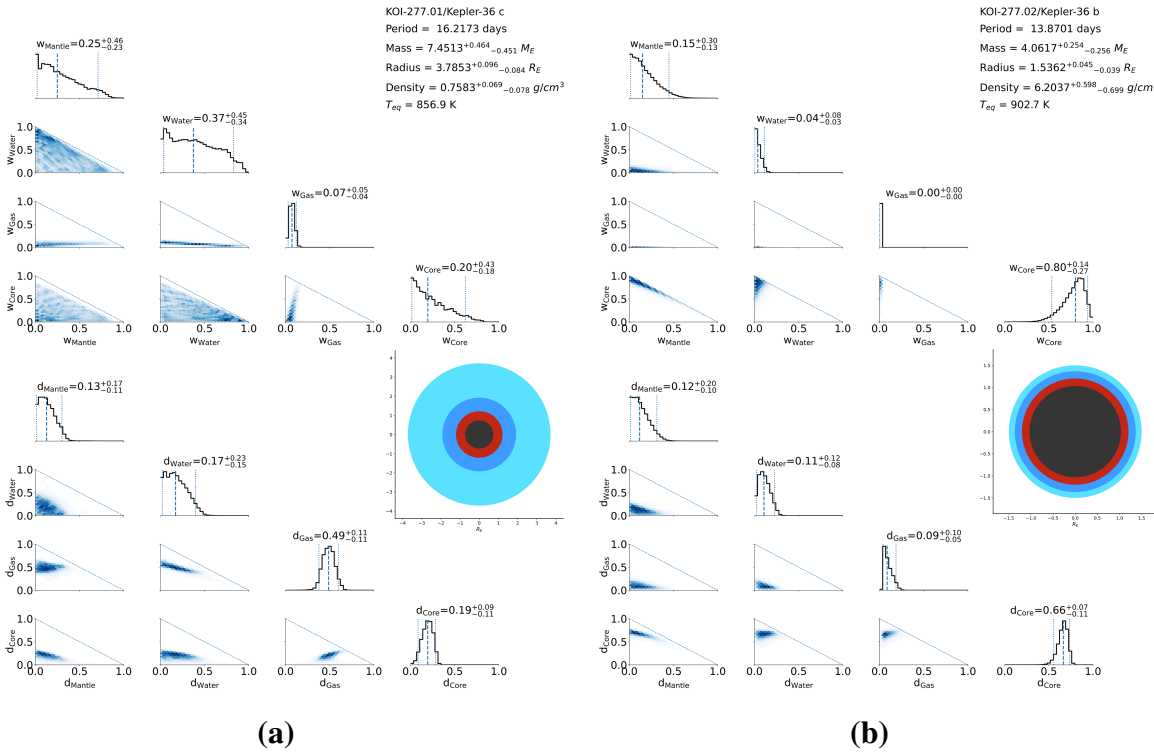


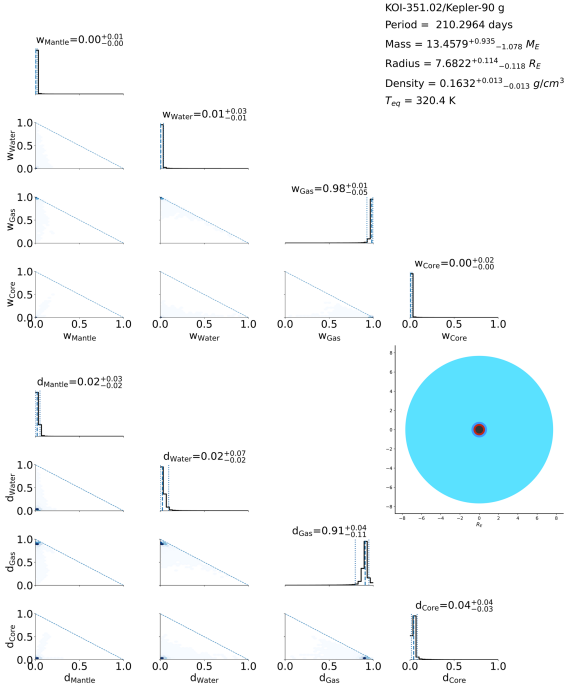
Figure 4.3 Posterior distributions for the four compositional layers for the two planets in KOI-277/Kepler-36. Included in each subplot are the posteriors for both the mass and radius fractions, the parameters used to model the system, and an image showing the median layer thicknesses plotted on top of each other to display a theoretical model of the planet. Since simply using the median values to generate a planetary model is not entirely accurate as they are not a self-consistent set of layer thicknesses, this image is more of a general illustration of the planets composition. In the case of Kepler-36 the two planets display considerable differences with planet c having a high probability for a significant water and nonzero gas layer. On the other hand, planet b is likely to be rocky with its composition almost entirely being made up of an iron core and a modest mantle component.

is significantly higher than the mass found in the KMDC which is $M_g = 13.5_{1.08}^{0.93} M_{\oplus}$. It is likely, that the mass given in the KMDC is in fact too low and that the true mass is closer to the one presented in Weiss et al. which when used would give a composition that is much more realistic. For Kepler-90d we find a composition consistent with the class of sub-Neptunes that are prominent in the Kepler database. What is interesting for this planet is the posterior distribution for the water mass fraction is essentially flat with a 1σ range of $w_{Water} = 0.46_{-0.41}^{+0.42}$ thus indicating that it would be possible for this planet to be a water world. When looking at the radius components though, it is seen that based off of the median values (which has a much more moderate water fraction) each layer contributes essentially an equal amount to the total radius. The results for Kepler-90e are surprisingly similar to those for planet d. There is a relatively flat posterior distribution for the mass fraction of water with a 1σ range of $w_{Water} = 0.51_{-0.46}^{+0.38}$ again indicating the possibility of this planet could be a water world. The median mass fraction of H/He is only 0.02, but despite this, the median atmospheric component in terms of radius is 0.25 showing that a large mass fraction of gas is not necessary in order to provide considerable enhancement to the planet's radius. Kepler-90c is one of the smaller interior planets in the system where the 1σ ranges for the four compositional components are $w_{Core} = 0.65_{-0.57}^{+0.34}$, $w_{Mantle} = 0.20_{-0.19}^{+0.58}$, $w_{Water} = 0.06_{-0.06}^{+0.29}$, $w_{Gas} = 0.00_{-0.00}^{+0.00}$. A clear difference between this planet and the others previously analyzed is at the 1σ level there is still a need for a nonzero mass core which gives supporting evidence that Kepler-90c is going more likely to be a rocky planet. Kepler-90f is once again a return to the sub-Neptune regime seen in the results for planets d and e, however planet f has essentially half the mass of planets d and e while maintaining a very similar radius. Due to this, its density measurement from the KMDC is substantially lower and ExoMDN estimates the median mass and radius fractions of the atmosphere at 0.07 and 0.47. Finally, Kepler-90i a bit more of an anomaly in the system. The masses and radii from the KMDC were $M_i = 2.32_{-1.624}^{+2.586} M_{\oplus}$ and $R_i = 1.05_{-0.15}^{+0.142} R_{\oplus}$ which results in a bulk density of $\rho_i = 12.02_{-8.553}^{+11.205} M_{\oplus}$. Given such a high density it is no surprise that the interior models favored

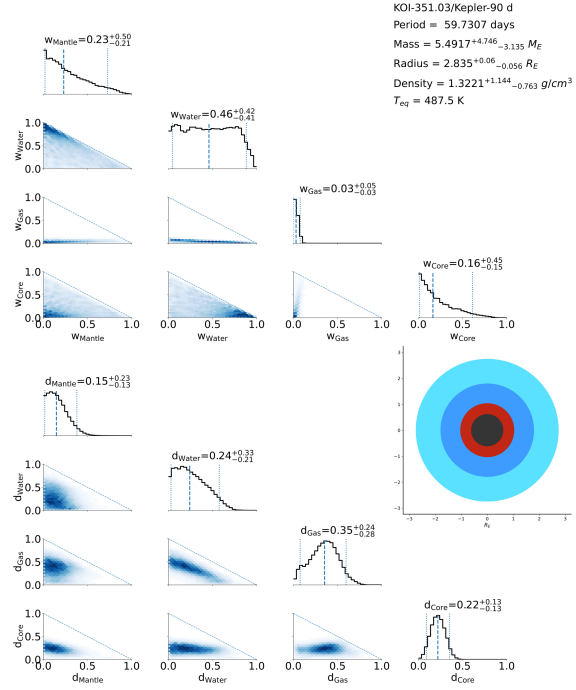
an iron core accounting for up to 100% of the planet's mass. However, it should also be noted that the wide uncertainty in this planet's densities means that it would not be unlikely for the planet to have a much lower density and thus a smaller mass contribution from the core.

KOI-3158/Kepler-444

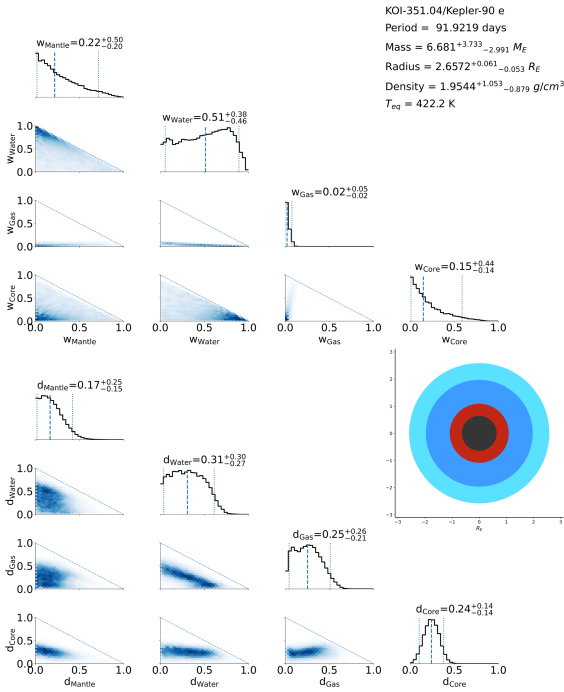
KOI-3158/Kepler-444 is a system of five mars sized planets orbiting within 0.1 AU from their host star. It has been studied by Weiss et al., Mills & Fabrycky (2017), and Hadden & Lithwick (2017), but mass measurements have only been locked in for planets d and e which orbit with periods of 6.18 and 7.74 days. Additionally, there has been no work done on modeling the interior structures of these planets. This is also a particularly noteworthy system as it has some of the tightest mass constraints out of any found in the KMDC. Based off our results from ExoMDN, which are shown in Figure 4.5, there are really only two classes of structures that exist for the four planets modeled. The inner two planets Kepler-444b and c were found to need to be composed 100% of an iron core. Then since Kepler-444 d and e have nearly identical masses and radii, their composition made up the second class which was a considerable core mass fraction of $w_{Core} = 0.78^{+0.16}_{-0.36}$ and a moderate mantle mass fraction of $w_{Mantle} = 0.2^{+0.38}_{-0.18}$. Each of these planets could also have a few percent of water by mass, but no measureable amount of H/He. It is worth noting though that the masses found by Mills & Fabrycky for planets d and e were smaller by a factor of 4. Therefore, if those masses were to be found to be more reliable than the ones published in the KMDC, which also agrees with those found in Hadden & Lithwick, the reduction in mass would allow for a substantially larger component of volatiles in either the form of water or H/He.



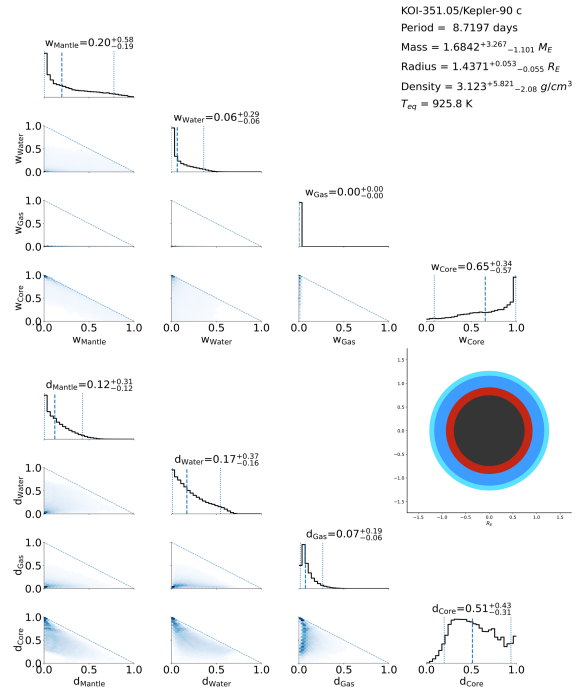
(a)



(b)



(c)



(d)

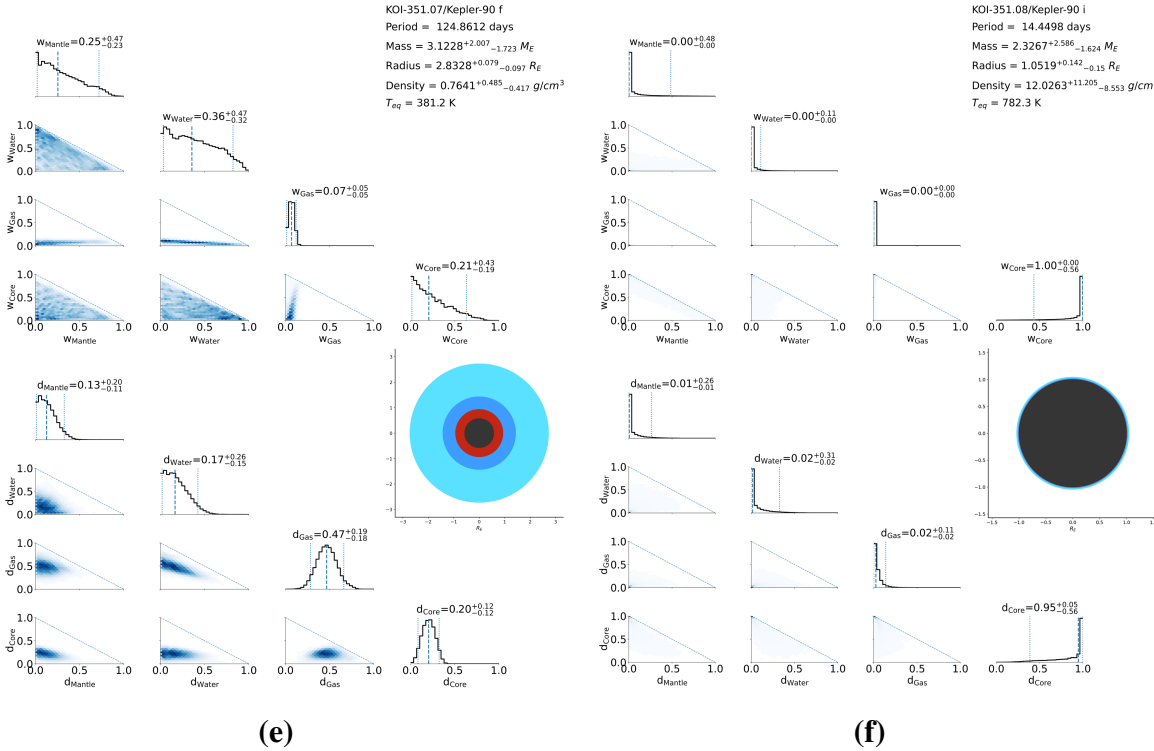
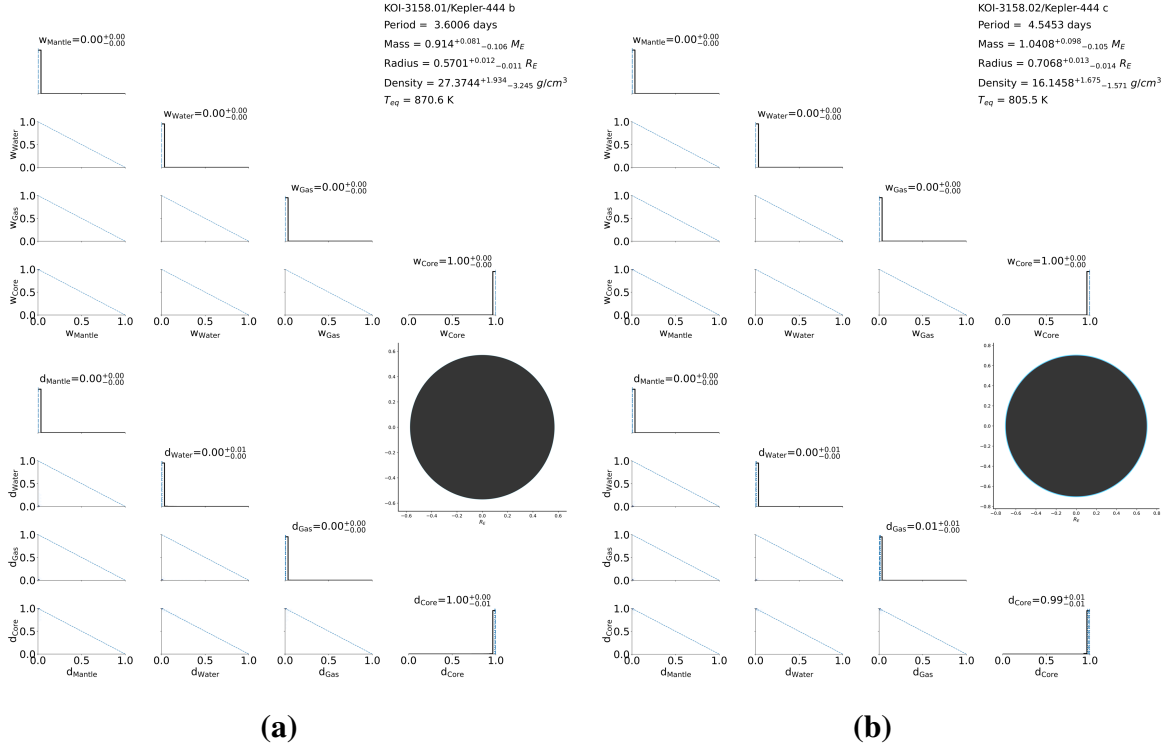


Figure 4.4 Posteriors, starting parameters, and compositional illustration for six of the planets in KOI-351/Kepler-90. Kepler-90 is an eight planet system of considerable interest since it has the largest number of known planets in the Kepler database. Only six of the eight planets are given here as the remaining two had considerable portions of their mass posterior that fell out of the allowed range for ExoMDN. It is interesting to see that between the six planets, we see four unique compositional structures. The first is the very low density planet Kepler-90g which based on the output of ExoMDN would need a composition almost entirely of H/He in order to explain its small mass and large radius, while this could be possible, it is worth noting that Weiss et al. found a much larger mass for this planet which in turn would allow for a larger contribution of heavier elements. Planets d, e, and f are fairly consistent with the class of sub-Neptunes as they have radii near $3R_{\oplus}$, considerable water/ice layers, and a significant amount of their radius is due to their H/He layer. Planet c is more likely a super-Earth given its slightly larger than Earth radius of $1.44R_{\oplus}$, but very small contribution of water by mass and essentially zero H/He. Finally, Kepler-90f is consistent with just containing an iron core that accounts for 100% of its mass. However, within the 1σ error range there is some possibility for a nonzero mantle and water component.



KOI	Period (d)	M_{Earth}	R_{Earth}	$T_{eq}(\text{K})$	m_{Core}	m_{Mantle}	m_{Water}	m_{Gas}	d_{Core}	d_{Mantle}	d_{Water}	d_{Gas}
41.02	6.89	$2.18^{+5.58}_{-0.73}$	$1.37^{+1.41}_{-1.34}$	990.0	$0.8^{+0.98}_{-0.42}$	$0.12^{+0.48}_{-0.01}$	$0.03^{+0.14}_{-0.00}$	$0.0^{+0.0}_{-0.0}$	$0.64^{+0.89}_{-0.36}$	$0.09^{+0.25}_{-0.01}$	$0.12^{+0.25}_{-0.01}$	$0.06^{+0.14}_{-0.03}$
41.03	35.34	$9.74^{+18.31}_{-2.95}$	$1.63^{+1.67}_{-1.58}$	648.05	$0.9^{+1.0}_{-0.53}$	$0.05^{+0.29}_{-0.00}$	$0.02^{+0.13}_{-0.00}$	$0.0^{+0.0}_{-0.0}$	$0.79^{+0.97}_{-0.48}$	$0.06^{+0.2}_{-0.01}$	$0.05^{+0.2}_{-0.01}$	$0.04^{+0.11}_{-0.01}$
46.02	6.03	$4.05^{+7.48}_{-1.38}$	$1.4^{+1.48}_{-1.32}$	990.0	$0.92^{+0.99}_{-0.59}$	$0.05^{+0.31}_{-0.00}$	$0.01^{+0.09}_{-0.00}$	$0.0^{+0.0}_{-0.0}$	$0.77^{+0.94}_{-0.49}$	$0.05^{+0.19}_{-0.01}$	$0.07^{+0.19}_{-0.01}$	$0.06^{+0.12}_{-0.02}$
70.02	3.7	$7.16^{+16.87}_{-2.09}$	$1.72^{+1.76}_{-1.69}$	974.26	$0.83^{+0.98}_{-0.4}$	$0.1^{+0.4}_{-0.01}$	$0.03^{+0.17}_{-0.00}$	$0.0^{+0.0}_{-0.0}$	$0.68^{+0.89}_{-0.39}$	$0.09^{+0.24}_{-0.02}$	$0.09^{+0.24}_{-0.02}$	$0.06^{+0.15}_{-0.03}$
70.05	19.58	$1.22^{+2.02}_{-0.59}$	$0.9^{+0.93}_{-0.87}$	562.37	$0.98^{+1.0}_{-0.78}$	$0.01^{+0.16}_{-0.00}$	$0.0^{+0.04}_{-0.00}$	$0.0^{+0.0}_{-0.0}$	$0.91^{+0.99}_{-0.68}$	$0.02^{+0.12}_{-0.00}$	$0.02^{+0.12}_{-0.00}$	$0.03^{+0.07}_{-0.01}$
72.01	0.84	$9.33^{+15.41}_{-3.32}$	$1.49^{+1.52}_{-1.46}$	990.0	$0.99^{+1.0}_{-0.79}$	$0.01^{+0.15}_{-0.00}$	$0.0^{+0.04}_{-0.00}$	$0.0^{+0.0}_{-0.0}$	$0.92^{+0.99}_{-0.64}$	$0.01^{+0.12}_{-0.00}$	$0.02^{+0.12}_{-0.00}$	$0.03^{+0.09}_{-0.00}$
82.01	16.15	$1.74^{+3.02}_{-0.76}$	$2.28^{+2.32}_{-2.25}$	497.2	$0.2^{+0.45}_{-0.05}$	$0.26^{+0.57}_{-0.08}$	$0.43^{+0.72}_{-0.15}$	$0.0^{+0.01}_{-0.00}$	$0.22^{+0.29}_{-0.14}$	$0.14^{+0.26}_{-0.04}$	$0.22^{+0.26}_{-0.04}$	$0.38^{+0.53}_{-0.24}$
82.02	10.31	$0.92^{+1.34}_{-0.63}$	$1.21^{+1.23}_{-1.19}$	570.89	$0.5^{+0.76}_{-0.2}$	$0.31^{+0.64}_{-0.09}$	$0.12^{+0.29}_{-0.03}$	$0.0^{+0.0}_{-0.0}$	$0.48^{+0.61}_{-0.33}$	$0.19^{+0.39}_{-0.06}$	$0.18^{+0.39}_{-0.06}$	$0.09^{+0.19}_{-0.03}$
82.03	27.45	$1.08^{+1.54}_{-0.64}$	$0.94^{+0.96}_{-0.92}$	411.9	$0.91^{+0.99}_{-0.68}$	$0.06^{+0.24}_{-0.01}$	$0.01^{+0.07}_{-0.00}$	$0.0^{+0.0}_{-0.0}$	$0.82^{+0.93}_{-0.63}$	$0.05^{+0.17}_{-0.01}$	$0.04^{+0.17}_{-0.01}$	$0.05^{+0.1}_{-0.03}$
85.01	5.86	$5.13^{+6.96}_{-3.54}$	$2.69^{+2.76}_{-2.62}$	990.0	$0.22^{+0.44}_{-0.06}$	$0.28^{+0.58}_{-0.08}$	$0.41^{+0.67}_{-0.18}$	$0.0^{+0.0}_{-0.0}$	$0.26^{+0.32}_{-0.17}$	$0.17^{+0.31}_{-0.05}$	$0.29^{+0.31}_{-0.05}$	$0.26^{+0.38}_{-0.15}$
85.02	2.15	$9.92^{+15.45}_{-4.96}$	$1.51^{+1.55}_{-1.46}$	990.0	$0.99^{+1.0}_{-0.88}$	$0.01^{+0.09}_{-0.00}$	$0.0^{+0.02}_{-0.00}$	$0.0^{+0.0}_{-0.0}$	$0.92^{+0.99}_{-0.73}$	$0.01^{+0.09}_{-0.00}$	$0.02^{+0.09}_{-0.00}$	$0.03^{+0.08}_{-0.01}$
85.03	8.13	$3.6^{+4.65}_{-2.65}$	$1.62^{+1.66}_{-1.57}$	990.0	$0.7^{+0.85}_{-0.44}$	$0.22^{+0.49}_{-0.06}$	$0.06^{+0.14}_{-0.02}$	$0.0^{+0.0}_{-0.0}$	$0.58^{+0.66}_{-0.47}$	$0.15^{+0.31}_{-0.05}$	$0.14^{+0.31}_{-0.05}$	$0.09^{+0.17}_{-0.05}$
89.02	207.61	$9.12^{+16.45}_{-3.49}$	$3.92^{+4.04}_{-3.83}$	472.52	$0.2^{+0.42}_{-0.05}$	$0.24^{+0.49}_{-0.07}$	$0.3^{+0.58}_{-0.08}$	$0.17^{+0.22}_{-0.12}$	$0.19^{+0.25}_{-0.12}$	$0.12^{+0.23}_{-0.04}$	$0.14^{+0.23}_{-0.04}$	$0.52^{+0.64}_{-0.42}$
94.02	10.42	$13.52^{+16.51}_{-11.26}$	$3.89^{+3.94}_{-3.84}$	990.0	$0.19^{+0.45}_{-0.05}$	$0.25^{+0.53}_{-0.07}$	$0.38^{+0.7}_{-0.11}$	$0.07^{+0.1}_{-0.04}$	$0.21^{+0.28}_{-0.14}$	$0.14^{+0.26}_{-0.05}$	$0.2^{+0.26}_{-0.05}$	$0.42^{+0.5}_{-0.35}$
94.03	54.33	$16.22^{+18.32}_{-13.67}$	$6.28^{+6.36}_{-6.2}$	593.38	$0.12^{+0.25}_{-0.03}$	$0.14^{+0.29}_{-0.04}$	$0.16^{+0.33}_{-0.04}$	$0.52^{+0.56}_{-0.46}$	$0.12^{+0.16}_{-0.08}$	$0.08^{+0.14}_{-0.03}$	$0.08^{+0.14}_{-0.03}$	$0.71^{+0.74}_{-0.67}$
94.04	3.74	$11.47^{+18.45}_{-3.73}$	$1.57^{+1.6}_{-1.55}$	990.0	$0.98^{+1.0}_{-0.75}$	$0.01^{+0.17}_{-0.00}$	$0.0^{+0.05}_{-0.00}$	$0.0^{+0.0}_{-0.0}$	$0.9^{+0.99}_{-0.61}$	$0.02^{+0.13}_{-0.00}$	$0.02^{+0.13}_{-0.00}$	$0.03^{+0.09}_{-0.01}$
102.02	4.07	$2.52^{+4.97}_{-0.66}$	$1.03^{+1.07}_{-0.99}$	110.0	$0.99^{+1.0}_{-0.44}$	$0.01^{+0.28}_{-0.00}$	$0.0^{+0.16}_{-0.00}$	$0.0^{+0.0}_{-0.0}$	$0.96^{+1.0}_{-0.48}$	$0.01^{+0.2}_{-0.00}$	$0.01^{+0.2}_{-0.00}$	$0.01^{+0.06}_{-0.00}$
108.02	179.6	$5.62^{+13.53}_{-1.7}$	$5.49^{+5.63}_{-5.35}$	371.49	$0.12^{+0.24}_{-0.03}$	$0.13^{+0.27}_{-0.04}$	$0.15^{+0.31}_{-0.04}$	$0.57^{+0.64}_{-0.41}$	$0.1^{+0.15}_{-0.07}$	$0.06^{+0.12}_{-0.02}$	$0.06^{+0.12}_{-0.02}$	$0.76^{+0.81}_{-0.68}$

Table 4.1 Planetary parameters and 1σ error ranges for model components using the hot dataset where the equilibrium temperature was set using Equation 4.1.

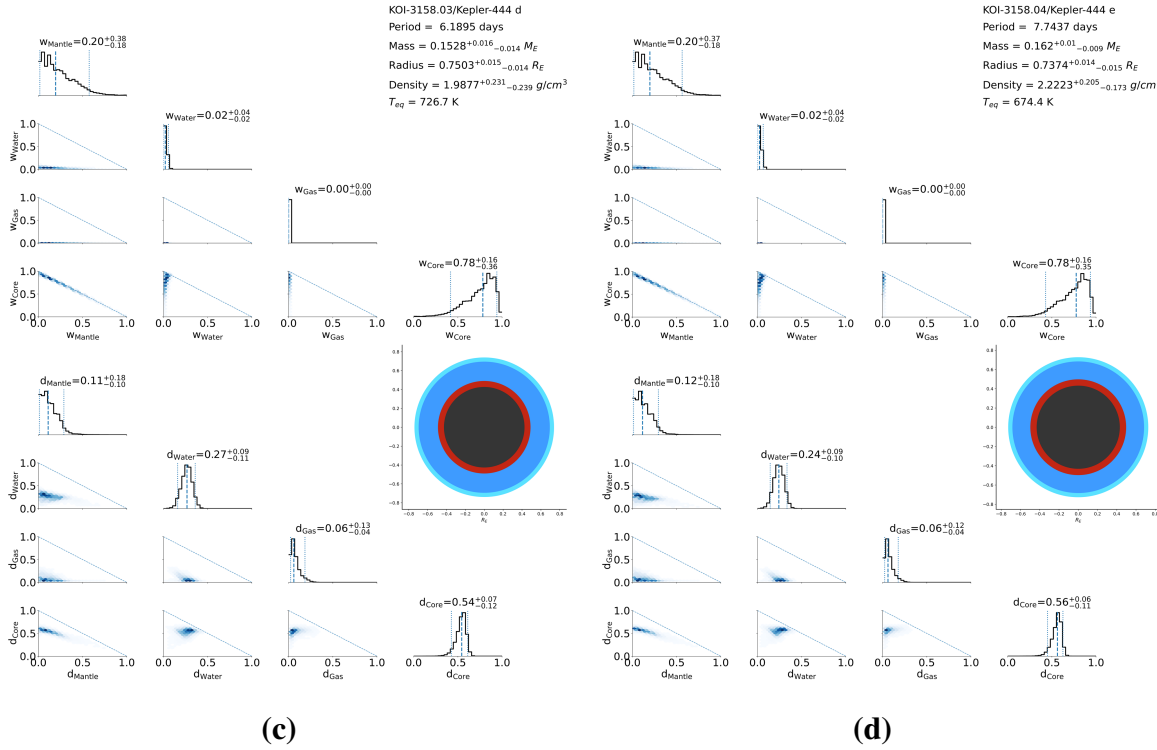


Figure 4.5 Posteriors, input parameters and relevant information, and theoretical composition illustration for four of the five planets in KOI-3158/Kepler-444. What makes Kepler-444 interesting enough to highlight here is that the planets had some of the tightest mass constraints out of all planets in the KMDC and therefore ExoMDN had the best chance of obtaining reliable interior models for these planets. What is seen is that despite the planets all having very similar radii, there are two distinct compositional structures. The first is for planets b and c which have masses near $1M_{\oplus}$ and radii between 0.57 and $0.7 R_{\oplus}$. This combination results in the iron ball configuration where the core mass fraction is required to be 100%. The second structure found for planets d and e is significantly less extreme where the core only accounts for 78% of the mass on average and there is also the possibility for nonzero components of a mantle and water layer. It is also worth noting that because the starting parameters for planets d and e are so close to each other, the models constructed by ExoMDN are virtually identical.

KOI	Period (d)	M_{Earth}	R_{Earth}	$T_{eq}(K)$	m_{Core}	m_{Mantle}	m_{Water}	m_{Gas}	d_{Core}	d_{Mantle}	d_{Water}	d_{Gas}
41.02	6.89	$2.18^{+5.58}_{-0.73}$	$1.37^{+1.41}_{-1.34}$	150.0	$0.51^{+0.91}_{-0.13}$	$0.19^{+0.52}_{-0.03}$	$0.14^{+0.5}_{-0.01}$	$0.0^{+0.0}_{-0.0}$	$0.52^{+0.84}_{-0.26}$	$0.15^{+0.35}_{-0.03}$	$0.15^{+0.35}_{-0.03}$	$0.04^{+0.2}_{-0.01}$
41.03	35.34	$9.74^{+18.31}_{-2.95}$	$1.63^{+1.67}_{-1.58}$	150.0	$0.84^{+0.99}_{-0.31}$	$0.08^{+0.35}_{-0.01}$	$0.04^{+0.25}_{-0.0}$	$0.0^{+0.0}_{-0.0}$	$0.77^{+0.96}_{-0.39}$	$0.08^{+0.26}_{-0.01}$	$0.07^{+0.26}_{-0.01}$	$0.01^{+0.06}_{-0.01}$
46.02	6.03	$4.05^{+7.48}_{-1.38}$	$1.4^{+1.48}_{-1.32}$	150.0	$0.74^{+0.97}_{-0.25}$	$0.13^{+0.42}_{-0.01}$	$0.06^{+0.32}_{-0.0}$	$0.0^{+0.0}_{-0.0}$	$0.7^{+0.92}_{-0.34}$	$0.11^{+0.3}_{-0.02}$	$0.1^{+0.3}_{-0.02}$	$0.02^{+0.09}_{-0.01}$
70.02	3.7	$7.16^{+16.87}_{-2.09}$	$1.72^{+1.76}_{-1.69}$	150.0	$0.63^{+0.92}_{-0.16}$	$0.15^{+0.46}_{-0.03}$	$0.1^{+0.45}_{-0.01}$	$0.0^{+0.01}_{-0.0}$	$0.6^{+0.85}_{-0.28}$	$0.14^{+0.33}_{-0.04}$	$0.13^{+0.33}_{-0.04}$	$0.02^{+0.13}_{-0.01}$
70.05	19.58	$1.22^{+2.02}_{-0.59}$	$0.9^{+0.93}_{-0.87}$	150.0	$0.95^{+1.0}_{-0.61}$	$0.03^{+0.25}_{-0.0}$	$0.01^{+0.1}_{-0.0}$	$0.0^{+0.0}_{-0.0}$	$0.89^{+0.99}_{-0.61}$	$0.03^{+0.19}_{-0.0}$	$0.02^{+0.19}_{-0.0}$	$0.02^{+0.06}_{-0.01}$
72.01	0.84	$9.33^{+15.41}_{-3.32}$	$1.49^{+1.52}_{-1.46}$	150.0	$0.94^{+1.0}_{-0.51}$	$0.03^{+0.26}_{-0.0}$	$0.01^{+0.14}_{-0.0}$	$0.0^{+0.0}_{-0.0}$	$0.88^{+0.99}_{-0.54}$	$0.04^{+0.21}_{-0.0}$	$0.03^{+0.21}_{-0.0}$	$0.01^{+0.04}_{-0.0}$
82.01	16.15	$1.74^{+3.02}_{-0.76}$	$2.28^{+2.32}_{-2.25}$	150.0	$0.2^{+0.45}_{-0.05}$	$0.24^{+0.51}_{-0.07}$	$0.33^{+0.64}_{-0.09}$	$0.12^{+0.17}_{-0.06}$	$0.22^{+0.3}_{-0.14}$	$0.13^{+0.26}_{-0.04}$	$0.16^{+0.26}_{-0.04}$	$0.45^{+0.58}_{-0.31}$
82.02	10.31	$0.92^{+1.34}_{-0.63}$	$1.21^{+1.23}_{-1.19}$	150.0	$0.29^{+0.55}_{-0.09}$	$0.35^{+0.65}_{-0.1}$	$0.27^{+0.53}_{-0.1}$	$0.0^{+0.0}_{-0.0}$	$0.4^{+0.54}_{-0.26}$	$0.24^{+0.44}_{-0.08}$	$0.23^{+0.44}_{-0.08}$	$0.07^{+0.16}_{-0.02}$
82.03	27.45	$1.08^{+1.54}_{-0.64}$	$0.94^{+0.96}_{-0.92}$	150.0	$0.85^{+0.97}_{-0.56}$	$0.09^{+0.32}_{-0.01}$	$0.03^{+0.12}_{-0.0}$	$0.0^{+0.0}_{-0.0}$	$0.79^{+0.92}_{-0.59}$	$0.08^{+0.23}_{-0.02}$	$0.05^{+0.23}_{-0.02}$	$0.03^{+0.07}_{-0.01}$
85.01	5.86	$5.13^{+6.96}_{-3.54}$	$2.69^{+2.76}_{-2.62}$	150.0	$0.19^{+0.44}_{-0.05}$	$0.24^{+0.51}_{-0.07}$	$0.34^{+0.65}_{-0.09}$	$0.11^{+0.17}_{-0.06}$	$0.25^{+0.32}_{-0.16}$	$0.16^{+0.29}_{-0.05}$	$0.19^{+0.29}_{-0.05}$	$0.37^{+0.48}_{-0.26}$
85.02	2.15	$9.92^{+15.45}_{-4.96}$	$1.51^{+1.55}_{-1.46}$	150.0	$0.94^{+1.0}_{-0.69}$	$0.03^{+0.2}_{-0.0}$	$0.01^{+0.09}_{-0.0}$	$0.0^{+0.0}_{-0.0}$	$0.88^{+0.99}_{-0.66}$	$0.04^{+0.18}_{-0.0}$	$0.03^{+0.18}_{-0.0}$	$0.01^{+0.03}_{-0.0}$
85.03	8.13	$3.6^{+4.65}_{-2.65}$	$1.62^{+1.66}_{-1.57}$	150.0	$0.36^{+0.58}_{-0.13}$	$0.36^{+0.67}_{-0.11}$	$0.22^{+0.42}_{-0.08}$	$0.0^{+0.0}_{-0.0}$	$0.46^{+0.57}_{-0.31}$	$0.27^{+0.47}_{-0.09}$	$0.21^{+0.47}_{-0.09}$	$0.03^{+0.07}_{-0.01}$
89.02	207.61	$9.12^{+16.45}_{-3.49}$	$3.92^{+4.04}_{-3.83}$	150.0	$0.15^{+0.33}_{-0.04}$	$0.17^{+0.37}_{-0.05}$	$0.21^{+0.44}_{-0.05}$	$0.34^{+0.54}_{-0.23}$	$0.18^{+0.25}_{-0.1}$	$0.1^{+0.21}_{-0.03}$	$0.11^{+0.21}_{-0.03}$	$0.56^{+0.7}_{-0.45}$
94.02	10.42	$13.52^{+16.51}_{-11.26}$	$3.89^{+3.94}_{-3.84}$	150.0	$0.18^{+0.38}_{-0.05}$	$0.21^{+0.43}_{-0.06}$	$0.26^{+0.52}_{-0.07}$	$0.26^{+0.31}_{-0.2}$	$0.21^{+0.27}_{-0.14}$	$0.13^{+0.25}_{-0.04}$	$0.15^{+0.25}_{-0.04}$	$0.49^{+0.56}_{-0.41}$
94.03	54.33	$16.22^{+18.32}_{-13.67}$	$6.28^{+6.36}_{-6.2}$	150.0	$0.02^{+0.04}_{-0.01}$	$0.02^{+0.04}_{-0.0}$	$0.02^{+0.05}_{-0.01}$	$0.93^{+0.96}_{-0.9}$	$0.08^{+0.1}_{-0.05}$	$0.04^{+0.07}_{-0.01}$	$0.04^{+0.07}_{-0.01}$	$0.83^{+0.86}_{-0.8}$
94.04	3.74	$11.47^{+18.45}_{-3.73}$	$1.57^{+1.6}_{-1.55}$	150.0	$0.93^{+1.0}_{-0.46}$	$0.04^{+0.27}_{-0.0}$	$0.01^{+0.16}_{-0.0}$	$0.0^{+0.0}_{-0.0}$	$0.86^{+0.98}_{-0.5}$	$0.05^{+0.21}_{-0.0}$	$0.04^{+0.21}_{-0.0}$	$0.01^{+0.04}_{-0.0}$
102.02	4.07	$2.52^{+4.97}_{-0.66}$	$1.03^{+1.07}_{-0.99}$	150.0	$0.99^{+1.0}_{-0.47}$	$0.01^{+0.27}_{-0.0}$	$0.0^{+0.14}_{-0.0}$	$0.0^{+0.0}_{-0.0}$	$0.96^{+1.0}_{-0.49}$	$0.01^{+0.19}_{-0.0}$	$0.01^{+0.19}_{-0.0}$	$0.01^{+0.07}_{-0.0}$
108.02	179.6	$5.62^{+13.53}_{-1.7}$	$5.49^{+5.63}_{-3.55}$	150.0	$0.03^{+0.08}_{-0.01}$	$0.03^{+0.08}_{-0.01}$	$0.03^{+0.1}_{-0.01}$	$0.9^{+0.95}_{-0.73}$	$0.07^{+0.12}_{-0.04}$	$0.04^{+0.08}_{-0.01}$	$0.04^{+0.08}_{-0.01}$	$0.84^{+0.9}_{-0.74}$

Table 4.2 Planetary parameters and 1σ error ranges for model components using the cold dataset where the equilibrium temperature was fixed at 150K. In general this led to an enhancement in the mass fraction of volatiles as their effect on the planet’s radius was mitigated due to the lower temperatures.

4.4 Discussion

Given the general high uncertainty in the mass and radius fractions shown in the posterior distributions for KOIs-277, 351, and 3158 as well as in Tables 4.1 and 4.2 it is difficult to draw definitive conclusions about the compositional structure of most of these planets. Combining this with the general lack of preexisting work that would allow for a side-by-side comparison of our results, there is not much more that can be said at this time on the level of individual planets.

However, it is still possible to look at our catalog of results and determine how well they agree with larger scale theories about the Kepler systems. Specifically, we look at what our interior models can say about the Kepler Radius Gap. As explained in Ho & Van Eylen (2023) is a well-established finding that within the Kepler dataset there is an under-density of planets with radii between $1.5-2 R_{\oplus}$ and periods less than 100 days. The current theory is that essentially all of the smaller rocky planets, super-Earths, and sub-Neptunes in the Kepler dataset have very similar compositions, but

the only major difference is the amount of gas the planets contain in their atmospheres. The small planets below the radius gap are expected to be rocky cores that have lost their atmosphere via either core powered mass loss or photoevaporation. With this in mind the idea then becomes planets below the gap do not have meaningful amounts of gas while those above the gap do have a significant gas component.

To begin our exploration of this theory, we have Figure 4.6. This figure shows all of our modeled planets plotted in period-radius space with points colored by their 1σ upper limits on their mass fraction of H/He (w_{Gas}) for both our hot and cold datasets. Already by eye, the radius gap is quite clear as it causes the points to separate into two distinct clusters. For both of our datasets the points below the gap have almost no or absolutely no H/He at the 1σ level. Above the gap the story is slightly different between the datasets. For the hot set, there are still a considerable number of planets that have no significant gas component until they start having radii above $5R_{\oplus}$ whereas for the cold dataset this transition occurs at a considerably lower radius value.

In order to see exactly where a theoretical split between no gas and a meaningful amount of gas would occur, we employed a support vector machine (SVM) classifier. To initialize the SVM we first needed to label all of the points/planets as either rocky or gaseous. Because there is no real defined value at which an atmospheric component becomes significant, we look at three cutoff values which are $w_{Gas} = 0.02, 0.05, 0.1$ and ran the SVM three times. For each iteration, planets with a 1σ upper limit on the mass fraction of their atmosphere that was less than the cutoff were labeled as rocky and those above were labeled gaseous. Then the SVM algorithm proceeded to find the line that maximized the difference between the boundary and the points chosen as support vectors per the standard SVM routine. The results of these analyses are also shown in Figure 4.6 with each labeled line indicating the cutoff value that was used to define the boundary. Unsurprisingly, the SVM boundaries for the hot dataset occur at higher radius values. Focusing just on the $w_{Gas} = 0.02$ boundary provides further evidence of how little gas the majority of the planets have since nearly

every planet below the boundary have less than 2% gas by mass. It is also interesting to see that in both the hot and cold datasets, the nominal division between the two clusters is overshoot by the SVM with the closest division coming from the cold dataset when the mass fraction boundary is 0.02. Additionally, we plot the boundary line found by Ho & Van Eylen to see how our results compare to what they found. In order to obtain their line, Ho & Van Eylen first updated the planetary parameters for 431 planets by making use of Kepler's short cadence data. With the parameters set, they proceeded to use a Gaussian Mixture Model to classify their data into two groups which could then be separated using an SVM. The boundary line found by Ho & Van Eylen performs significantly better in matching the radius gap for the hot dataset. In the case of the cold dataset, it does not necessarily perform better than the 0.02 mass fraction line, but does differ in the fact that it has a more significant negative slope, leading to a small difference in classifications.

While it is encouraging that the cold dataset allows for a definition of the radius gap that roughly agrees with other findings, it is not satisfying because the cold dataset is much more an experimental investigation into the effect that temperature has on compositional models than a dataset that is believed to accurately represent the Kepler planets. To further examine the radius gap, in Figure 4.7 we present the same data as in Figure 4.6 except that this time the points are colored by the radius fraction of the gas instead of their mass fraction. By doing this, three results immediately become clear. First, the difference between the cold and hot dataset almost vanishes as is evident by both the point colors and the location of the SVM boundaries. These boundaries were obtained in the same manner except the cutoffs were based off having a radius fraction at or above 0.05, 0.1, and 0.2. The boundaries in the cold dataset (right) are slightly more compressed with each other but the overall placement of the boundary and slope of the lines is consistent with those found in the hot dataset (left). The second result is that the boundaries agree much better with where a visual inspection would define the radius gap. Finally, given that these boundary lines agree better with the radius gap as a whole, they also agree considerably better with the boundary found by Ho &

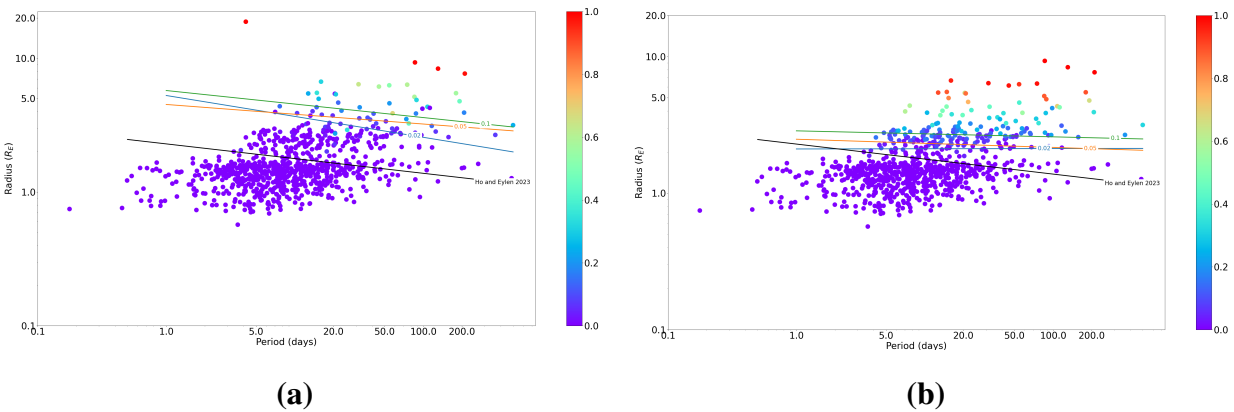


Figure 4.6 Period-radius plot for the 830 planets modeled in this work. The points are colored by the atmospheric mass fraction w_{Gas} . Subplot a shows the results for the hot dataset while subplot b shows the same for the cold dataset. In both figures the Kepler Radius Gap is visible as an under density of points/planets. We attempt to model this division in the Kepler data using an SVM classifier. The two classes are either rocky or gaseous with the distinction made in the 1σ upper limit on each planet’s H/He mass fraction. Three potential boundaries are found based off of what cutoff (.02, 0.05, 0.1) was used to divide rocky from gaseous planets. Along with this, we plot the boundary line that was found by Ho & Van Eymen to describe the Kepler gap. For the hot dataset (subplot a) we find that the boundary occurs at much too high of radii in all three cases. Alternatively for the cold dataset, the boundary occurs at much lower radii, but still does not provide an acceptable match to the one given by Ho & Van Eymen especially when factoring the cold dataset does not accurately model the equilibrium temperatures for these planets.

Van Eylen. Ho & Van Eylen’s boundary line still has a larger slope than the one found by our SVM method, but given the low density of points in this region of parameter space, it only changes the classification for a small (20-30) number of points/planets out of the 830 in the figures. Although it is not the point of this work to say which boundary is more accurate, it is worth noting that the one found in this work was obtained by first classifying the planets based off of a physical characteristic (either mass or radius fraction of H/He) whereas the one given in Ho & Van Eylen was found by simply applying a clustering algorithm and labeling points based off of which cluster they were assigned to and then in both cases the SVM was used to find the boundary. While this is a small difference in methods, and certainly the clustering algorithm was picking up a feature of the data caused by a physical difference in the planets, our method does benefit from the fact that it was more solidly rooted in the physical characteristics of the planets.

Some additional attention is needed to address the difference between Figures 4.6 and 4.7. When looking at mass fractions, there is a clear difference between the hot and cold datasets which is not visible when looking at the radius fractions. This is clear evidence of how temperature can affect a planet’s radius if it has any amount of atmosphere. The planets in the hot and cold datasets have the same radii and, as seen in Figure 4.7 they also have essentially the same radius fractions. Thus, due to stellar heating, the same inflation of a planet’s radius can be achieved with a very minimal fraction of H/He. This leads to a couple of implications. First, the planets above the radius gap that have sizes in the super-Earth to sub-Neptune range may have considerably less gas by mass than would have nominally been expected. As seen many planets above the gap had a mass fraction less than 0.2 and it must also be reemphasized that that was their 1σ upper limit on the mass fraction. Next, it really is possible that the planets in this $R < 5R_{\oplus}$ regime have nearly identical compositions, and it is only a slight enhancement of gas that makes them appear considerably different in radius space. Finally, it would not require much mass loss via photoevaporation or core powered to make a planet drop from above to below the radius gap.




Figure 4.7 The same plots as Figure 4.6 except that now points are colored based off of the radius fraction of the H/He layer (d_{Gas}) and the classes for the SVM boundary are defined as either rocky or gaseous if the radius fraction is greater than 0.05, 0.1, and 0.2. This figure immediately differs from Figure 4.6 in two important ways. First, the difference between the hot and cold datasets are almost completely gone. This is because while the hot dataset has lower mass fractions of gas, the contribution to the overall radius is the same. Second, the level of agreement between the boundary lines we found and the one given by Ho & Van Eylon is significantly greater with the best being between their line and our 0.05 radius fraction line. Our boundary lines have a considerably less negative slope, but due to the under density of points in this region, it does not lead to a difference in classification of many points.

4.5 Conclusion

The KMDC is currently the largest homogeneously derived catalog of exoplanetary parameters. Use of this catalog makes it possible to go from analyzing exoplanets on a system/cluster of systems level to a scale that includes nearly all of the Kepler multis. In this work we applied the catalog to generate interior models of 830 planets using the software tool ExoMND. This modeling software allowed us to construct Bayesian posterior distributions for a four-layer planetary compositional structure consisting of an iron core, silicate mantle, water/ice, and a H/He atmosphere. These results are presented in this work in part with the full data files and posterior plots being available upon request while the data tables giving the 1σ error ranges on the model parameters are included as supplemental material in a machine readable format.

After modeling the 830 planets, several findings became quite clear:

1. The KMDC contained many planets that could only be modeled as a 100% ball of solid iron. These planets are unlikely to be physical and are likely a result of missing features in the photodynamical model that generated the KMDC or starting parameters that need to be updated. For example, if stellar dilution is not properly modeled or the stellar radius is off, then the radius of the planet is underestimated and the result is that the density is overestimated.
2. The combined posterior distributions for the model parameters all have a peak at a mass fraction of 0 indicating that it would be possible to build a consistent model with the input parameters that had at most three layers instead of four. This speaks to the inherent degeneracy of the problem being investigated here. From a straight modeling point of view, it is just as possible for a planet to have a small iron core and large gaseous atmosphere as it is for it to have a moderate mantle and water layer. Both of these solutions can return a planet that matches the input parameters, and thus it is a difficult task to determine which is more correct.

3. Unsurprisingly, when only three parameters are used as input for ExoMDN, all three parameters are important and significant changes in one will lead to significant differences in model compositions. This was seen when comparing the differences between our hot dataset that used a physically motivated calculation of the equilibrium temperature and the cold dataset where the temperature was set to 150K for every modeled planet. When modeled with a lower temperature, it was possible for the planets to have higher fractions of volatiles because they were not inflating the radius as they would at higher temperatures.
4. Given the high equilibrium temperatures that many of the Kepler planets have as a result of orbiting close to their host star, a very small amount of H/He by mass (<1%) can account for approximately 30% or more of their radius.
5. By plotting the modeled planets in period-radius space, the Kepler Radius Gap is easily seen as an under-density of data points. Using an SVM and labeling the points based off of their gas mass fraction, it is difficult to find a boundary line that separates the two groups as expected regardless of whether the cold or hot dataset is used. However, when switching to labeling the points by the fraction of the radius which the gas contributes, an SVM classifier is able to match expected results, and the findings are consistent between the hot and cold datasets.

With these in mind, the biggest question that remains is the overall level of uncertainty in the planet models. The first source of this uncertainty is the starting parameters themselves. The KMDC was meant to be a large scale endeavor to solve for the parameters of all the Kepler multis, but did this at the expense of not necessarily being about to obtain the tightest parameter constraints for each individual planet. If this catalog was used in combination with other results or future results that improved upon those given in the KMDC, it may be possible to provide tighter compositional constraints. Additionally, as stated previously, the albedo used when calculating the equilibrium

temperatures was set at 0.5 for all planets, which is not expected to be the best value across the board. If some additional work were to be done to find a suitable way to more accurately set this parameter, it would help to improve the models found and published here. Beyond improving the starting parameters, there are also enhancements that could be made in ExoMDN itself. For example, Baumeister & Tosi specifically mention that they used a simplified atmospheric model, as their main interest was in the machine learning capabilities of ExoMDN more so than in capturing all the correct physics. It is therefore reasonable to predict that if more advanced models were incorporated into ExoMDN, the models would benefit at least from being derived from a more accurate representation of exoplanetary atmospheres. Finally, it is very likely the case that the three parameter inputs used are not sufficient to constrain the four parameter compositional model. Baumeister & Tosi were able to significantly improve their solutions and reduce their uncertainties in the models when the fluid Love number k_2 was included. Once included, their models for Earth were accurately constrained to the point that core composition needed to be included beyond the default pure iron used in ExoMDN. While Love numbers for the Kepler planets are highly unknown, if in the future there were a method for determining them, their incorporation into interior models would significantly decrease the compositional uncertainty. In the meantime, it is also plausible that other more readily obtainable parameters, such as stellar parameters, could be used to constrain the models beyond their current state.

Another challenge facing this work as a whole is the lack of context in which to place it. Determining accurate masses for the Kepler planets is still an active area of research, and modeling their interiors is one step down the line. With the exception of Baumeister & Tosi and there are not many results that we could compare on a large scale. On smaller scales, analyses have been performed such as what has been found for KOI-277 by Carter et al., but it would require an investigation beyond the scope of this work to compare all the compositional models done on single planets/systems with the results presented here. The goal, however, is that between the KMDC and

the results given in this paper, a significant step forward can be taken towards understanding the compositional structures of the Kepler multis.

4.6 Future Work

With such a range of exoplanetary interior compositions, there are many potential avenues for additional research. A more detailed exploration of the occurrence-weighted compositions is likely to reveal information about the true composition distribution of exoplanets for the first time. A Hierarchical Bayesian Model can also explore population level assumptions. For example, are the data consistent with $w_{\text{Water}} = 0$ for all Kepler systems? That is, can the data be (mostly) explained by rocky-iron cores with varying gas fractions? This is a hypothesis that has been proposed in the literature and is essential for assessing whether Kepler-like planets formed *in situ* (e.g., inside the ice-line and thus accreted no water) or formed further out (with significant components of water) and migrated inwards. Another commonly proposed hypothesis that could be tested with Hierarchical Modeling is that the rock to iron ratio for small planets is similar to Earth's or potentially related to the [Si]/[Fe] ratio for the stars.

Another avenue of exploration is how compositions are correlated (or not) within individual systems. One of the cardinal results of solar system compositions is that the interior terrestrial planets have much less water and gas than exterior giant planets. Are there similar compositional trends within Kepler-like systems?

Overall, the goal of future work is to reduce the importance of the significant degeneracies found in individual planets by exploring the ensemble of planets either within planetary systems or across all Kepler discoveries. With hundreds of homogeneously-analyzed planets, the dataset provided here provides an excellent start to these and many other questions.

4.7 Acknowledgments

This material is based upon work supported by the National Science Foundation under Award No. 2143195. In addition, we would like to recognize the contributions of Raymond Kelly, who is currently a undergraduate researcher, who assisted in generating some of the plots and figures presented here. Finally, we acknowledge the Brigham Young University Office of Research Computing, the supercomputing resources available through this department allowed for a significant enhancement in data generation and analysis speed.

Bibliography

Agol, E., Steffen, J., Sari, R., & Clarkson, W. 2005, MNRAS, 359, 567

Armitage, P. J. 2024, arXiv e-prints, arXiv:2412.11064

Armstrong, D. J., et al. 2025, mnras, 537, 3175

Baumeister, P., & Tosi, N. 2024, in European Planetary Science Congress, EPSC2024–493

Berger, T. A., Huber, D., van Saders, J. L., Gaidos, E., Tayar, J., & Kraus, A. L. 2020, VizieR
Online Data Catalog: Gaia-Kepler stellar properties catalog.I. KIC stars (Berger+, 2020), VizieR
On-line Data Catalog: J/AJ/159/280. Originally published in: 2020AJ....159..280B

Bishop, C. 1994

Borsato, L., et al. 2019, MNRAS, 484, 3233

Brinkman, C. L., et al. 2024, arXiv e-prints, arXiv:2410.00213

Buchhave, L. A., et al. 2016, AJ, 152, 160

Cabrera, J., et al. 2014, ApJ, 781, 18

Carter, J. A., et al. 2012, Science, 337, 556

Castro-González, A., et al. 2024, mr-plotter: Mass-radius diagrams plotter, Astrophysics Source
Code Library, record ascl:2412.027

- Catanzarite, J., & Shao, M. 2011, *ApJ*, 738, 151
- Cochran, W. D., et al. 2011, *ApJS*, 197, 7
- Coughlin, J. L., et al. 2016, *ApJS*, 224, 12
- Cranmer, M., Tamayo, D., Rein, H., Battaglia, P., Hadden, S., Armitage, P. J., Ho, S., & Spergel, D. N. 2021, *Proceedings of the National Academy of Science*, 118, e2026053118
- Dainese, S., & Albrecht, S. H. 2025, arXiv e-prints, arXiv:2503.02451
- Deck, K., Holman, M. J., Agol, E., Carter, J. A., Payne, M. J., Winn, J. N., Lissauer, J. J., & Ragozzine, D. 2013, in *AAS/Division of Dynamical Astronomy Meeting*, Vol. 44, *AAS/Division of Dynamical Astronomy Meeting #44*, 100.02
- Egger, J. A., Alibert, Y., Haldemann, J., & Venturini, J. 2022, in *European Planetary Science Congress*, EPSC2022–320
- Fabrycky, D. C., et al. 2014, *ApJ*, 790, 146
- Foreman-Mackey, D. 2017, *Autocorrelation Time Estimation*
- Fressin, F., et al. 2011, *ApJS*, 197, 5
- Gautier, III, T. N., et al. 2012, *ApJ*, 749, 15
- Gladman, B. 1993, *Icarus*, 106, 247
- Goldstein, H., Poole, C., & Safko, J. 2002, *Classical mechanics*
- Graham, A., & Ragozzine, D. in prep.
- Hadden, S., & Lithwick, Y. 2017, *AJ*, 154, 5

—. 2018, *AJ*, 156, 95

He, M. Y., Ford, E. B., Ragozzine, D., & Carrera, D. 2020, *AJ*, 160, 276

Ho, C. S. K., & Van Eylen, V. 2023, *MNRAS*, 519, 4056

Holczer, T., et al. 2016, *VizieR Online Data Catalog: Kepler TTVs. IX. The full long-cadence data set (Holczer+, 2016), VizieR On-line Data Catalog: J/ApJS/225/9. Originally published in: 2016ApJS..225....9H*

Holman, M. J., & Murray, N. W. 2005, *Science*, 307, 1288

Holman, M. J., et al. 2010, *Science*, 330, 51

Hsu, D. C., Ford, E. B., Ragozzine, D., & Ashby, K. 2019, *AJ*, 158, 109

Huang, C., Rice, D. R., & Steffen, J. H. 2022, *MNRAS*, 513, 5256

Hussain, N., & Tamayo, D. 2020, *MNRAS*, 491, 5258

Jones, D., Ragozzine, D., & Fabrycky, D. in prep.

Jontof-Hutter, D., Dalba, P. A., & Livingston, J. H. 2022, *AJ*, 164, 42

Jontof-Hutter, D., Wolfgang, A., Ford, E. B., Lissauer, J. J., Fabrycky, D. C., & Rowe, J. F. 2021, *AJ*, 161, 246

Judkovsky, Y., Ofir, A., & Aharonson, O. 2022, *AJ*, 163, 91

—. 2024, *AJ*, 167, 103

Kane, M., Ragozzine, D., Flowers, X., Holczer, T., Mazeh, T., & Relles, H. M. 2019, *AJ*, 157, 171

Langford, Z., & Agol, E. 2024, arXiv e-prints, arXiv:2410.03874

- Laskar, J. 1997, *A&A*, 317, L75
- Leleu, A., Chatel, G., Udry, S., Alibert, Y., Delisle, J. B., & Mardling, R. 2021, *A&A*, 655, A66
- Leleu, A., et al. 2023, *A&A*, 669, A117
- Lissauer, J. J., Dawson, R. I., & Tremaine, S. 2014a, *Nature*, 513, 336
- . 2014b, *Nature*, 513, 336
- Lissauer, J. J., & Gavino, S. 2021, *Icarus*, 364, 114470
- Lissauer, J. J., Rowe, J. F., Jontof-Hutter, D., Fabrycky, D. C., Ford, E. B., Ragozzine, D., Steffen, J. H., & Nizam, K. M. 2024, *psj*, 5, 152
- Lissauer, J. J., et al. 2011, *ApJS*, 197, 8
- . 2013, *ApJ*, 770, 131
- Lithwick, Y., & Wu, Y. 2012, *ApJ*, 756, L11
- Lithwick, Y., Xie, J., & Wu, Y. 2012, *ApJ*, 761, 122
- Luque, R., & Pallé, E. 2022, *Science*, 377, 1211
- MacDonald, M. G., Shakespeare, C. J., & Ragozzine, D. 2021, *AJ*, 162, 114
- MacDonald, M. G., et al. 2016, *AJ*, 152, 105
- Mandel, K., & Agol, E. 2002, *ApJ*, 580, L171
- Marchal, C., & Bozis, G. 1982, *Celestial Mechanics*, 26, 311
- Marcy, G. W., et al. 2014, *ApJS*, 210, 20
- Mayor, M., et al. 1995, *IAU Circ.*, 6251, 1

- Millholland, S. C., & Winn, J. N. 2021, *ApJ*, 920, L34
- Mills, S. M., & Fabrycky, D. C. 2017, *ApJ*, 838, L11
- Mills, S. M., Fabrycky, D. C., Migaszewski, C., Ford, E. B., Petigura, E., & Isaacson, H. 2016, *Nature*, 533, 509
- Miralda-Escudé, J. 2002, *ApJ*, 564, 1019
- Murray, C. D., & Dermott, S. F. 1999, *Solar System Dynamics*
- NASA. 2024, NASA Exoplanet Archive
- Neil, A. R., & Rogers, L. A. 2020, *ApJ*, 891, 12
- Obertas, A., Van Laerhoven, C., & Tamayo, D. 2017, *Icarus*, 293, 52
- Ofir, A., Yoffe, G., & Aharonson, O. 2025, *AJ*, 169, 90
- Pál, A. 2012, *MNRAS*, 420, 1630
- Parviainen, H., & Aigrain, S. 2015, *MNRAS*, 453, 3821
- Parviainen, H., Luque, R., & Palle, E. 2024, *MNRAS*, 527, 5693
- Payne, M. J., Ford, E. B., & Veras, D. 2010, *ApJ*, 712, L86
- Petit, A. C., Pichierri, G., Davies, M. B., & Johansen, A. 2020, *A&A*, 641, A176
- Petrovich, C. 2015, *ApJ*, 808, 120
- Pu, B., & Wu, Y. 2015, *ApJ*, 807, 44
- Pyle, T. 2011, *Six Worlds for Kepler-11*
- Ragozzine, D., & Holman, M. J. 2010, arXiv e-prints, arXiv:1006.3727

- Ragozzine, D., Mills, S., Fabrycky, D., & Sharma, V. in prep.
- Rogers, J. G., Dorn, C., Aditya Raj, V., Schlichting, H. E., & Young, E. D. 2025, *ApJ*, 979, 79
- Rowe, J. F., & Thompson, S. E. 2015, arXiv e-prints, arXiv:1504.00707
- Rowe, J. F., et al. 2014, *ApJ*, 784, 45
- Schlaufman, K. C., & Halpern, N. D. 2021, *ApJ*, 921, 24
- Schmitt, J. R., et al. 2014, *AJ*, 148, 28
- Shahaf, S., Mazeh, T., Zucker, S., & Fabrycky, D. 2021, *MNRAS*, 505, 1293
- Shallue, C. J., & Vanderburg, A. 2018, *AJ*, 155, 94
- Siegel, J. C., & Fabrycky, D. 2021, *AJ*, 161, 290
- Smith, A. W., & Lissauer, J. J. 2009, *Icarus*, 201, 381
- Swift, J. J., Johnson, J. A., Morton, T. D., Crepp, J. R., Montet, B. T., Fabrycky, D. C., & Muirhead, P. S. 2013, *ApJ*, 764, 105
- Tamayo, D., Gilbertson, C., & Foreman-Mackey, D. 2021a, *MNRAS*, 501, 4798
- Tamayo, D., Murray, N., Tremaine, S., & Winn, J. 2021b, *AJ*, 162, 220
- Tamayo, D., et al. 2020, *Proceedings of the National Academy of Science*, 117, 18194
- Ter Braak, C. J. F. 2006, *Statistics and Computing*, 16, 239
- Thompson, S. E., et al. 2017, ArXiv e-prints
- Tuchow, N. W., Ford, E. B., Papamarkou, T., & Lindo, A. 2019, *MNRAS*, 484, 3772
- VanWyngarden, M., & Cloutier, R. 2024, *AJ*, 168, 154

Weiss, L. M., et al. 2024, *ApJS*, 270, 8

Wikimedia. 2007, *Orbital Elements*

Winn, J. N. 2010, in *Exoplanets*, ed. S. Seager, 55–77

Wisdom, J. 1980, *AJ*, 85, 1122

Wolszczan, A., & Frail, D. A. 1992, *Nature*, 355, 145

Youdin, A. N., & Goodman, J. 2005, *ApJ*, 620, 459

Zhu, W., & Dong, S. 2021, *ARA&A*, 59, 291