Spring 2022

A Dynamical Synthesis of Planetary Systems

Malena Rice
Yale University Graduate School of Arts and Sciences, malena.rice@yale.edu

Abstract

A Dynamical Synthesis of Planetary Systems

Malena Rice

2022

Over the past three decades, complementary lines of evidence have each provided tantalizing hints about the underlying mechanisms driving the diverse set of planetary system architectures. This dissertation leverages dynamics to synthesize the various components of planetary systems, including stars, planets, and minor planets. My work progresses at the intersection of subfields, drawing evidence from both solar system and exoplanet studies to advance a cohesive picture of planetary system evolution.

This dissertation is fundamentally focused on interactions between the components of planetary systems. As a result, it is organized into three segments detailing the relationship between these components. A brief summary is provided as follows.

- Part I (Chapter 2): The Star-Minor Planet Connection. This chapter explores the use of occultation measurements, in which foreground asteroids briefly block out the light of background stars, as a mechanism to precisely probe the positions of minor planets within the solar system. We demonstrate that this method can be applied to constrain the presence of neighboring masses, including the predicted "Planet Nine", in the distant solar system.

- Part II (Chapters 3-4): The Planet-Minor Planet Connection. These two chapters examine how minor planets can inform our understanding of planets more broadly. In Chapter 3, we describe a novel algorithm developed to directly search for distant solar system objects relevant to the Planet Nine hypothesis using data from the Transiting Exoplanet Survey Satellite (TESS). Then, in Chapter 4, we demonstrate that the long-period Neptune-mass exoplanet population suggested by protoplanetary disk

images can also efficiently eject neighboring minor planets, accounting for the high rate of observed interstellar objects passing through the solar system.

- Part III (Chapters 5-7): The Star-Planet Connection. These three chapters investigate the relationship between stars and planets in two distinct ways: through compositional studies and through dynamical analyses. In Chapter 5, we describe the development of a machine learning algorithm that rapidly extracts stellar parameters, including 15 elemental abundances, from input optical stellar spectra. In Chapter 6, we introduce the Stellar Obliquities in Long-period Exoplanet Systems (SOLES) survey to investigate the origins of exoplanet spin-orbit misalignments. Finally, in Chapter 7 we conduct a population study of the stellar obliquity distribution that provides evidence for high-eccentricity migration and tidal damping as the two key mechanisms crafting the dynamical evolution of hot Jupiters.

A Dynamical Synthesis of Planetary Systems

A Dissertation
Presented to the Faculty of the Graduate School
of
Yale University
in Candidacy for the Degree of
Doctor of Philosophy

by
Malena Rice

Dissertation Director: Gregory P. Laughlin

May 2022

Contents

Acknowledgments

A dissertation is never a particularly easy project to complete, and one completed in the midst of a global pandemic is certainly no exception. I would like to begin by thanking the wonderful mentors and collaborators who have made my graduate experience such a joy.

It is extraordinarily difficult to put into words my gratitude for my advisor, Greg Laughlin. Greg has always had my best interests at heart and has provided unending, unconditional support throughout my graduate career. I am extremely fortunate to have had such an encouraging and optimistic supporter during the entirety of my PhD.

Songhu Wang and John Brewer have been both close collaborators and wonderful friends to me over the course of my dissertation. I am eternally grateful for their unending support, advice, and positivity.

I am grateful to my dissertation committee at Yale, which includes Greg Laughlin, Debra Fischer, Héctor Arce, and Dan Foreman-Mackey, as well as my external collaborators and letter writers Gabriella Contardo, Smadar Naoz, Matt Holman, Giovanna Tinetti, and George Ricker. I have been incredibly lucky to receive such unwavering support throughout graduate school.

I would also like to thank my student collaborators for the wonderful experience of working with them and for bringing so much laughter and light to everyday work. In particular, I would like to acknowledge Hanna Adamski, Jude Gussman, Keduse Worku, Alexandra Cruz, Kaitlyn Sarkissian, and Rachel Feng.

I would like to thank Yale Astronomy's administrators for all of their hard work to keep the department running smoothly, even through such unusual times. I would especially like to thank Robyn Lisone, Teena Griggs, and Kim Monocchi for bringing so much lighthearted fun to the department.

My community has been the greatest joy of my graduate school experience. I have been incredibly fortunate to meet many wonderful individuals who have inspired and encouraged me, and who have truly enhanced my experience in the field. I would like to acknowledge a few individuals who have made a special impact on me.

First, I want to extend an enormous thank you to Imad Pasha for convincing me to try out astronomy on a whim over a late-night conversation nine years ago, and for being by my side through thick and thin ever since. Thank you for being my rock and for the friendship of a lifetime.

I would like to thank Christopher Agostino for his incredible and unfaltering support from across state borders. You are a true light.

I am grateful to my lifelong friends who have always provided their love and support from afar and who have truly been by my side through it all. I would especially like to acknowledge Hanna Rivera, Caitlin Smith, Sabrina Berger, Camille Biscarrat, Aaron Kusik, Ruth Latragna, Darrah Covert, and Mariah Clayton.

My incredible cohort – Tim Miller, Aritra Ghosh, and Uddipan Banik – is comprised of individuals who I will always count among my dearest friends. I thank them for our many monthly cohort dinners and the unconditional support and positivity that we have always expressed towards each other. We decided to be friends from the start, and I cherish my relationship with each of you.

I am grateful to Konstantin Gerbig and Chris Lindsay for bringing so much laughter and fun into my life over the past two years, and for motivating me to be stronger in both body and spirit. Thank you for always looking out for me during the final years of my PhD.

I am grateful to my downstairs neighbors, Seamus Dwyer and Jeong Yeon Lee, for bringing so much warmth and kindness into my final year at Yale. Your friendship has been a tremendous gift.

My Poorvu Center community has provided an unwavering external support system to me outside of the department, and I am eternally grateful for the kindness and positivity of this group that I have come to think of as family. I would especially like to acknowledge Suzanne Young and Gina Hurley and to thank them for a wonderful four years with the Teaching Fellows team. I would also like to thank Julia Istomina and Elena Kallestinova for a fulfilling 1.5 years at the Poorvu Center's Graduate Writing Lab.

I am grateful to my astro[sound]bites team, Alex Gagliano and Will Saunders, for many great laughs and for all of their support over the past three years. It has been incredibly rewarding to build our podcast from the ground up together, and I am so proud of all that we have created.

I am so grateful for my beautiful, supportive friendships that have grown over the course of my graduate career. I would especially like to thank Yasmeen Asali, Sarah Millholland, Sanah Bhimani, Emily Kuhn, Luna Zagorac, Ava Polzin, Abby Mintz, Juan Guerra, Joel Ong, Claire Dickey, Suchi Narayanan, Gary Chew Kiat Shing, Fabio Pacucci, Stella Zhong, Kelly Backes, Emre Engin, and Chelsea Li for all of the fun moments, laughter, happiness, and support that you have provided. It is such a joy to have so many incredible, brilliant, inspirational people in my life.

I am lucky to have met wonderful colleagues and friends at conferences, workshops, and summer schools over the course of my dissertation, who have provided so much laughter and happiness. I cannot name them all, but I would like to acknowledge a few who have been particularly special to me: Alex Teachey, Ella Holme, Quang Tran, Jade Checlair, and Sanjana Curtis.

Last, but certainly not least, I am grateful to my family for their continual encouragement throughout the entirety of my career. I am grateful to have had strong supporters believing in me from the very beginning.

Dynamics is a powerful reminder that systems are constantly changing and evolving. The study of planets is a powerful reminder of how exquisite our own Earth is. It is a privilege to work in field so deeply intertwined with hope, gratitude, and inspiration.

List of Figures

List of Tables

Foreword

All scientific chapters presented within this dissertation are provided in their complete, published forms. We have elected to maintain this format such that each chapter stands alone and can be read independently, as a component of a cohesive whole. The purpose of the introduction, therefore, is to provide context as to how these chapters fit together, as well as motivation for the set of projects conducted over the course of this PhD.

We have chosen to use lowercase type for the phrase "solar system", as a matter of personal preference. The argument for an uppercase Solar System is that the term may be considered a proper noun in reference to Sol, the Sun. In the context of this work, we consider "solar" to be a descriptor, in which we treat the solar system as merely one of many planetary systems. This is a tricky dichotomy, as, by virtue of our vantage point, evidence obtained from the solar system is categorically different from and substantially more detailed than that which can be extracted from extrasolar systems. Nevertheless, we prefer the lower-case version to combat a subconscious tendency towards treating the solar system as one that is unlike any others.

Chapter 1

Introduction

1.1 Dynamics as a Probe of Planetary Systems

The projects encompassed by this dissertation are each inspired by or directly draw upon dynamics: a broadly applicable framework that describes how bodies gravitationally interact. While dynamics spans size and mass scales across astrophysics, the focus of this work will be at the scales of planetary systems, from minor planets through stars. We will begin with some brief general notes about dynamics and why it is an appropriate tool to develop a cohesive theory of planetary system formation.

The power of dynamics lies in its emphasis on the relationship between the objects in a system. Dynamics uniquely enables us to probe how bodies interact with each other as a function of time and proximity. Thus, dynamics intrinsically describes how systems evolve, from their initial formation to their currently observed state. This past and future evolution can be understood through physical models, ideally while invoking the fewest possible low-likelihood requirements.

While dynamical evolution may, in some cases, be directly probed through observations, the timescales of dynamical processes tend to be much longer than a typical PhD (or, for that matter, a typical human lifetime). In this way, dynamics is, at its core, a theoretical field. Systems in space are captured at a given moment in time, and it remains for the dynamicist to disentangle why and how they look as they do, as well as what the implications are for our understanding of astrophysics. Dynamicists are fundamentally storytellers, piecing together snapshots of the Universe to develop a cohesive narrative of our cosmic history.

Systems are always shifting, even if they appear static in a given moment. Dynamics provides a language to explain this shift on astronomical scales, making it the ideal tool to describe the past and future history of planetary systems.

1.2 Planetary System Formation and Evolution

The goal of this dissertation is to advance models of how planetary systems form and evolve through the lens of dynamics. To provide broad context and motivation for the studies encompassed by this work, we describe here the prevailing models for this evolutionary process, as well as the underlying evidence that these models draw from. We summarize two complementary lines of evidence that have informed existing models: the detailed evidence drawn from the solar system (Section 1.2.1), as well as the broader, multivariate evidence provided by the thousands of known extrasolar systems (Section 1.2.2).

1.2.1 Insights from the Solar System

Until just a few decades ago, the solar system was the only known example upon which all models of planetary system formation were based. As a result, the historical context for planet formation theories comes primarily from models of the solar system's evolution. We review the key solar system formation theories here, including the constraints drawn from the solar system that have been leveraged to understand planet formation more generally.

1.2.1.1 The Nebular Hypothesis

In 1755, the philosopher Immanuel Kant proposed the "nebular hypothesis" for the solar system (Kant, 1755), suggesting that the Sun and its surrounding planets were created from disperse gas and dust brought together by a force of attraction. This theory was later refined and quantified by mathematician Pierre-Simon Laplace (Laplace, 1796). In the nebular hypothesis, the Sun formed from a contracting, spinning clump of gas and dust. This clump radiated away energy, causing the gas and dust to contract further and to spin faster to ultimately produce a star. As the Sun contracted, rings of gas and dust were left behind around the star, from which planets and their moons later accreted.

While many of the quantitative details were not fully worked out in the 1700s, Kant and Laplace laid the groundwork for what remains the prevailing model for the formation of planetary systems. Stars are initially born from giant molecular clouds in the interstellar medium, and these molecular clouds gravitationally contract when sufficient overdensities stochastically form. Those contracting clumps of overdensities accrete additional gas from their surrounding environment to ultimately produce stars. As material falls inwards, the forming stars spin more rapidly to conserve angular momentum. A disk of gas and dust, or a protoplanetary disk, forms around the nascent star's equator, and planets form from rings of material within that disk.

1.2.1.2 The Nice Model

After the solar system's protoplanetary disk dispersed and the solar system planets fully formed, the planets continued to interact before reaching their currently observed orbits. The most widely accepted model for this phase of evolution, known as the Nice Model, suggests that all the solar system planets were formed in a compact configuration spanning a radius of roughly 17 au, with an external massive disk of material (Tsiganis et al., 2005). After a period of slow migration, Jupiter and Saturn crossed the 1:2 mean motion resonance, at which point the eccentricities and inclinations of both planets were excited to their currently observed values. Uranus and Neptune were also scattered outwards at this time, disrupting the outer planetesimal disk and scattering minor planets inwards to populate the Jovian Trojan reservoirs (Morbidelli et al., 2005) while inducing the Late Heavy Bombardment (Gomes et al., 2005).

The evidence for the Nice Model is largely drawn from the solar system minor planet populations, as well as the current orbits of the solar system planets. Collectively, the Nice Model is able to account for the nonzero eccentricities and inclinations of the solar system planets; the spike in craters on the Moon that occurred about 700 million years

after the planets had formed (the Late Heavy Bombardment); and the resemblance between Jovian Trojans and comets at visible wavelengths (Fernández et al., 2003; Barucci et al., 2002). We note that some potential caveats to the Nice Model remain: for example, recent observations have revealed that the Jovian Trojan asteroids are systematically redder than the Kuiper belt from which the population should have originated under the Nice Model framework (Jewitt, 2018).

The Nice Model was later modified to include viscous stirring and a late instability due to the crossing of several weak secular resonances (Levison et al., 2011). Further modifications also include the "jumping-Jupiter" scenario in which an ice giant entered a close encounter with Jupiter to produce a rapid separation of Jupiter's and Saturn's orbits (Brasser et al., 2009). This prevents the smooth migration of Jupiter and Saturn's orbits, which would have led to sweeping secular resonances across the inner solar system that would have excited the eccentricities of the terrestrial planets beyond their current values. Current models of the solar system are statistically best matched when assuming that the young solar system began with a fifth ice giant planet that produced the jumping-Jupiter scenario and that would have likely been ejected from the solar system after this encounter with Jupiter (Nesvorný, 2011). The Nice Model, together with these modifications, comprises the most widely accepted current model for the solar system's formation.

1.2.1.3 The High-Inclination Solar System

While the classical models for solar system formation capture many elements of the currently observed system, several outstanding mysteries remain. Minor planet populations lay the foundation for leading models of the solar system's early dynamical evolution, as demonstrated by the Nice Model. Most objects within these populations are roughly coplanar with the ecliptic plane of the solar system, with small inclinations excited by dynamical interactions with the solar system planets.

However, several dynamically interesting populations exist in the high-inclination solar system that challenge classical formation theories. For example, minor planets known as Centaurs, defined as small bodies with $a_J < a < a_N$ and $a_J < q < a_N$ (Jewitt, 2009), have been discovered on chaotic, high-inclination orbits. Here, a is semimajor axis, q is perihelion distance, and the J and N subscripts represent Jupiter and Neptune, respectively. Centaurs have relatively short dynamical timescales of only $\sim 10^7$ years due to their interactions with the giant planets of the solar system (Holman & Wisdom, 1993). This raises the question of how they were dynamically excited to their current orbits – some of which are retrograde – and from what reservoir of small bodies the population is continually replenished.

The long-term stable detached Kuiper belt objects, a population of long-period minor planets with perihelia outside of the orbit of Neptune, have also been found with unexpectedly high eccentricities and inclinations (e.g. Brown et al., 2004; Trujillo & Sheppard, 2014). Because this population does not directly interact dynamically with any of the eight known solar system planets, it is unclear how the elevated eccentricities and inclinations of these objects were attained, leaving them on their current stable orbits. These "extreme trans-Neptunian objects" (eTNOs) may have originated from either the dynamically hot scattered disk population of the Kuiper belt (Levison et al., 2008; Nesvorný, 2018) or from the inner Oort cloud, with a more comet-like origin (Vokrouhlický et al., 2019; Batygin & Brown, 2021).

The eTNO population raises a further mystery: Trujillo & Sheppard (2014) found that the periastra of the observed eTNOs, quantified by the mean longitude of perihelion $\overline{\omega}$, appear to be preferentially clustered in one direction on the sky. Batygin & Brown (2016) later demonstrated that the longitude of ascending node Ω is clustered for the long-term stable eTNOs, as well, demonstrating that the orbits of these objects are grouped in both their direction on the sky and their mean orbital inclination.

Both of these observed effects are unexpected if they are present – that is, if they are not an

artifact of observational biases. The orbits of the eTNOs should precess at different rates that would cause them to spread out across the sky over a timescale of about a billion years (Batygin et al., 2019). Objects in the solar system do not orbit in a perfectly Keplerian potential, since the solar system is not a 2-body system, and this deviation leads to orbital precession. The deviation V from a Keplerian potential can be approximated by summing over the perturbations induced by the four outer giant planets ($j = 5 - 8$) in the secular approximation:

$$V = -\frac{1}{8}\frac{GM_\odot}{a}\frac{3\cos^2 i - 1}{(1-e^2)^{3/2}}\sum_{j=5}^{8}\frac{m_j a_j^2}{M_\odot a^2} \tag{1.1}$$

$$\frac{d\overline{\omega}}{dt} = -\frac{\sqrt{1-e^2}}{e\sqrt{GM_\odot a}}\frac{\partial V}{\partial e} = \frac{3}{4}\sqrt{\frac{GM_\odot}{a^3}}\frac{1}{(1-e^2)^2}\sum_{j=5}^{8}\frac{m_j a_j^2}{M_\odot a^2} \tag{1.2}$$

$$\frac{d\Omega}{dt} = -\frac{-1}{\sqrt{GM_\odot a(1-e^2)}\sin i}\frac{\partial V}{\partial i} = \frac{3}{4}\sqrt{\frac{GM_\odot}{a^3}}\frac{\cos i}{(1-e^2)^2}\sum_{j=5}^{8}\frac{m_j a_j^2}{M_\odot a^2} \tag{1.3}$$

As a result of this deviation, the orbits of the eTNOs do not perfectly close, and they should appear randomly distributed across the sky today in the absence of additional perturbers. The collective properties of high-inclination solar system objects, including the orbital properties of both Centaurs and eTNOs, pose an unsolved problem for current models of the solar system.

1.2.1.4 A Proposed Ninth Solar System Planet

A ninth solar system planet has been proposed as a single, relatively simple solution to all of the problems outlined in Section 1.2.1.3. A massive, eccentric, and highly inclined planet in the distant solar system could provide the restoring torque necessary to prevent the differential precession of eTNOs. This same planet could also excite high-inclination

a	e	i	mass	Reference
700 au	0.6	30°	$10M_\oplus$	Batygin & Brown (2016)
400-800 au	0.2-0.5	15-25°	5-$10M_\oplus$	Batygin et al. (2019)
380 au	0.2	16°	$6M_\oplus$	Brown & Batygin (2021)

Table 1.1: Predicted Planet Nine parameters over time.

and eccentric orbits, such as those of many Centaurs and eTNOs.

As dynamical models have been refined over recent years, the parameters expected for this "Planet Nine" have undergone several iterations. We include a few sets of predicted Planet Nine parameters in Table 1.1 for reference, as well as a visualization of the expected orbital parameters from Batygin et al. (2019) in Figure 1.1. The most updated prediction to date for Planet Nine's parameters places it on a wide ($a \sim 380$ au), eccentric ($e \sim 0.2$), and inclined ($i \sim 16°$) orbit as a relatively massive ($M \sim 6M_\oplus$) planet (Brown & Batygin, 2021). Although this most recent iteration of Planet Nine parameters would make the object lower-mass than the originally predicted version, the planet's closer average distance would produce a brighter overall signal (V magnitude $\sim 19 - 22$).

While a sub-Neptune-mass Planet Nine would have a relatively large radius ($\sim 2 - 4R_\oplus$), it would not have necessarily been discovered by previous surveys due to its wide orbit and its accordingly dim expected magnitude. The light observed from Planet Nine should be dominated by reflected sunlight, with a flux that rapidly falls off with distance as d^{-4}. Because of this steep decrease in reflected light flux with distance, a wide swath of parameter space remains for undiscovered outer solar system planets such as Planet Nine. If Planet Nine exists, it would pose an additional challenge for solar system formation models: while the existence of the planet could solve several mysteries related to the high-inclination solar system minor planet populations, a new mystery would be put forth regarding how Planet Nine itself reached its current eccentric and high-inclination orbit.

We emphasize that, in addition to the proposed sub-Neptune-sized Planet Nine, one or more

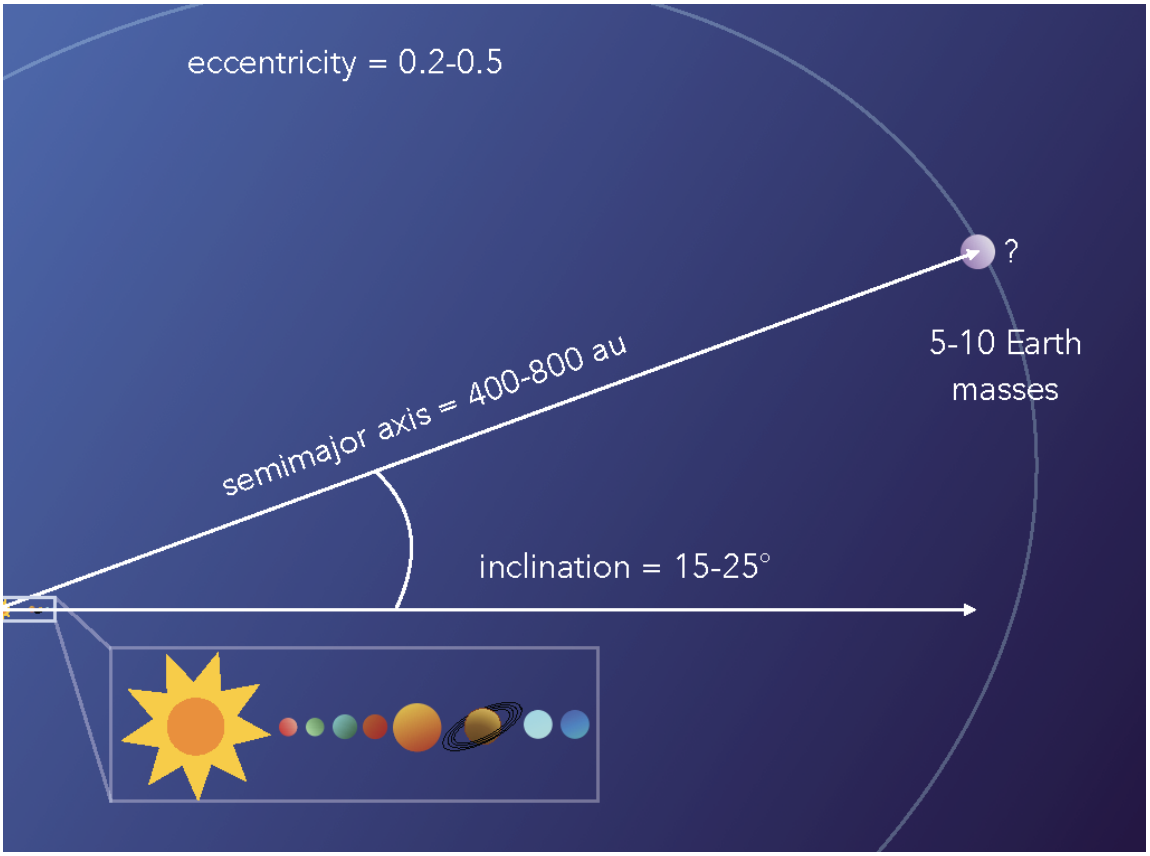

Figure 1.1: Planet Nine parameter ranges from Batygin et al. (2019), with the expected distance to Planet Nine shown to scale for reference.

Mars- to Earth-sized planets have also been suggested to reside in the distant solar system. These propositions have been made in part due to the expected formation mechanism of the solar system, in which multiple rocky planet embryos are expected to have formed from the outer planetesimal disk of the early solar system. Based on dynamical simulations, one or more of these embryos, if they were initially produced, is likely to remain in the outer solar system (Silsbee & Tremaine, 2018). This population of embryos may have produced the stirring needed to excite high-inclination and high-eccentricity orbits among the observed eTNOs.

A separate, more direct line of evidence for a distant, Mars-mass perturber is an observed warp in the mean plane of the Kuiper belt, at semimajor axes $a = 40 - 42$ au, that would be well-explained by a relatively low-mass planet with $a \lesssim 100$ au (Volk & Malhotra, 2017). Although lower-mass, relatively nearby planets have not necessarily been targeted by recent and ongoing Planet Nine searches, they may also feasibly exist and similarly would not yet have been discovered through existing solar system surveys. Perhaps surprisingly, over 400 years after the advent of telescopic astronomy, a vast parameter space remains in the distant solar system for undiscovered planets.

1.2.2 Insights from Extrasolar Systems

The discovery of the first exoplanets has extended theories of planetary systems to, for the first time, be constrained by more than a single data point. This section provides an overview of the complementary new insights attained through studies of extrasolar systems.

1.2.2.1 The Diverse Architectures of Planetary Systems

Exoplanets have been discovered on sideways and upside-down orbits, around the breadth of stellar types, and with system architectures entirely unlike that of the solar system. The growing set of recently discovered exoplanets provides constraints on the multivariate

formation and evolution pathways for planets, including many planet types that are not found in the solar system. Several key insights drawn from the diverse set of exoplanet system architectures are summarized below.

1. *Planets Exist At a Wide Range of Orbital Separations:* A wide diversity of planets has been found orbiting not only at star-planet separations similar to those found in the solar system, but also at much shorter and longer orbital periods. The regularity with which these types of planets form provides an important clue to understand the prevalence of migration patterns in the early evolution of planetary systems.

   At the time of writing, the vast majority of confirmed exoplanets have been discovered using the transit method, which is biased towards short-period planets that are relatively likely to pass in front of their host star during their orbit. Many confirmed, transiting exoplanets lie far interior to the analog of Mercury's orbit ($P = 88$ days) in the solar system, often with multiple planets present around a given star in a compact configuration. Most known transiting exoplanets to date were discovered using photometry from the space-based *Kepler* mission (Borucki et al., 2010) that observed a small patch of the sky continuously for four years, from 2009 to 2013. Ongoing ground-based surveys have continued to search for short-period transiting exoplanets, while the space-based Transiting Exoplanet Survey Satellite (TESS; Ricker et al., 2015) has, since 2018, scanned across the full sky in search of exoplanets around bright stars.

   Radial velocity surveys with long temporal baselines (up to $\sim$25 years) have increasingly constrained the occurrence rates of longer-period planets out to several astronomical units, providing a complementary exoplanet search method that is less biased against long-period exoplanets. These surveys measure the Doppler shift of a star as it wobbles on its orbit due to gravitational interactions with a neighboring, massive perturber. Radial velocity demographic surveys suggest that Jovian-mass

$(1-13M_J)$ planets may be found at moderate $(1-20$ au) separations around roughly half of all main sequence stars (Knutson et al., 2014).

At a wider separation from the host star, direct imaging observations have detected the infrared radiation emitted by young (<100 Myr) super-Jupiter planets out to tens of astronomical units (e.g. Marois et al., 2010; Macintosh et al., 2015). These planets appear to be relatively rare, with an occurrence rate $9^{+5}_{-4}\%$ for planets ranging from $5-13M_J$ at $a=10-100$ au (Nielsen et al., 2019). Evidence from the Disk Substructures at High Angular Resolution Project (DSHARP; Andrews et al., 2018; Zhang et al., 2018) that resolves the detailed profiles of protoplanetary disks, as well as evidence from microlensing surveys that directly detect exoplanets (Suzuki et al., 2016), has arisen in parallel to suggest a higher occurrence rate for Neptune-mass planets at wide (∼tens of au) separation from their host star.

2. *Planets Orbit a Wide Range of Stellar Types:* Exoplanets have been found around an enormous range of stellar types. This includes not only those spanning the Main Sequence, but also younger stars (e.g. Clarke et al., 2018); older, post-Main-Sequence stars (e.g. Almenara et al., 2015; Van Eylen et al., 2016; Chontos et al., 2019); and even pulsars (e.g. Wolszczan & Frail, 1992). Free-floating planets have been detected with no stellar host (Bennett et al., 2014; Freeman et al., 2015), as well, providing direct evidence that planets can be ejected through dynamically violent processes during the evolution of their systems. Exoplanets appear to be omnipresent across the Milky Way.

   Some trends of planet occurrence patterns with stellar type have been discovered and well-characterized; for example, hot Jupiters are preferentially found around metal-rich stars (Fischer & Valenti, 2005). Stars with planets have been found to typically be enriched in silicon and nickel (Robinson et al., 2006), and compact multiplanet systems – such as the "peas-in-a-pod" type systems (Millholland et al.,

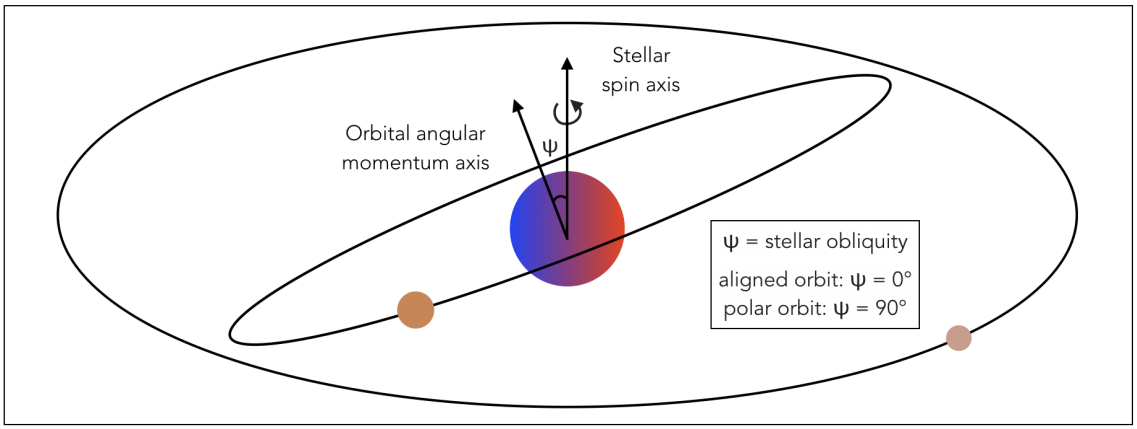

Figure 1.2: Geometry of an example planetary system. Aligned planets orbit within the plane defined by the host star's equator, in the same direction as the stellar spin. Polar planets orbit in a perpendicular plane, and retrograde planets orbit backwards across the host star's equator. Rossiter-McLaughlin and Doppler tomography observations provide λ, the 2D sky-projection of the 3D obliquity Ψ.

2017; Weiss et al., 2018) commonly detected by the *Kepler* spacecraft – are typically found around metal-poor stars (Brewer et al., 2018). These compositional trends provide tantalizing clues regarding the initial conditions that produce the range of planetary system types observed today.

3. *Planets are Found at a Wide Range of Projected Spin-Orbit Angles:* The obliquity of a star, or the degree of alignment between the star's spin axis and its companion planets' net orbital angular momentum axis (see Figure 1.2), provides crucial insights into the dynamical evolution of the surrounding system. The solar system planets are all near alignment with the invariable plane of the solar system, with a maximum deviation of $\sim 7^\circ$ for Mercury. However, this coplanarity with the host star's equator has not been found to be universal across extrasolar systems.

   Most stellar obliquity measurements to date have been made for hot Jupiter systems because of the relatively strong signal attainable across a single transit. The sky-projected obliquities determined through the Rossiter-McLaughlin (Rossiter, 1924; McLaughlin, 1924) and Doppler tomography (e.g. Cameron et al., 2010) methods,

which have provided the majority of current obliquity constraints, have demonstrated that hot Jupiters may span a wide range of spin-orbit angles, with orbits ranging from prograde to polar and retrograde.

Trends in stellar obliquity can constrain the prevalence of various dynamical processes crafting the observed distribution of exoplanet orbits. Imprints of planetary systems' early dynamical evolution are embedded within the currently observed stellar obliquity distribution. However, the current distribution is not necessarily representative of the primordial set of stellar obliquities. Star-planet tidal interactions can act to realign misaligned planets over time, providing an additional factor that can confuse the interpretation of measured obliquities (e.g. Winn et al., 2010; Albrecht et al., 2012; Winn & Fabrycky, 2015). As a result, the observed set of stellar obliquities is likely not representative of the maximum obliquities typically reached by planetary systems.

The high observed frequency of misalignments suggests that strong dynamical interactions may be common in the early lifetimes of hot Jupiter systems, providing an important clue regarding how the observed planets reached their current orbits. Several potential formation routes for hot Jupiters are outlined in Figure 1.3, each of which can produce the relatively high obliquities that are observed today.

1.2.2.2 Insights from Protoplanetary Disks

In recent years, high-resolution images of protoplanetary disks have provided revolutionary new insights enabling a more detailed understanding of the planet formation process. Some of this evidence has been direct – forming planets, embedded within protoplanetary disks, were directly imaged for the first time within the PDS 70 system (Keppler et al., 2018; Isella et al., 2019), for example. These nascent planets have provided the first direct look at young planetary systems actively accreting material within a protoplanetary disk, where

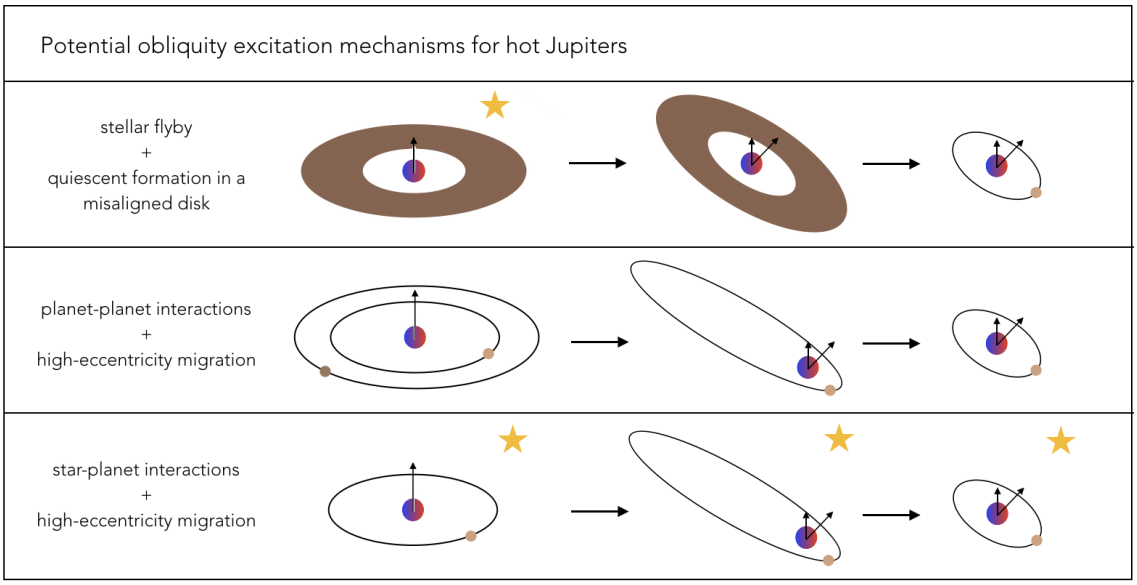

Figure 1.3: Potential hot Jupiter formation mechanisms; each can induce misalignments in hot Jupiter systems, but the stage at which the misalignment is obtained varies.

the surrounding gas has not yet dissipated. Furthermore, candidate circumplanetary disks were discovered around both PDS 70 b (Christiaens et al., 2019) and PDS 70 c (Benisty et al., 2021), providing the first direct evidence for this new type of disk around planets, from which moons are predicted to form.

Beyond direct evidence for forming planets, protoplanetary disks also provide further insights into the environment within which planets form. More massive stars typically host more massive protoplanetary disks (Pascucci et al., 2016), and recent observations have demonstrated a positive correlation between not only $M_{\rm dust}$ and M_*, but also $M_{\rm gas}$ and M_* (Ansdell et al., 2016). This correlation helps to explain the lower frequency of giant planets observed around low-mass stars (Endl et al., 2006). Oddly, protoplanetary disks have been observed to contain systematically less dust mass than is required to produce the observed populations of exoplanets, suggesting either very early formation of massive planetary cores or a continual inflow of material into the protoplanetary disk system (Manara et al., 2018).

Furthermore, protoplanetary disks have been discovered with an array of substructures, suggesting the dynamical influence of embedded planets that are not yet directly resolvable. These substructures are especially notable in results from the DSHARP survey (Andrews et al., 2018) of high-resolution radio images taken with the Atacama Large Millimeter Array (ALMA). The DSHARP survey suffers from selection effects, such that it is biased towards high-mass stars and relatively nearby systems for which protoplanetary disk substructures would be easily resolvable. However, surveys that have focused on lower-mass stars more representative of the "typical" system have also found evidence for the universality of disk substructures in protoplanetary disks (Long et al., 2018).

These substructures have been detected both through ALMA image data and through velocity maps that demonstrate non-Keplerian flows coincident with the radial location of axisymmetric gaps in resolved protoplanetary disks (Pinte et al., 2019). These complementary lines of evidence have been proposed as a potential new exoplanet detection technique, where the requirements to claim the existence of an embedded planet include the detection of a gap, enhanced line broadening at the gap, and a disturbance in the velocity map that is both coincident with the center of the line and that is observed in more than one molecular line (Armitage et al., 2020).

While protoplanetary disk substructure analysis does not yet comprise a standard exoplanet detection method, it is a promising subfield that provides highly complementary evidence to that obtained from the primary existing exoplanet detection techniques. Disk substructures are uniquely capable of probing a new parameter space of relatively low-mass, Neptune-sized planets at wide host star separation.

1.2.2.3 From Protoplanetary Disks to Planets

The exact processes involved in the transition from a protoplanetary disk to a system of discrete planets are poorly understood, and detailed formation models, primarily studied

through computational hydrodynamics simulations and laboratory experiments, comprise an active field of study. Evidence from extrasolar systems has so far suggested two primary channels for planet formation.

The first proposed channel is gravitational instabilities in a protoplanetary disk (Boss, 1997). Gravitational instabilities can occur in self-gravitating disks with sufficiently short cooling timescales that enable disk fragmentation as energy is dissipated by shocks and turbulence (Gammie, 2001). This mechanism leads to runaway growth as the disk fragments and accretes into one or more giant gaseous protoplanets, which ultimately collapse into massive planets.

The second proposed mechanism is core accretion, in which planetesimals accrete and grow over time into successively more massive cores. These cores may either make up rocky planets such as those of the inner solar system, or, after surpassing a mass threshold, they may undergo runaway accretion to produce gas giant planets (Pollack et al., 1996).

While core accretion is traditionally favored, several caveats arise when the process is modeled in detail. For example, silicates begin to bounce off of each other, rather than coagulating, after they reach a size of a few millimeters (the "bouncing barrier"; Zsom & Dullemond, 2008; Güttler et al., 2009). It is also difficult to form planetesimals larger than roughly one meter in size (the "meter-size barrier") without having them either fragment or rapidly spiral inwards to be engulfed by their host star (Weidenschilling, 1977).

Although a general consensus has not yet been reached, potential solutions have been proposed for each of these issues. Recent work has found that the bouncing barrier may be overcome by electrical charging that occurs when millimeter-size grains collide (Steinpilz et al., 2020). A mechanism known as the "streaming instability" has been suggested as a solution to the meter-size barrier, as well. In the streaming instability, solid and gaseous particles in a disk are assumed to be only moderately coupled, producing a pressure

gradient (Youdin & Goodman, 2005). Mutual interactions between gas and dust particles then prevent rapid inward radial drift, enabling growth beyond the meter-size barrier (e.g. Johansen & Youdin, 2007; Johansen et al., 2009). These mechanisms are broadly applicable across both the solar system and extrasolar systems.

It is possible that each of these mechanisms forms planets in certain cases. For example, it is difficult to form compact multiplanet systems through gravitational instabilities, which tend to produce more massive planets through runaway accretion. On the other hand, core accretion may be particularly inefficient in the outer regions of protoplanetary disks, where a given planetesimal will complete an orbit far less frequently than in the inner disk, making gravitational instabilities more compelling.

The prevalence of each mechanism is obfuscated by additional poorly constrained factors, such as the regularity with which planets migrate. A combination of compositional studies and exoplanet demographics will provide further constraints for planet migration over the coming years.

1.2.2.4 The Discovery of Interstellar Objects

Interstellar objects are a challenging type of small body to categorize, in that they are discovered passing through the solar system but originated from extrasolar systems. We place them in the "extrasolar" systems category because this population is scientifically compelling as an informant of how extrasolar systems form, what materials they are comprised of, and how much material they eject. With that said, interstellar objects are observed within the solar system and are therefore possible to study with similar detail to that of the other solar system minor planet populations.

Until recently, extrasolar systems were studied exclusively from afar. In 2017, however, the first interstellar object – a small body originating in an extrasolar system and passing through the solar system on a hyperbolic orbit – was discovered tumbling on its way out

of the solar system (Meech et al., 2017). This object, 1I/'Oumuamua, displayed several unusual properties that remain to be fully explained based on the sparse data that was collected during the short window of time between its discovery and its rapid dimming as it exited the solar system.

Observations of 'Oumuamua taken with the Gemini South telescope, the Keck II telescope, the Very Large Telescope (VLT), the Canada France Hawaii Telescope (CFHT), and the United Kingdom InfraRed Telescope (UKIRT) demonstrated dramatic light curve variations between $g \sim 22 - 25$, suggesting that the interloper was a flattened and/or elongated object tumbling through space with a 7.34 hour period (Meech et al., 2017). The current best model of 'Oumuamua's light curve indicates that it is most consistent with a flattened, pancake-like shape tumbling through space, although a cigar-shaped model also fits the data relatively well (Mashchenko, 2019). Objects of these extreme dimensions are not known to exist within the solar system; however, the Kuiper belt object 2014 MU69 (Stern et al., 2019), observed in a flyby through the New Horizons extended mission (Stern et al., 2018), was also found to be a surprisingly flattened contact binary.

Perhaps the most puzzling property of 'Oumuamua is that the object was found to be accelerating out of the solar system at a rate faster than could be accounted for based on gravitational effects (Micheli et al., 2018). No CO or CO_2 outgassing was detected based on observations taken with the *Spitzer* space telescope (Trilling et al., 2018), and the sublimation of water ice at an undetected rate likely could not fully account for 'Oumuamua's observed non-gravitational acceleration (Sekanina, 2019). This puzzling observation has led to several propositions that 'Oumuamua may not have been cometary in nature, and that it may have instead originated as a flattened, icy aggregate (Moro-Martín, 2019), an asteroidal object (Portegies Zwart et al., 2018), or through another mechanism. The origins of 'Oumuamua's non-gravitational acceleration have not been agreed upon and likely will not be resolved until additional interstellar objects have been detected and observed in

greater detail.

With these mysteries at hand, interstellar objects have provided a compelling new direction of study for extrasolar systems. 'Oumuamua's flattened shape and non-gravitational acceleration suggest that free-floating minor planets may have an interesting history that is not well-captured by current models of planet formation. In any case, interstellar objects provide an unprecedented opportunity to directly study the material from extrasolar systems up close, making them a compelling new class of object to study, as well as an interesting class of potential targets for upcoming missions such as the European Space Agency's proposed *Comet Interceptor* (Snodgrass & Jones, 2019).

1.3 Motivation for a Cohesive Study

Tremendous strides have been made in our understanding of both the solar system and extrasolar systems, as summarized above. These lines of reasoning are entirely complementary: the solar system provides a single, exquisitely detailed picture of planetary systems, while extrasolar systems offer an unprecedented opportunity to understand the breadth of physical processes that can produce a wide range of planetary system types.

There remains the question of why planetary systems cover such a tremendous range of parameters, and what conditions cause planets to form at all. By carrying out studies to characterize both the solar system and extrasolar systems, we progress towards answering these questions through the junction of multiple interconnected lines of information. We draw primarily from dynamics but do not limit ourselves to a single approach; instead, we maintain flexibility in our methods and utilize a combination of theory, computation, and observation.

This work builds upon several lines of evidence to probe the evolutionary history of planetary systems. As illustrated in Figure 1.4, we focus on the interactions between stars,

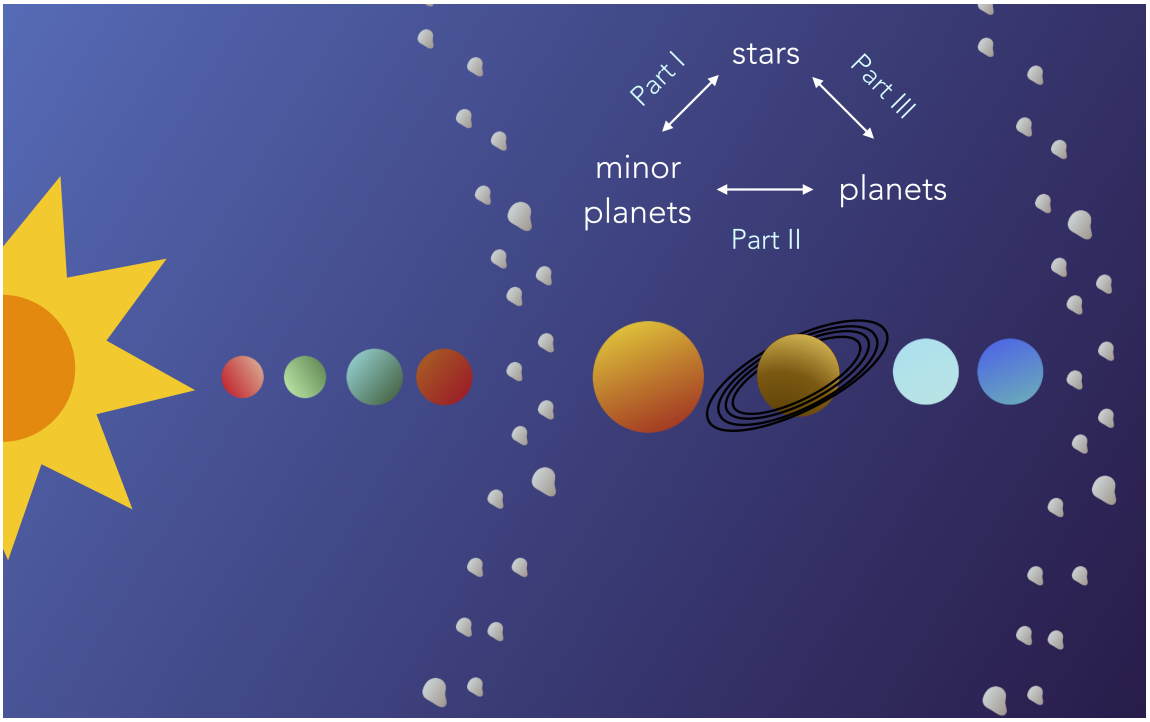

Figure 1.4: Schematic overview of the parts encompassed by this dissertation.

planets, and minor planets to disentangle how planetary systems as a whole evolve, drawing together evidence from several related subfields to develop a cohesive picture.

We begin with the star-minor planet connection in Part I (Chapter 2), exploring how occultation measurements of asteroids over background stars can provide detailed new evidence unveiling the properties of the solar system. Then, in Part II (Chapters 3-4), we describe the planet-minor planet connection that has motivated our search for Planet Nine, as well as how the observed frequency of interstellar objects can provide information regarding the occurrence rates of extrasolar Neptune analogues. In Part III (Chapters 5-7), we describe how both stellar compositions and planetary system architectures can be leveraged to better understand how stars interact with and shape the properties of their surrounding planetary systems. Finally, we conclude in Chapter 8 with future directions and final thoughts.

Part I

The Star-Minor Planet Connection

Chapter 2

The Case for a Large-Scale Occultation Network

Originally published as:

Rice, M. & Laughlin, G. 2019, The Astronomical Journal, 158, 19

2.1 Abstract

We discuss the feasibility of and present initial designs and approximate cost estimates for a large ($N \sim 2000$) network of small photometric telescopes that is purpose-built to monitor $V \lesssim 15$ Gaia Mission program stars for occultations by minor solar system bodies. The implementation of this network would permit measurement of the solar system's tidal gravity field to high precision, thereby revealing the existence of distant trans-Neptunian objects such as the proposed "Planet Nine." As a detailed example of the network capabilities, we investigate how occultations by Jovian Trojans can be monitored to track the accumulation of gravitational perturbations, thereby constraining the presence of undetected massive solar system bodies. We also show that the tidal influence of Planet Nine can be discerned from that of smaller, nearer objects in the Kuiper belt. Moreover, ephemerides for all small solar system bodies observed in occultation could be significantly improved using this network, thereby improving spacecraft navigation and refining Solar System modeling. Finally, occultation monitoring would generate direct measurements of size distributions for asteroid populations, permitting a better understanding of their origins.

2.2 Introduction

Stellar occultations are a long-standing observational technique used in a wide variety of applications. The rings of Uranus were serendipitously discovered during a stellar occultation of the planet that revealed five pre- and post-occultation dips in light produced by five of Uranus's narrow rings (Elliot et al., 1977). Soon thereafter, the presence of Pluto's atmosphere was inferred from the detection of a symmetric, gently sloping lightcurve of the dwarf planet during a stellar occultation (Hubbard et al., 1988; Elliot et al., 1989). The demonstration that Pluto's surface is not a mere airless expanse significantly augmented the motivation for an eventual mission. More recently, with the advent of *Kepler* (Borucki et al.,

2010) and a multitude of ground-based exoplanet detection surveys (e.g. HAT (Bakos et al., 2004), WASP (Pollacco et al., 2006), NGTS (Wheatley et al., 2017)), occultation-based transit photometry has constituted the most successful exoplanet detection method yet; data from exoplanet transit events has been used to measure planetary radii, atmospheric compositions, and system multiplicity (e.g. Winn & Fabrycky, 2015). Further applications of occultation measurements range from ring seismology to study normal-mode planetary oscillations (e.g. Baillié et al., 2011; Mankovich et al., 2019) to developing spatial maps of small solar system bodies (e.g. Elliot et al., 2010; Sicardy et al., 2011; Alvarez-Candal et al., 2014; Buie et al., 2015; Leiva et al., 2017).

An intriguing use of occultations was recently demonstrated in support of the New Horizons Kuiper Belt Extended Mission, which included a flyby of Kuiper Belt Object (KBO) 486958 2014 MU69 (hereafter referred to as MU69) on January 1st, 2019 (Stern et al., 2018; Moore et al., 2018). In order to coordinate this close-range flyby, it was necessary for New Horizons' navigation team to constrain the orbit of MU69 to exquisite precision.

To accurately delineate the orbit, the team first set up a "picket line" of twenty-five portable telescopes perpendicular to MU69's path to observe a predicted occultation on June 2nd-3rd, 2017 (Porter et al., 2018). This initial campaign was unsuccessful; astrometric data of MU69 obtained with the Hubble Space Telescope (HST) in June and July indicated that the telescopes had been displaced by $\sim$80 km relative to the occultation path. Using this improved astrometry to update the predicted track, the team successfully observed an occultation on July 10th, 2017 using the Stratospheric Observatory For Infrared Astronomy (SOFIA; Young et al., 2012) aircraft (Porter et al., 2018). Shortly thereafter, on July 17th, 2017, the team employed another "picket line" of twenty-four 16-inch telescopes to capture an additional occultation, leading to further refinement of MU69's orbital elements and a direct measure of its size. Of the 24 telescopes used for the July 17th event, five successfully observed occultation chords across MU69's solid body (Porter et al., 2018),

while simultaneous HST photometry of the occulted star provided upper limits on additional material within the KBO's Hill sphere (Kammer et al., 2018). A fourth occultation was subsequently predicted through orbital fitting in Porter et al. (2018), resulting in yet another successful observation on August 3rd, 2018.

The frequent occurrence of small telescope-observable occultation events (in this case, 3 events in 2 months) by an object of equivalent diameter $\sim$20 km that subtends only 600 μas, as well as the demonstrated success of precise astrometric measurements using arrays of small ground-based telescopes, motivates us to further investigate the technique's prospects. These observations of MU69 were by no means the first successful application of occultation measurements to study small bodies in the solar system; asteroid occultations have been captured with increasing frequency over time as stellar astrometry and asteroid ephemerides have grown more precise. The case of MU69 is, however, a particularly compelling proof-of-concept that, with a combination of precise stellar parallaxes from the Gaia Mission (Gaia Collaboration et al., 2016) and improved CCD and CMOS (Saint-Pé et al., 2017) detector technologies, a relatively inexpensive network can be used to monitor occultations by distant minor planets with unprecedented precision.

In this paper, we explore the benefits of deploying a network of small, roboticized telescopes across the United States to monitor stellar occultations. Our proposed telescope grid is envisioned to use the infrastructure – including site locations, access roads, communications provisioning, etc. – associated with the "cISP" ultra-low latency microwave network that has been proposed for the continental United States (Singla et al., 2015; Bhattacherjee et al., 2018), thereby lowering costs and minimizing environmental impact. With a small telescope placed at each site on the cISP (see Figure 2.1), the resulting network would span a vast area ($\sim$4000 km $\times$ 2500 km). In our view, the timing for such a project is right, with impetus amplified by the myriad of small solar system bodies that will soon be detected with the Large Synoptic Survey Telescope (LSST; Jones et al., 2016).

We begin by providing a description of the network and its properties in §2.3. Then, §2.4 outlines how the network can act as a gravitational probe for the putative "Planet Nine" (Batygin & Brown, 2016). In the course of this overview, we explore the motivation behind a tidal acceleration search for Planet Nine (§2.4.1), as well as the expected rate of observable occultation events (§2.4.2), the refinement of orbital elements over time (§2.4.3), and methods to distinguish the gravitational effect of Planet Nine from that of as-yet undetected bodies in the Kuiper belt (§2.4.4). §5.8 includes a discussion of various additional science cases for this network, including taking synergistic observations to complement upcoming NASA missions (§2.5.1); studying the morphology and size distribution of Jovian and Neptunian Trojan asteroids (§2.5.2); developing an acceleration map of the outer solar system using observations of Trans-Neptunian Objects (TNOs; §2.5.3); and refining and/or obtaining asteroid diameter measurements (§2.5.4). §2.6 involves an overview of our proposed instrumental setup based on comparisons with existing facilities, while §2.7 addresses alternate potential versions of the network with fewer telescopes. We conclude with final remarks and a summary of our results in §7.7.

2.3 Network Overview

The proposed cISP network, most recently described by Bhattacherjee et al. (2018), consists of line-of-sight microwave relays connecting the 120 largest population centers in the continental United States. It is designed to provide an ultra-low-latency infrastructure for small packet traffic on the Internet, with individual links transmitting at data rates of order a gigabit per second. The cISP takes advantage of both optimized routing topology and the ∼40% faster propagation speed of electromagnetic waves in air by comparison to glass fiber optic cable. The cISP grid's design adopts only *existing* communication towers in the FCC Antenna Structure Registration database,[1] and its sites are strategically

[1] http://wireless.fcc.gov/antenna/index.htm?job=home

placed at topographic high points chosen to confer optimal line-of-sight. Construction of a cISP-based occultation-detection network would leverage access to these locations as well as the basic facilities – including weather stations for relaying real-time meteorological conditions – at each site.

In the system we envision, one roboticized telescope equipped with a high-cadence imager will be placed at each of $\sim$2000 points on the cISP network, resulting in an array that densely samples occultation tracks. For asteroids with diameter $\sim$10 km or larger, multiple telescopes will typically be activated as the shadow crosses the United States. An example is demonstrated in Figure 2.1, which displays two sample trajectories of the binary Jovian Trojan 617 Patroclus-Menoetius – a target of the *Lucy* Jovian Trojan flyby mission (Levison et al., 2016) – overlaid on the current cISP network design, which includes 1913 sites. The uneven site spacing provides a balance between a large coverage area and higher-density regions necessary for detailed studies involving multiple observations of individual occultation events. Occultation tracks run predominantly from East to West from the perspective of ground-based observations near opposition; as a result, telescopes separated in the North-South direction by $D \gtrsim 2\,\mathrm{km}$ will act as statistically independent accumulators of events. The Research and Education Collaborative Occultation Network (RECON; Buie & Keller, 2016) is an excellent example of a successful, currently-operating network of coordinated telescopes designed for occultation observations (e.g. Benedetti-Rossi et al., 2016), albeit on a much smaller scale than is proposed here and without roboticized operations.

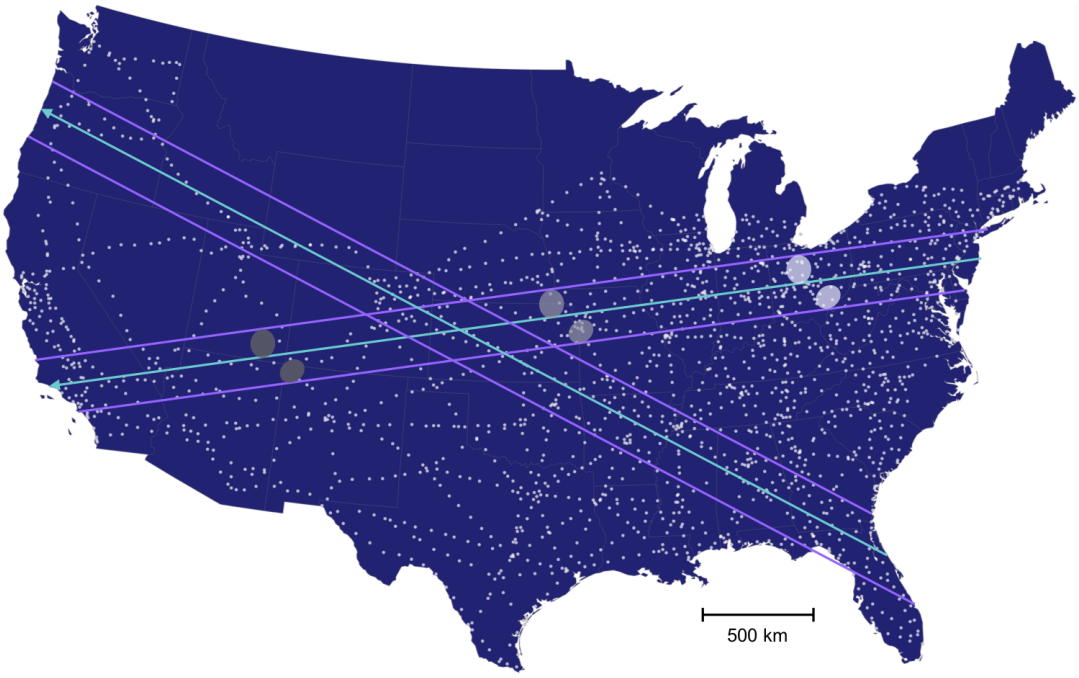

Figure 2.1: Schematic of binary Jovian Trojan asteroid 617 Patroclus-Menoetius occulting the cISP network. Two possible occultation paths are shown to scale, bounded by purple lines with cyan arrows marking the central trajectories. Each white point on the map denotes a cISP network site, where the current cISP network design includes 1913 sites. State boundaries are shown in gray. Relative sizes of the two occulting bodies are based on results from Buie et al. (2015), which reports axial sizes 127 x 117 x 98 km for Patroclus and 117 x 108 x 90 km for Menoetius from previous stellar occultation measurements.

2.4 On the Detection of Planet Nine

2.4.1 The Case for a Tidal Acceleration Search

Authors, notably including Trujillo & Sheppard (2014), Batygin & Brown (2016), and Malhotra et al. (2016), have raised the possibility that the solar system contains an as-yet undetected planet orbiting far beyond Neptune (typical models often invoke an $M \sim 10M_{\oplus}$ body on an eccentric orbit with $P \sim 2 \times 10^4\,\mathrm{yr}$). Evidence for such a "Planet Nine" is indirect, however, and rests primarily on the larger-than-expected degree of orbital alignment among the most distant Kuiper belt objects. Additional clues are provided by the observed population of Centaur-like objects with high inclinations that are difficult to explain through alternative mechanisms (Batygin & Brown, 2016). Numerical simulations by Holman & Payne (2016a), Millholland & Laughlin (2017), Becker et al. (2017), and Batygin & Morbidelli (2017) (among others) have steadily refined the dynamical evidence and now point toward a planet with $m \sim 5-10M_{\oplus}$ with semimajor axis $a \sim 400-800\,\mathrm{AU}$. Parameter ranges based on current constraints are listed in Table 2.1 (Batygin et al., 2019), along with the assigned values used throughout this work.

Direct optical searches for Planet Nine in reflected sunlight are rendered difficult by the d^{-4} decrease in brightness with distance, and they become extremely challenging for $V \gtrsim 23-24$, especially in regions that are crowded with stars. Directed searches at infrared wavelengths (e.g. Meisner et al., 2018) for self-luminous trans-Neptunian planets are stymied by the detection limits of wide-field space-based surveys (e.g. WISE (Wright et al., 2010) and NEOWISE (Mainzer et al., 2011)) unless a planet exists that is much larger and much warmer than expected. Because Planet Nine's emission is expected to peak at submillimeter wavelengths (50-100 μm), cosmic microwave background (CMB) surveys provide an alternative method to search for Planet Nine through its parallactic

Table 2.1: **Planet Nine**

Parameter	Allowed Range	Assigned Value
a	400-800 AU	600 AU
e	0.2-0.5	0.35
i	15-25°	20°
m	5-10$M_\oplus$	7.5$M_\oplus$

motion (Cowan et al., 2016). However, only a few current and upcoming CMB surveys have the requisite sensitivity (Cowan et al., 2016), and the effectiveness of this search method relies upon assumptions about Planet Nine's flux while also requiring a specific observing mode which repeatedly maps the same large regions of the sky at a cadence of a few months. Synergies between multiple detection methods may exist (Baxter et al., 2018), but a definitive assessment remains elusive. Furthermore, in the case that a ninth planet *does not* exist, it is difficult to rule out its presence simply through the lack of a direct detection.

Fortunately, however, indirect detection via the planet's tidal gravity signature may be feasible. Shortly after Batygin & Brown (2016) issued their detailed prediction of the Planet Nine orbit, Fienga et al. (2016) combined a high-resolution solar system dynamical model with telemetry from the Cassini Mission to assess whether the data and model are consistent with the presence of Planet Nine at different locations along Batygin & Brown (2016)'s fiducial orbit. They found no constraint on the planet near apoastron while simultaneously determining that its presence near periastron is inconsistent with the data. This work has since been revisited by Folkner et al. (2016), who found that, under certain assumptions regarding the mass distribution in the Kuiper belt, perturbing bodies larger than $10 M_\oplus$ at $1000\,\mathrm{AU}$ are disallowed. Holman & Payne (2016b,a) further analyze the proposed Planet Nine's effects on the known solar system ephemeris, concluding that models are best fit where Planet Nine is placed at or near apoastron in its orbit.

The instantaneous tidal differential acceleration a_t induced by a perturbing object of mass m at a distance d from the Sun is given by

$$a_t = \frac{x_B G m}{d^3}\,, \tag{2.1}$$

where x_B is the effective baseline across which the tidal acceleration is applied. For an as-yet undetected 7.5 $M_\oplus$ object at $d = 810\,\mathrm{AU}$ (the apoastron Sun-Planet Nine distance corresponding to the Planet Nine properties adopted from Table 2.1) and measured with a baseline $x_B \sim 10\,\mathrm{AU}$ appropriate to Jupiter's Trojan asteroids, $a_t \sim 3 \times 10^{-13}\,\mathrm{cm\,s^{-2}}$, corresponding to a displacement $dx \sim \frac{1}{2} a_t t^2 \sim 30\,\mathrm{m}$ over a 5-year time scale.

A net displacement of order $dx \sim 30\,\mathrm{m}$ incurred over 5 years may, at first glance, appear to be insignificant, but it has the same order of precision to which solar system ephemerides are currently determined. For example, the range residuals to the Cassini orbiter's telemetry reported by Fienga et al. (2016) are of order $dx \sim 10 - 100$ m, whereas the Martian range has been determined to within 3.6 m using just over one year of Mars Express data, and the lunar range to within 2 cm using 10 years of Apollo retro-reflector data from the Apache Point Observatory and 2 years of infrared lunar laser ranging data from the Grasse station (Viswanathan et al., 2017, 2018).

Occultation monitoring can be used to measure the accumulated effects of the tidal acceleration imparted by massive outer solar system objects, including Planet Nine. To fix ideas, we investigate the use of the aggregate Jovian Trojan populations, located at the L4 and L5 Jupiter-Sun Lagrange points, as a probe of tidal gravity. This population provides an abundance of available objects to track – a significant statistical sample that permits triangulation to distinguish perturbations of the Earth's orbit from those of the Trojans' orbit – as well as a relatively long baseline over which the differential tidal acceleration from an external perturber can be measured ($\sim$10 AU across a Jovian Trojan orbit). Over 7000

Jovian Trojans are currently catalogued in the JPL Small-Body Database,[2] and this number is anticipated to surge with the advent of LSST. Both Jovian Lagrange points are well-populated, with number asymmetry $N(\mathrm{L}_4)/N(\mathrm{L}_5) = 1.34$ reported from WISE/NEOWISE observations (Marzari et al., 2002; Grav et al., 2012). Hereafter in §2.4 we use the terms "Trojan" and "Jovian Trojan" interchangeably.

An accurate measurement of tidal acceleration must combine precise astrometry with a rigorous dynamical model of the full solar system to integrate the trajectories of occulting objects. Together, the combination of measurements and model enables the prediction of future sky positions that can be compared with future observations. The requisite stellar astrometric precision is provided by the Gaia Mission, while solar system ephemeris models (e.g. Giorgini et al., 1996; Fienga et al., 2014; Folkner et al., 2014; Pitjeva & Pitjev, 2018) combined with occultation measurements provide ever-stricter constraints on the projected positions of foreground asteroids over time.

The aforementioned Gaia Mission is obtaining accurate sky positions and sky motion models for $N \sim 10^9$ stars with $G < 20.0$, where Gaia's G filter is similar to the standard V-band filter (Gaia Collaboration et al., 2016). For this large aggregate of stars, the original mission specifications mandated that 50% of sources have $G < 19$ and that 2.27% ($\sim$26 million stars) have $G < 14.5$. The actual astrometric precision to which Gaia's observed stars can be fixed varies inversely with brightness, with design expectations ranging from $\omega = 7\,\mu$as at $V = 10$ to $\omega = 12 - 25\,\mu$as at $V = 15$ and $\omega = 100 - 300\,\mu$as at $V = 20$, where the ranges at a given magnitude depend upon the intrinsic brightness of the star (de Bruijne et al., 2005). Gaia Collaboration et al. (2018) have recently published Gaia's Data Release 2 (DR2), an interim compendium of measurements based on 22 months of spacecraft operations.

Most small bodies in the solar system have ephemerides determined from photometric

[2] https://ssd.jpl.nasa.gov/

measurements which, in the absence of radar range measurements, have accuracy tied to the pixel scale, p, of the employed imaging detector.[3] For individual observations, this is typically of order $p = 1''/\mathrm{px}$, corresponding to sky plane displacements of $\sim 3,700\,\mathrm{km}$ at the distance of the Trojan asteroids and $\sim 30,000\,\mathrm{km}$ for objects in the Kuiper belt. When a minor planet occults a star, however, its instantaneous ephemeris is radically improved. An observed blocking of starlight indicates that a chord traversing the line of sight to the object also traversed the star, thereby establishing a time accuracy $\Delta t \sim \delta/v_T$ where δ is the angle subtended by the transit chord and v_T is the angular velocity of the body on the sky. The spatial accuracy on the plane of the sky is of order ω, the astrometric accuracy of the star's position. With the advent of Gaia DR2, the improvement factor from previous Hipparcos astrometric uncertainties is often measured in the thousands (Perryman et al., 1997; Lindegren et al., 2018). A Trojan asteroid observed to occult a bright Gaia star generates a one-time sky plane positional accuracy of just a few tens of meters, comparable to the order-of-magnitude tidal displacement $dx \sim 30\,\mathrm{m}$ induced by a Planet Nine-like object over a half-orbit time scale for the Trojan occulter. This startling potential accuracy is the primary motivator for the investigation presented here.

2.4.2 Occultation Rates

The population of Jovian Trojan asteroids is understood to be extensive, with estimates ranging from $N \sim 1 \times 10^4$ to $N \sim 2 \times 10^5$ objects with $D > 2\,\mathrm{km}$ in the combined L_4 and L_5 swarms (Jewitt et al., 2000; Yoshida & Nakamura, 2005; Fernández et al., 2009; Wong & Brown, 2015). Variations in N stem from assumptions regarding the underlying albedo distribution. The Gaia DR2 data release, moreover, contains of order 3×10^7 stars of apparent brightness $V < 15$ whose sky positions have been determined to accuracy $\omega \sim 20\mu\mathrm{as}$ (de Bruijne et al., 2005). At typical Trojan distance $a = 5.2\,\mathrm{AU}$, this is

[3] https://www.minorplanetcenter.net/

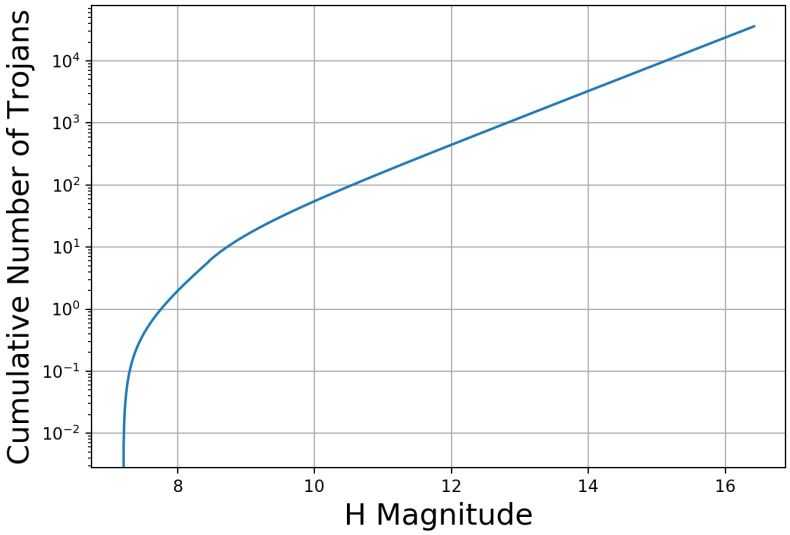

Figure 2.2: Cumulative number of Trojans as a function of H magnitude, estimated from Equation 2.2 using parameters obtained in Wong & Brown (2015).

equivalent to sky-plane positional accuracy $dx \sim 75\,\mathrm{m}$ for the portion of the asteroidal body responsible for the occultation.

To estimate the rate of observable Trojan occultations, we must first assume a Trojan size distribution. We adopt the broken power law differential magnitude distribution

$$\frac{dN}{dH} = \begin{cases} 10^{\alpha_1(H-H_0)} & H < H_b \\ 10^{\alpha_2 H + (\alpha_1-\alpha_2)H_b - \alpha_1 H_0} & H \geq H_b \end{cases} \tag{2.2}$$

and associated best-fitting parameters $\alpha_1 = 0.91$, $\alpha_2 = 0.43$, $H_0 = 7.22$, and $H_b = 8.46$, obtained in Wong & Brown (2015) by fitting the distribution of bright Trojans. The resultant magnitude distribution is shown in Figure 2.2. We combine this magnitude distribution with Equation 2.11 to obtain the corresponding size distribution, where we use the process described in Appendix 2.10 to set albedo values.

We then sum the diameters of all inferred Trojans with $D > 2$ km ($7.2 < H < 16.42$) to obtain a total length 1.4×10^5 km for all 3.6×10^4 large Trojans lined up side-by-side. Given $P \sim 12\,\mathrm{yr}$ orbital periods, each Trojan sweeps out ∼30 degrees yr^{-1}, and the aggregate of Trojans sweeps out projected area ∼0.31 square degrees yr^{-1}. For a uniform distribution of $V < 15$ stars on the sky, this corresponds to $N \sim 225$ potentially observable occultations per year at each location on Earth. Using a network of $N \sim 2000$ small telescopes, approximately 4.5×10^5 events are thus observable annually, some of which may not be unique – i.e. multiple telescopes may observe a single occultation event. Assuming an $f = 0.6$ duty cycle to account for airmass, solar angle, and overlapping targets, as well as an additional $f = 0.4$ duty cycle when considering the day-night cycle, one plausibly expects $\sim 1.1 \times 10^5$ unique occultation events per year on the network, of which a subset would be observable by multiple telescopes.

Focusing on the catalog of known Trojans in the JPL Small-Body Database, we provide a more detailed assessment to compare with the foregoing estimates. The occultation rate for an individual Trojan is

$$r = n_* \sigma v_T, \tag{2.3}$$

where n_* is the areal number density of background stars, σ is the linear cross-section for occultations, and v_T is the angular speed of the Trojan relative to the observer at the time of occultation. For the purpose of assessing rates, rather than planning specific observations, we adopt a simplified solar system model containing only the Earth, the Sun, and the Trojan of interest. We ignore Earth's axial precession and assume that background stars are point sources, which is an appropriate approximation at the scale of this problem. For example, a G2 solar-type dwarf star located at $d = 1000\,\mathrm{pc}$ has $V \sim 15$ and angular size $\delta \sim 9.3\,\mu$as, smaller than the projected end-of-mission $\omega = 12 - 25\,\mu$as astrometric precision of Gaia.

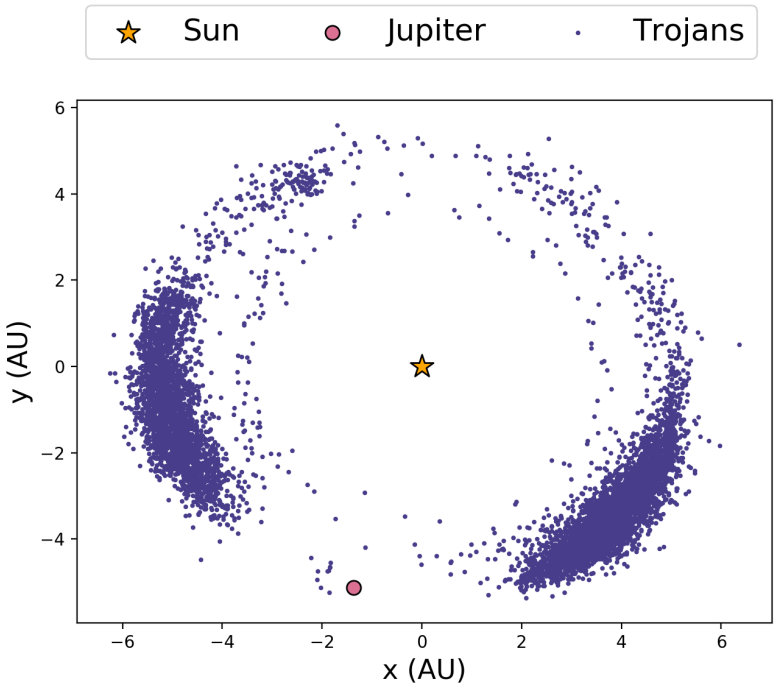

Figure 2.3: Distribution of all modeled Jovian Trojan asteroids at time t=0 in our occultation rate calculation.

To study the frequency of observable occultation events, we simulate the Sun-Earth-Trojan system for each of the 7207 Jovian Trojans listed in the JPL Small-Body Database on November 1st, 2018. The starting spatial distribution of these Trojans relative to Jupiter and the Sun in our simulations is shown in Figure 2.3. Combining the `REBOUND` orbital integrator (Rein & Liu, 2012) with standard coordinate frame transformations, we obtain the RA/Dec coordinates for each Trojan at 2000 equally spaced time steps over 24 years – approximately the duration of two full Trojan orbits. We calculate the expected occultation rate for each Trojan at each time step, with each component of the rate equation obtained using the following procedure:

- n_*: We categorize the total number of stars along the ecliptic plane (to rough approximation, the orbital plane of the Jovian Trojans) into 360 RA bins that are evenly spaced such that each bin encompasses 1° in RA space. We choose to parameterize over RA space in order to encompass regions predominantly passing through the

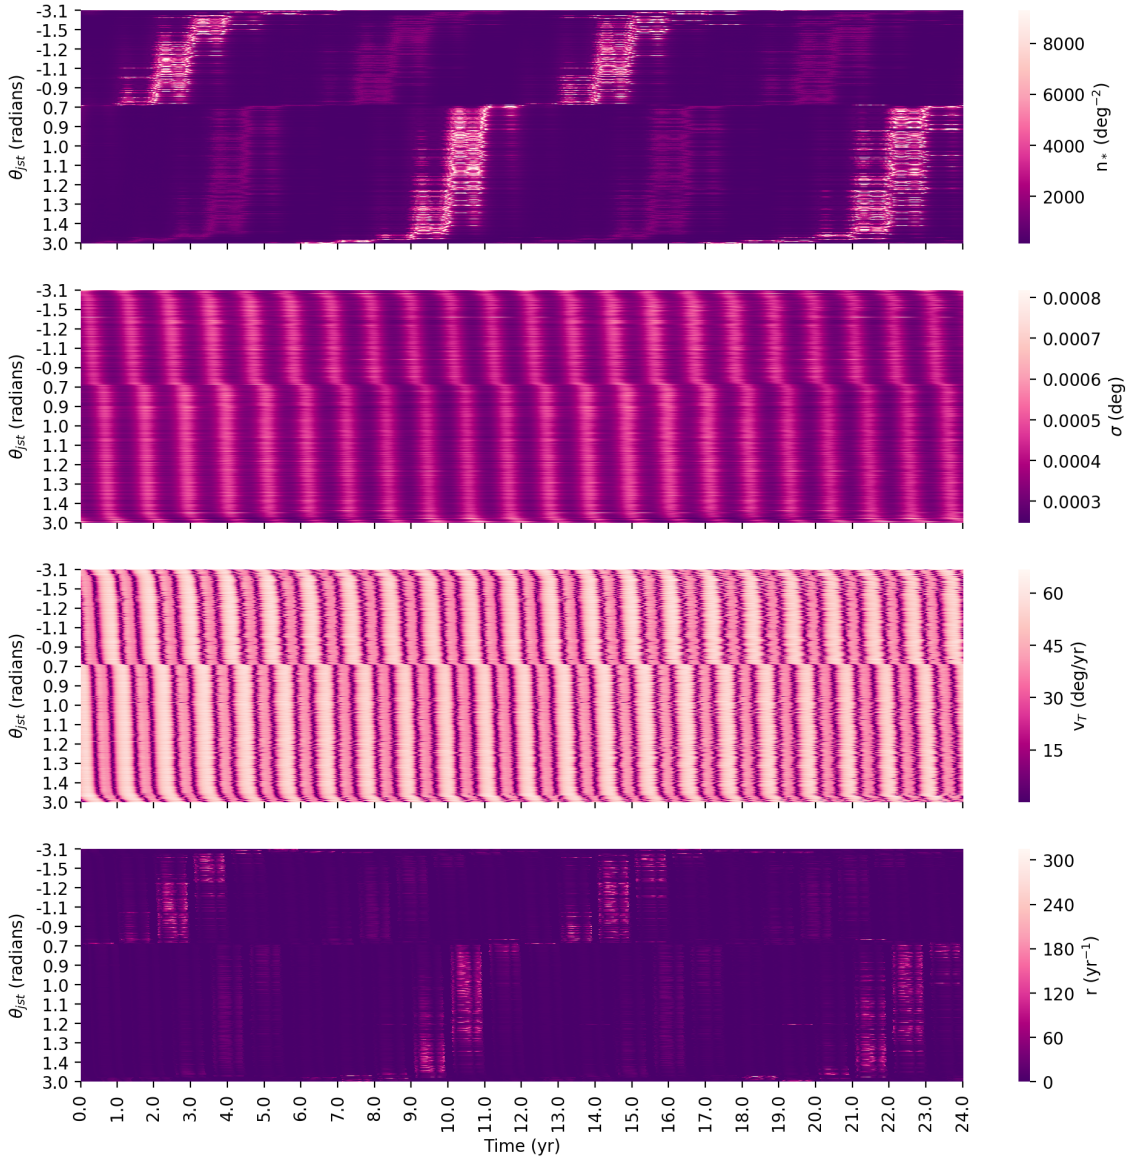

Figure 2.4: n_*, σ, v_T, and occultation rate r displayed as a function of time for all 7207 Trojan asteroids in the JPL Small-Body Database. Trojans are ordered by angle θ_{jst} between the Sun-Jupiter and Sun-Trojan vectors. The distribution of occultation rates in the bottom map consists of the top 3 maps multiplied together, with an additional condition that regions with $\theta_{set} < 38^\circ$ have occultation rate set to zero, where θ_{set} is the angle between the Earth-Sun and Earth-Trojan vectors.

Galactic plane within individual bins.

To accomplish this, we advance Jupiter's orbit at 360 evenly spaced time steps (corresponding to slightly varying step sizes in RA space) using `REBOUND`. Between each set of adjacent RA values, we query the Gaia DR2 database within $\pm 2.5^\circ$ from the midpoint declination value to obtain a star count within each bin. Using these total star counts and the size of each bin on the sky, we calculate the stellar density in each bin. From this point, we linearly interpolate our results to obtain stellar densities in evenly spaced RA bins. We obtain n_* at each time step by selecting the RA bin within which the Trojan falls at that time, then retrieving the corresponding stellar density.

- σ: The effective occultation cross-section for each Trojan can be estimated as the 2545 km latitudinal extent of the continental United States subtended at the instantaneous Earth-Trojan distance. This cross-section takes into account the full range of telescopes extending across the country such that nearly all asteroids with $D \gtrsim 2$ km passing across the United States in the East-West direction can be seen in occultation by one or more network telescopes.

  Among the latitudinal distribution of telescopes, gaps larger than the diameter of the occulter may prevent the detection of an occultation, since small asteroids can slip through these gaps. Few large latitudinal gaps (33 gaps larger than 10 km in extent) exist within the cISP network design that we employ, and the full distribution of gaps is displayed in Figure 2.5. For the initial estimates presented here, we do not explore this in further detail but acknowledge that these gaps may lead to an underestimate of the total occultation rate. If desired, these gaps could be addressed by adding new, strategically placed telescope sites to the current network design, increasing the total number of sites from $N = 1913$ to $N \sim 2000$.

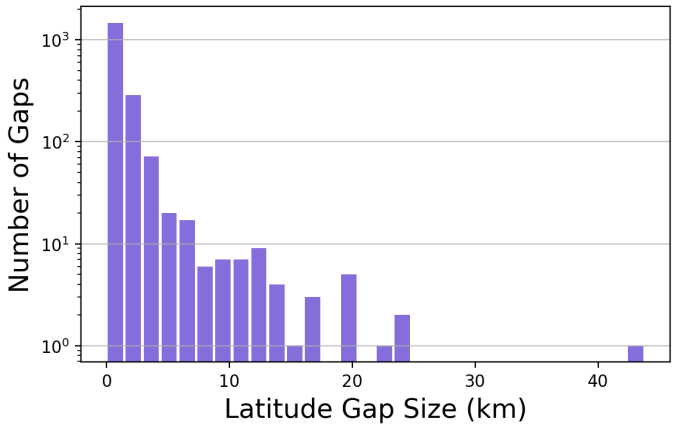

Figure 2.5: Distribution of all latitudinal gaps within the current cISP network setup. These gaps are obtained by ordering all network site locations by latitude and finding the gap size between all adjacent sites in this listing.

- v_T: This is the effective angular velocity of the Trojan, which is given by the absolute value of the sky-plane velocity of the Trojan relative to that of the Earth ($|\mathbf{v}_E$ - $\mathbf{v}_T|$) at each time step.

Our resulting n_*, σ, and v_T values, as well as the corresponding occultation rates r, are displayed in Figure 2.4. All four maps order the Trojans by θ_{jst}, the angle between the Sun-Jupiter and Sun-Trojan vectors, where Figure 2.3 provides the orientation of the full system. We use the convention that negative angles correspond to Trojans lagging behind Jupiter in its orbit (the L5 swarm), while positive angles correspond to Trojans leading Jupiter (the L4 swarm).

The structure in the top panel of Figure 2.4 reflects the periodicity with which the two Trojan swarms pass through the Galactic midplane. Bright maxima in the background stellar density n_* of each swarm occur with a $\sim$12-year periodicity corresponding to the frequency with which they pass through the Galactic midplane in the direction of the central Milky Way. Dimmer local maxima occur 6 years apart from each of these brighter maxima, corresponding to passage through the outer region of the Galactic plane with a lower stellar density. The offset between central peaks in L4 and L5 is approximately 4

years in accordance with the spatial separation of the two swarms along Jupiter's orbit.

The second panel shows the linear cross-section σ of the United States as viewed by each Trojan asteroid as it moves through its orbit. Alternating bright and dark streaks, each with a duration of ~ 6 months, reflect the Earth's periodic movement towards and away from the two Trojan swarms during its orbit.

The third panel displays the sky-plane velocity v_T of each Trojan as a function of time, where dark streaks correspond to times during which the Trojan's motion is perpendicular to Earth's sky-plane (approximately twice per Earth orbit). Between these dark streaks are alternating high- and medium-velocity local maxima. High-velocity local maxima correspond to times during which the Earth and the Trojan move in opposite directions within their orbits (for example, when they are on opposite sides of the Sun), increasing the Trojan's relative velocity as viewed from Earth. Conversely, medium-velocity local maxima occur where the Earth and the Trojan move in the same direction, such as when they are both on the same side of the Sun, reducing the Trojan's relative sky-plane velocity.

Finally, our fourth panel is produced by multiplying together the first three panels and incorporating an additional condition that all regions with $\theta_{set} < 38\,^\circ$ have occultation rate set to zero, where θ_{set} is the angle between the Earth-Sun and Earth-Trojan vectors. This final condition accounts for solar angle and airmass considerations, where we require that, for a Trojan to be observable during the night, (1) the solar angle must be -8° or lower relative to the horizon, slightly below civil twilight (Nichols & Powers, 1964), when the Trojan is in the sky, and (2) the asteroid must be at airmass $z < 2$, or placement in the sky at least 30° above the horizon. We assume that all Trojans are observable from the continental United States due to the orientation of their orbits roughly along the solar system ecliptic plane.

The overall structure of occultation rates in Figure 2.4 follows that of n_*, with overlaid dark

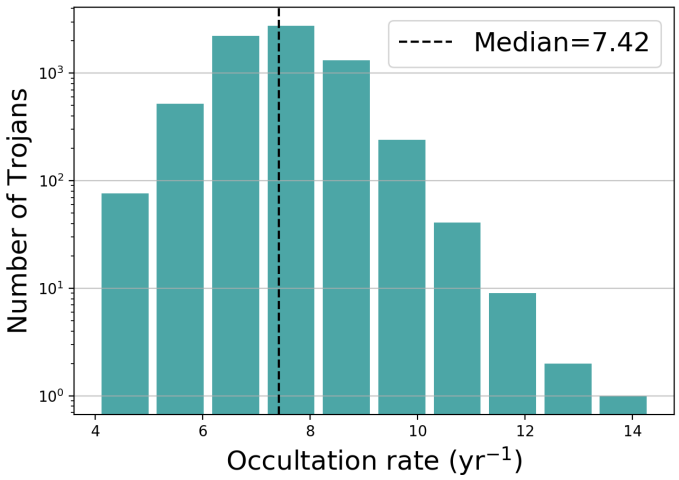

Figure 2.6: Average occultation rates obtained for each Trojan over 2000 sampled time steps spanning 2 years. The median number of total occultations per year is 7.42.

streaks reflecting the influence of σ and v_T. Longer-lasting dark streaks are superimposed as well due to our requirement that $r = 0$ where $\theta_{set} < 38\,^\circ$.

From this distribution, we extract a median rate of 7.42 occultations per Trojan per year, with the rate distribution displayed in Figure 2.6. Extrapolating to the full Jovian Trojan population with $D > 2$ km (N $\sim 3.6 \times 10^4$), this corresponds to $\sim 2.7 \times 10^5$ total occultations per year. Some of these events will occur during the daytime and thus will not be observable. Approximating that, on average, 9/24 hours per day can be spent observing, we estimate a total of 1.0×10^5 observable occultation events per year, in excellent agreement with our previous estimate. This is a sufficiently high occultation rate to regularly track a large number of Trojan asteroids in search for signatures of a massive perturber. We envision that, in conjunction with the proposed program to track Trojan asteroids, the network will be kept continually busy with observations of various small-body occultations throughout the solar system.

Although this estimate does not account for weather conditions, we argue that, for all but the smallest targets, the extended nature of this network will mitigate weather problems.

Table 2.2: 3794 Sthenelos

Parameter	Value	Uncertainty
a	5.204 AU	1.246e-7 AU
e	0.147	6.766e-8
i	0.106°	1.163e-7°
Ω	5.990°	9.742e-7°
ω	6.215°	1.052e-6°
M_0	5.355°	5.122e-7°
D	34.531 km	0.364 km

An occulting body will cross numerous network sites as it traverses the continent, making it unlikely for weather conditions to prevent observations by *all* telescopes along the path. In particular, the characteristic Rossby parameter ($\beta = 2\omega \cos\phi/a$, where ϕ is the latitude, ω is Earth's angular rotation speed, and a is Earth's radius) at mid-latitudes generally ensures nonuniform weather conditions across the country. For the smallest targets, which occult only 1-2 sites, however, weather conditions will play a larger role.

2.4.3 Orbital Element Refinement

A set of three occultation measurements over a time baseline of several months or more permits, in idealized principle, the determination of all six osculating Keplerian orbital elements to the precision of the measured sky-plane positions (Bernstein & Khushalani, 2000). To numerically evaluate the precision to which we can pin down the orbit of a Trojan with a given number of occultation measurements, we simulate a series of successive occultations and study their effect on refining the orbital elements of the occulter. We report results for medium-sized Jovian Trojan 3794 Sthenelos from the JPL Small-Body Database, with properties summarized in Table 2.2.

We first assume that the fiducial reported orbit from the JPL Small-Body Database is the 'true' orbit. We select N_{occ} points along the orbit at which we simulate occultations, and

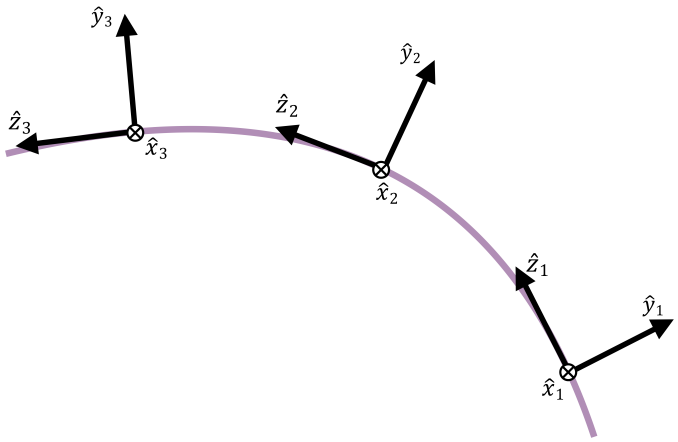

Figure 2.7: Schematic of coordinate frame transformations for three occultations. Each $\hat{x}$ is in the observer-Trojan direction, $\hat{z}$ is in the direction of the Trojan's sky-plane velocity, and $\hat{y}$ is in the mutually perpendicular direction.

for each of these points we obtain a separate coordinate transformation with $\hat{x}$ along the Earth-Trojan line of sight and $\hat{y}$ and $\hat{z}$ in the sky-plane as measured from Earth. We set $\hat{z}$ in the direction of the Trojan's sky-plane velocity, while the perpendicular sky-plane direction is set as $\hat{y}$ (see Figure 2.7).

The precision to which an asteroid's position can be determined within the sky-plane is dependent upon two factors: (1) the astrometric uncertainty associated with the background star and (2) the network's observing cadence. Occultations do not provide a constraint in the $\hat{x}$ direction for any individual measurement; information from multiple events, however, can be combined to provide constraints in all three directions.

In the $\hat{y}$ direction, the positional uncertainty is equivalent to the occulted star's astrometric uncertainty. In our model, we set all astrometric uncertainties to 20 μas, corresponding to the projected end-of-mission Gaia astrometric performance for a star in the range $V = 12 - 15$ mag (Gaia Collaboration et al., 2016; de Bruijne et al., 2005). This is equivalent to a sky-plane positional uncertainty $\sigma_y \sim 75$ m at a distance 5.2 AU; in our simulation, we set uncertainties at each occultation based on the fiducial Earth-Trojan distance at that time.

In the along-track direction within the sky-plane, $\hat{z}$, the asteroid's positional uncertainty is limited by the observing cadence needed to obtain the requisite signal-to-noise ratio (SNR) for the occultation event. We aim for an observing cadence sufficiently high to reach the limits of the Gaia astrometric precision in the along-track direction.

Integrating over visible and near-infrared wavelengths from 400-800 nm, $\sim$4900 photons from a $V = 15$ star strike the aperture of a 16-inch telescope each second. To achieve 20 μas precision at $d = 5.2$ AU, we must observe at 225 fps, which results in 21 photons per frame before accounting for inefficiencies in instrumentation. Achievement of this observing cadence is thus permitted by physics and will hinge on the available equipment. Frame rates for CCD and CMOS sensors are limited by the effective number of pixels used by the camera (e.g. El-Nozahi et al., 2009). Because high resolution is not required for our purposes, in which each telescope observes a single star at a given time, we propose that the requisite precision can be attained by binning pixels to meet the appropriate frame rates needed for each star. As a result, we set our uncertainty in the $\hat{z}$ direction to 20 μas as in the $\hat{y}$ direction.

To quantify the orbital element improvement with improved positional constraints, we employ a Markov Chain Monte Carlo (MCMC) analysis using an affine-invariant ensemble sampler from the `emcee` Python package (Foreman-Mackey et al., 2013). We consider a set of 5 occultations measured in 1-year intervals, directly reflecting the expected orbital improvement over a 5-year baseline across which Planet Nine's perturbational signal should be observable. We note that we have adopted a conservative occultation measurement rate of 1 yr^{-1}, and in practice it therefore should be possible to constrain the orbital elements to even higher precision with more frequent measurements. For comparison purposes, we also consider the case in which 5 occultation measurements are spread evenly across 12 years, corresponding to $\sim$ 1 full Trojan orbit.

In each case, we use 100 walkers to sample 6 parameters (the 6 orbital elements), initialized

with Gaussian priors where the fiducial orbital elements and associated uncertainties are set as the mean and standard deviation, respectively. At each step of the MCMC, we transform every proposed set of orbital elements into the coordinate frame of each occultation to find where that set of parameters places the asteroid at the measured occultation time. Then, models are accepted or rejected based on the log likelihood function in Equation 2.4, where x_n and σ_n represent the fiducial asteroid position and associated uncertainty at the nth occultation, respectively, while μ_n represents the model position. We sum over the $\hat{y}$ and $\hat{z}$ directions for each occultation using associated positional uncertainties σ_y and σ_z.

$$\mathcal{L}(\mu|x,\sigma) = -\frac{1}{2}\sum_{i=1}^{n}\left(\frac{(x_n-\mu_n)^2}{\sigma_n^2} + \ln(2\pi\sigma_n^2)\right) \tag{2.4}$$

We run this process for 1,000 iterations, and we find that the walkers stabilize after ∼200 iterations. We discard our first 800 iterations – well beyond the point of walker stability – as our "burn-in" and draw posteriors only from the final 200 iterations. Our results are displayed in Figure 2.8. The posterior distribution demonstrates that 5 occultation measurements across 5 years can constrain each orbital element to fractional uncertainty $\sigma/\mu \sim 1.5 \times 10^{-9}$ or better. With occultation measurements that are more spread out in time, the width of each posterior shrinks further, reflecting the exquisite precision to which orbits can be determined from occultation measurements over an extended temporal baseline.

2.4.4 The Dynamical Effect of Planet Nine vs. the Kuiper Belt

By incorporating a two-dimensional model of the Kuiper belt into the EPM2017 (Ephemerides of the Planets and Moon – 2017 version) planetary ephemerides, Pitjeva & Pitjev (2018) found that perturbations to the giant outer solar system planets due to the presence of the Kuiper belt are comparable in magnitude to those induced by a $10M_\oplus$ planet at a

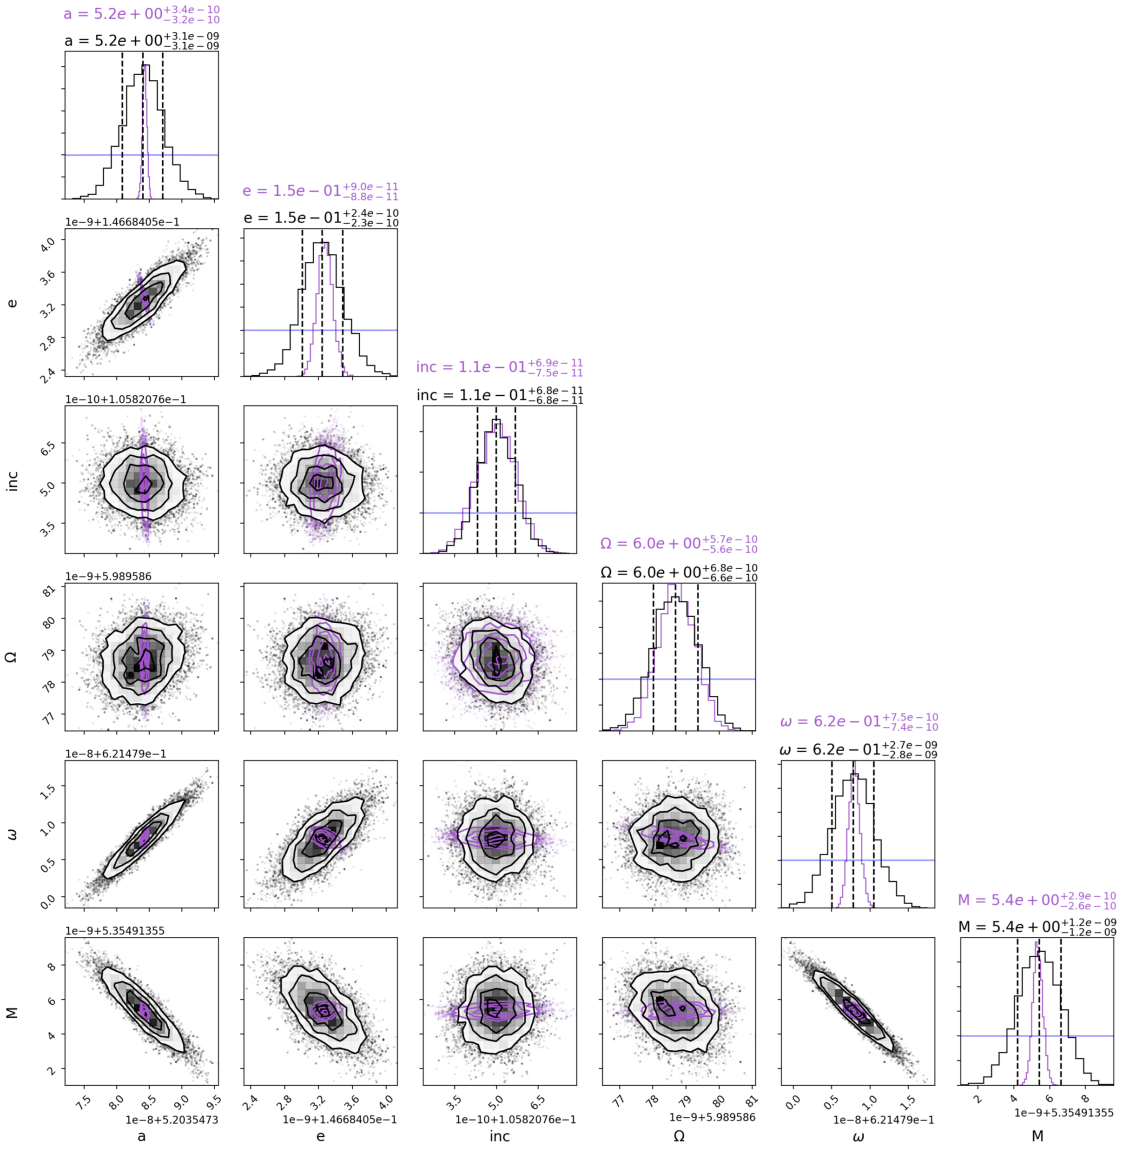

Figure 2.8: Corner plot showing the distribution of walkers after burn-in for each orbital element after 5 occultations taken with a 1-year spacing. The 16th, 50th, and 84th percentiles of the marginalized posterior probability distribution functions are displayed as dotted lines. In purple, we also include the posteriors where 5 occultations are instead taken at evenly-spaced intervals across the ∼12-year Trojan orbit. The blue lines in each histogram display the Gaussian distributions associated with the initial orbital element uncertainties prior to any occultation measurements. The original uncertainties were all improved by over an order of magnitude; as a result, the blue distributions appear flat on this scale. Plot created using the `corner` Python package (Foreman-Mackey, 2016).

solar distance $d = 550$ AU. Here, we investigate how the two sources of influence can be disentangled by tracking their respective perturbations on Jupiter's Trojan asteroids.

We examine the change in each orbital element with time due to the presence of an external perturbing body. We complete this study in several steps. First, we compare analytical expressions for the force applied to the Trojan due to a ring and due to a point source (§2.4.4.1). We then follow the formulation detailed in Murray & Dermott (1999) (originally derived in Burns (1976)), which derives analytical expressions for the derivative of each orbital element due to an external perturber (§2.4.4.2). Lastly, we simulate the effect of a point source vs. ring perturber in `REBOUND` to assess net perturbations over an integrated orbit of the Trojan asteroid (§2.4.4.3). Throughout these analyses, we use values listed in Table 2.1 and set $\Omega = \omega = 0$ and $M = \pi$ to place Planet Nine at apoastron in its orbit.

2.4.4.1 Forces Imparted by an External Perturber

Evaluation of the instantaneous force imparted by each external perturber illuminates the comparative influences from Planet Nine and from the Kuiper belt, which we approximate as a 2D ring of uniform material. Planet Nine can be treated as a point mass (m_{P9}), and thus the acceleration, or force per unit mass imparted upon another point mass a distance d from Planet Nine, is $a_{\rm P9} = -Gm_{\rm P9}/d^2$.

We consider two collinear configurations for the Trojan: one in which it is at the farthest point in its orbit from Planet Nine (Figure 2.9: top), and another in which it is at its close approach to Planet Nine (Figure 2.9: bottom). The acceleration imparted on the Trojan is a_{min} = 2.01 x 10^{-11} cm/s^2 in the former configuration and a_{max} = 2.06 x 10^{-11} cm/s^2 in the latter.

For comparison, we model the Kuiper belt as a uniform circular ring with radius extending from R_{min} to R_{max}, where the angle between Planet Nine and any point in the Kuiper belt is given by θ (see Figure 2.10). The acceleration imparted upon our test Trojan by the

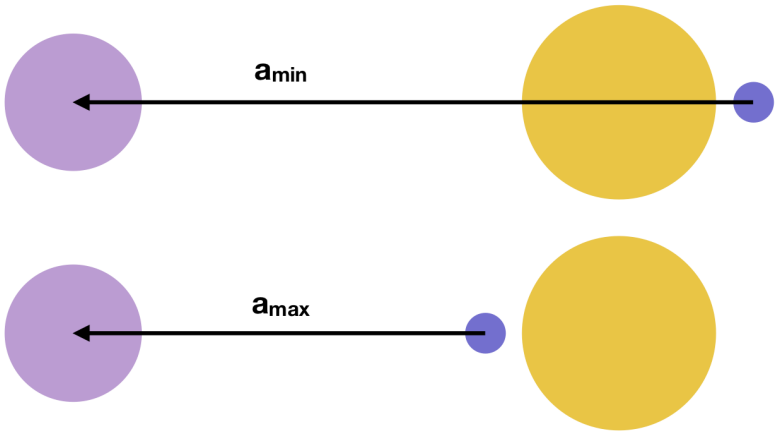

Figure 2.9: Orbital orientations corresponding to the minimum (top) and maximum (bottom) acceleration imparted upon the test Trojan by Planet Nine. Planet Nine (purple), the Trojan (blue), and the Sun (yellow) are shown with logarithmically scaled relative sizes and distances.

2D Kuiper belt is given by Equation 2.5, where we use the law of cosines to integrate the force per unit mass along the full angular and radial extent of the ring. Here, σ_{KB} denotes the 2D mass density of the Kuiper belt. The location of the Trojan is given as $(r, 0)$ in polar heliocentric coordinates, where we define the axis representing $\theta = 0$ along the line connecting the Sun and the Trojan.

$$a_{\mathrm{KB}} = -G\sigma_{KB} \int_{R_{min}}^{R_{max}} \int_{0}^{2\pi} \frac{R(r - R\cos\theta)}{(r^2 + R^2 - 2rR\cos\theta)^{3/2}} d\theta dR \tag{2.5}$$

We place our test Trojan at $r = 5.2$ AU and apply Equation 2.5 to study the instantaneous perturbation to the Trojan imparted by the Kuiper belt. We obtain surface density σ_{KB} by dividing the total mass of the Kuiper belt used in Pitjeva & Pitjev (2018), m = 1.97 $\times$ $10^{-2}\mathrm{M}_{\oplus}$, by the full surface area of our 2D disk. We integrate Equation 2.5 with limits R_{min} = 39.4 AU and R_{max} = 47.8 AU to emulate the disk model from Pitjeva & Pitjev (2018), resulting in a_{KB} = 1.10 $\times$ 10^{-10} cm/s^2. This is approximately an order of magnitude larger than the acceleration obtained from Planet Nine. Due to the symmetry of the Kuiper belt,

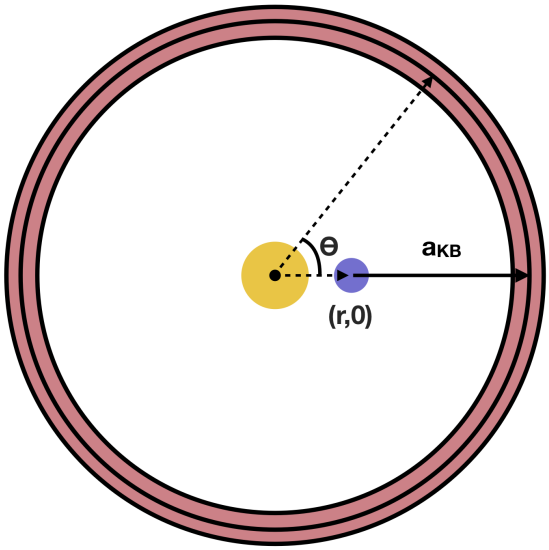

Figure 2.10: Kuiper belt model used with Equation 2.5 to integrate acceleration imparted by the ring. We integrate over all θ to obtain the total a_{KB} = 1.10 x 10^{-10} cm/s^2 in the direction of the nearest Kuiper belt edge.

the magnitude of this acceleration is unaffected by the starting location of the Trojan, though the direction changes throughout the Trojan orbit. This result underscores the findings of Pitjeva & Pitjev (2018) and motivates further study to disentangle the perturbational effect of Planet Nine from that of the Kuiper belt.

2.4.4.2 Analytical Perturbations to Orbital Elements

To explore this problem in greater detail, we also consider the effect of each structure on the orbital element evolution of our fiducial Trojan (Roy, 2005). We consider a force $F = \bar{R}\hat{r} + \bar{T}\hat{\theta} + \bar{N}\hat{z}$ imparted by a perturber with mass M_p, where $\bar{R}$, $\bar{T}$, and $\bar{N}$ are components of force in the $\hat{r}$, $\hat{\theta}$, and $\hat{z}$ cylindrical coordinate directions, respectively. We set all mutual inclinations to zero and thus have $\bar{N} = 0$. Zero force applied in the normal direction necessitates that $\dot{I}$ and $\dot{\Omega}$ must be equivalently zero. Expressions for all other orbital element derivatives, following the formulation of Murray & Dermott (1999), are provided in Equations 2.6-2.10. We report changes in τ, the time of perihelion passage,

rather than M_0, where the two are related by $M_0 = -n\tau$.

$$\mu = GM_p \tag{2.6}$$

$$\dot{a} = 2\frac{a^{3/2}}{\sqrt{\mu(1-e^2)}}[\bar{R}e\sin f + \bar{T}(1+e\cos f)] \tag{2.7}$$

$$\dot{e} = \sqrt{\frac{a(1-e^2)}{\mu}}[\bar{R}\sin f + \bar{T}(\cos f + \cos E)] \tag{2.8}$$

$$\dot{\omega} = \sqrt{\frac{a(1-e^2)}{\mu e^2}}\Big[-\bar{R}\cos f + \bar{T}\sin f\frac{2+e\cos f}{1+e\cos f}\Big] \tag{2.9}$$

$$\begin{aligned}\dot{\tau} = \Big[&3(\tau-t)\sqrt{\frac{a}{\mu(1-e^2)}}e\sin f \\ &+a^2(1-e^2)\mu^{-1}\Big(\frac{-\cos f}{e}+\frac{2}{1+e\cos f}\Big)\Big]\bar{R} \\ &+\Big[3(\tau-t)\sqrt{\frac{a}{\mu(1-e^2)}}(1+e\cos f) \\ &+a^2\mu^{-1}(1-e^2)\Big(\frac{\sin f(2+e\cos f)}{e(1+e\cos f)}\Big)\Big]\bar{T}\end{aligned} \tag{2.10}$$

We first study the effect of a single point mass perturber to characterize the gravitational impact of Planet Nine. We neglect the tangential and normal components of the force ($\bar{T} = \bar{N} = 0$) and focus solely on Planet Nine's radial force. For our purposes, this simplified model is sufficient to obtain an estimate of the expected change in orbital elements over time, since Planet Nine's large predicted separation ensures that its radial force is substantially stronger than forces in the other two reference directions. We calculate

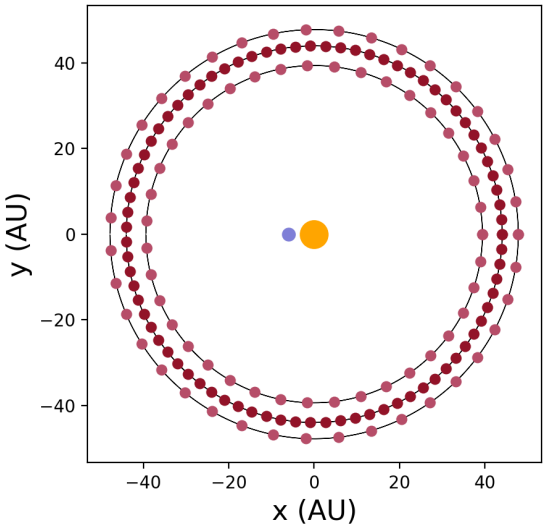

Figure 2.11: Discretized model of the Kuiper belt used throughout §2.4.4.2 and §2.4.4.3. A sample Trojan location is shown in blue, while the Sun is in orange (not to scale). There are 80 point masses encompassing the central ring of the Kuiper belt model (dark red), while there are 40 in each the inner and outer rings (light red).

$\bar{R}$ using Planet Nine's apoastron distance from the Sun. We report results for Trojan 3794 Sthenelos, again using the orbital elements from Table 2.2. The resultant orbital element derivatives are provided in Table 2.3.

We also complete this exercise for a circular, 2D ring in the Trojan's orbital plane to obtain the instantaneous change in orbital elements induced by the Kuiper belt. This model includes a tangential force but again has zero normal force. Following the procedure of Pitjeva & Pitjev (2018), our model of the Kuiper belt includes 3 rings, where the inner and outer ring are located at $R_{min} = 39.4$ AU and $R_{max} = 47.8$ AU, respectively, corresponding to the 3:2 and 2:1 orbital resonances with Neptune. The inner and outer ring each contain 40 discrete, evenly spaced point masses. The central ring contains 80 point masses and is located at $r = 44.0$ AU. The total mass of the Kuiper belt, set as 1.97 x 10^{-2} $M_\oplus$ (Pitjeva & Pitjev, 2018), is evenly distributed among these 160 points comprising the Kuiper belt. The model described here is visualized in Figure 2.11, and the instantaneous changes in orbital elements resulting from the Kuiper Belt are listed in Table 2.3.

	Planet Nine	Kuiper Belt
$\dot{a}$	-8.50e-10	1.69e-7
$\dot{e}$	-5.44e-10	1.65e-9
$\dot{\omega}$	-4.77e-9	-3.82e-7
$\dot{\tau}$	-3.47e-9	-1.06e-6

Table 2.3: Instantaneous change in orbital elements induced by Planet Nine and by the Kuiper belt. All values are given in units of the orbital element per year, where a is in AU and ω is in radians.

The magnitude of the instantaneous change in each orbital element is significantly smaller in the case of Planet Nine than in the case of the Kuiper belt. However, a continuous force in a single direction – from the effectively static Planet Nine at apoastron – should have a different effect over time from a ring imparting a symmetric force throughout the Trojan's orbit, where the effect would in part cancel out. We explore this further in §2.4.4.3.

2.4.4.3 Numerical Integration with `REBOUND`

We next compound the perturbative effects of both Planet Nine and the Kuiper belt on a single Trojan using the `REBOUND` orbital integrator. We consider two test systems and compare the perturbations upon the Trojan in each: (1) a system including only the Sun, the test Trojan, and Planet Nine, and (2) a system including the Sun, the test Trojan, and the same model of the Kuiper belt described in §2.4.4.2 (Figure 2.11).

All components of each model are treated as point masses, and we integrate over one full orbit of the test Trojan, tracking perturbations from the initial Keplerian orbit over time. For simplicity, we set all inclinations in both models to zero such that all bodies lie within the same orbital plane. We note that our analogous tests which incorporate the inclination of Planet Nine ($\sim$20°) behave qualitatively in the same way within the plane of perturbations ($\Delta r \times \Delta \phi$) that we consider here.

For consistency, we again provide results for Trojan 3794 Sthenelos, with properties listed

in Table 2.2 with the exception that we set inclination $i = 0^\circ$. We also run the same REBOUND simulations for several other Trojans of various sizes and confirm that there are no qualitative differences in results.

The results of our integrations are shown in Figures 2.12 and 2.13. We use the convention $\Delta\phi$=$\phi_p - \phi_T$ and $\Delta r = r_p - r_T$, where the p subscript refers to the system with a perturber, while the T subscript refers to the system including only the Sun and the Trojan, with no external perturber. Thus, a positive value for Δr, as is the case for the Kuiper belt where $e_T = 0$, indicates that the perturber has pushed the Trojan outwards in its orbit.

In Figure 2.12, we display the perturbations induced by Planet Nine for a range of initial angles between the Sun-Planet Nine and Sun-Trojan vectors (θ_{pst}). The perturbation induced by the Kuiper belt is not clearly visible on this scale and is thus omitted. At the beginning of each orbit, Planet Nine pushes the Trojan asteroid radially outwards (positive Δr) for $\theta_{pst} < \pi/2$ and radially inwards (negative Δr) for $\theta_{pst} > \pi/2$. $\Delta\phi$ is set by the location of the Trojan in its orbit: positive as the Trojan moves towards Planet Nine, and negative as it moves away. Thus, all four quadrants of this space are populated by perturbations, with the perturbative direction set by the Planet Nine-Sun-Trojan orientation.

However, the *differential* perturbation induced by Planet Nine – the measurable perturbation from within the solar system – is substantially smaller than the *global* perturbation shown in Figure 2.12. This is because Figure 2.12 does not incorporate the acceleration also imparted upon the Sun by Planet Nine.

In Figure 2.13, we subtract the Sun's motion from the system to show the displacement of the Trojan relative to the Sun. To elucidate the behavior of the Trojan in multiple scenarios, we display the perturbations produced in two cases: (1) where the Trojan eccentricity e_T is set to zero (left and central panels in Figure 2.13), and (2) where we use the fiducial, nonzero Trojan eccentricity extracted from the JPL Small-Body Database (right panel in

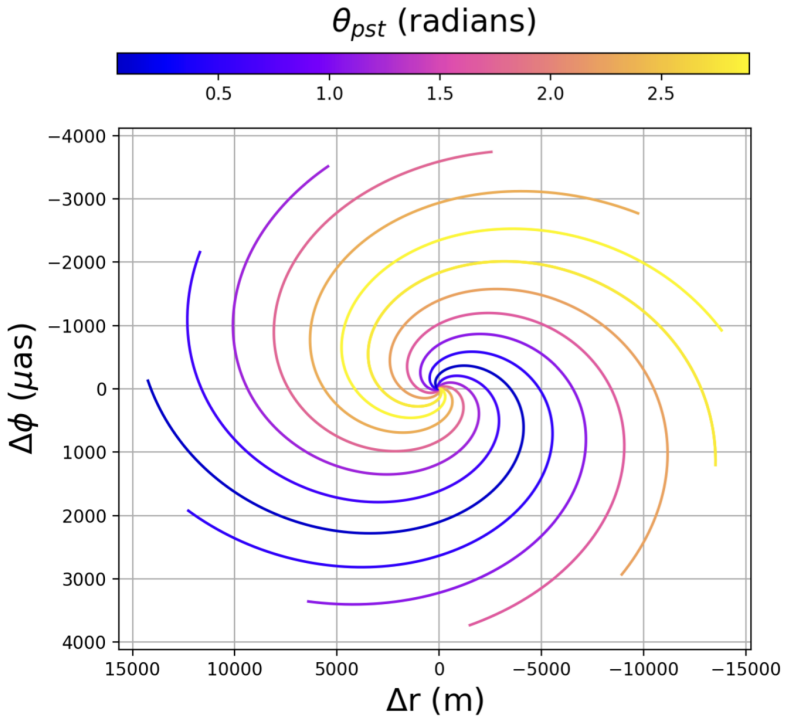

Figure 2.12: Perturbations induced by Planet Nine, given in spherical coordinates. Each line originates at (0,0) and traces out perturbations over one full Trojan orbit. The color bar indicates different starting points in the Trojan's orbit, where θ_{pst} provides the Planet Nine-Sun-Trojan angle at the start of the integration.

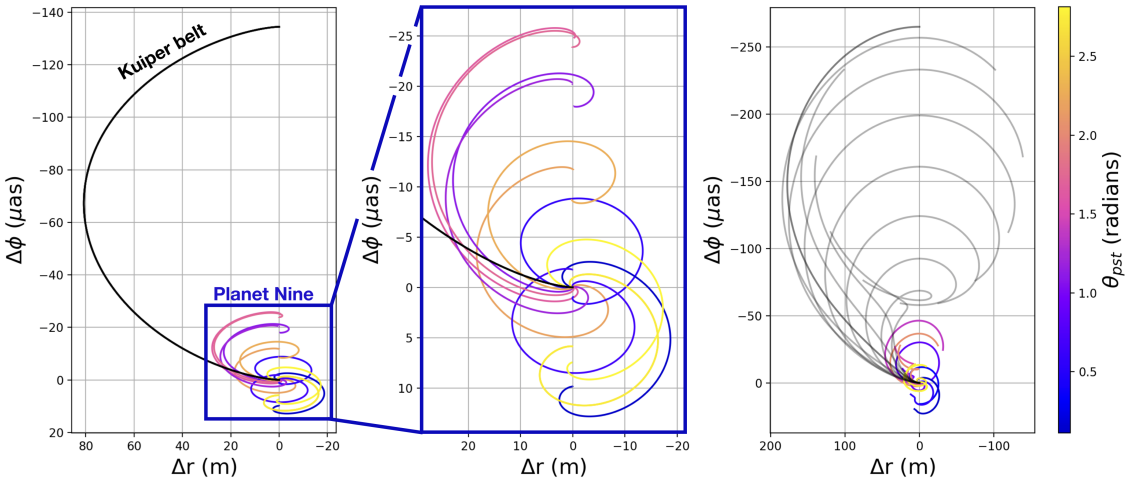

Figure 2.13: Perturbations induced by Planet Nine (in color) and by the Kuiper belt (in gray), with the Sun's motion subtracted from the system. As in Figure 2.12, each line originates at (0,0) and traces out perturbations over one full Trojan orbit, while the color scale provides the starting Planet Nine-Sun-Trojan orientation. Results are shown for the case where $e_T = 0$ for the Trojan orbit (left and corresponding central panel) and where $e_T = 0.147$ from the JPL Small-Body Database (right). In the $e_T = 0$ case (left), which is a direct transformation of Figure 2.12, the Kuiper belt has the same effect irrespective of the initial Trojan location due to the system's symmetry. Incorporating nonzero e_T removes this perturbative symmetry.

Figure 2.13).

For nonzero Trojan eccentricities, the Kuiper belt can produce perturbations in the $+\Delta r$, $-\Delta r$, and $-\Delta\phi$ directions, but never in the $+\Delta\phi$ region of parameter space. In contrast, Planet Nine can produce perturbations in all four quadrants. We obtain similar results when, instead of using the Kuiper belt model in Figure 2.10, we model all but the three most massive TNOs listed in the JPL Small-Body Database (Pluto, Eris, and Makemake). Despite the lack of perfect symmetry in this case, we still find only small ($\sim$5 μas) deviations into the $+\Delta\phi$ space, suggesting that, in the absence of additional undiscovered Pluto-sized objects, the symmetry of the Kuiper belt is sufficient such that Planet Nine should be gravitationally distinguishable from the Kuiper belt. Given the presence of an additional Pluto-sized perturber in the Kuiper belt, this network would provide strong evidence for its existence and constraints upon its location – an interesting result in itself.

It is possible to distinguish the effects of Planet Nine and the Kuiper belt, therefore, by studying occultations of the Trojans in regions of the sky which correspond to positive $\Delta\phi$ for a Planet Nine perturber, but which still result in a negative $\Delta\phi$ from the Kuiper belt. These include regions in which the Trojan is moving towards Planet Nine in its orbit and thus is pulled further forward by the perturber, as illustrated in Figure 2.14. The perturbational signature of the putative Planet Nine will deviate most from the noise produced by the Kuiper belt during the half of the Trojan orbit in which perturbations are in the $+\Delta\phi$ direction. The peak of this $+\Delta\phi$ deviation occurs when $\theta_{pst} = \pi/2$ and the Trojan is moving towards Planet Nine. As a result, the angular orientation of Planet Nine can be extracted upon detection of the predicted perturbation from a distribution of Trojans.

The perturbation size detectable by our network is limited by the angular precision of the occulted star (ω=12-25 μas at $V = 15$ from the Gaia catalog). This is much more precise than standard astrometric measurements, and even astrometry with HST typically results in uncertainties of order $\sim$1 mas or greater (e.g. Bellini et al., 2011; Porter et al., 2018) and can only reach precision ~ 0.1 mas for the very brightest targets (Benedict et al., 2002). Thus, the extremely high precision necessary to observe Planet Nine's perturbational effect can be decisively obtained through the occultation method.

By measuring a large number of Trojan positions over time and searching for this $+\Delta\phi$ signature of a point perturber, we can extract the Planet Nine signal and separate it from that of the Kuiper belt. For an individual small Trojan, we have shown that ephemeris refinement from successive occultations through the course of a full orbit generates a final positional accuracy of order $\delta\mathbf{x} = 75\,\mathrm{m}$. Figure 2.13, on the other hand, suggests that systematic deflections on the Trojan orbit due to Planet Nine are of order 25 meters, corresponding to a detection of $\sigma_i = 1/3\,\sigma$ significance. As a consequence, the tracking of a single Trojan is no more useful for detecting Planet Nine than was the Cassini telemetry. With a large occultation network, however, thousands of bodies can be monitored, with

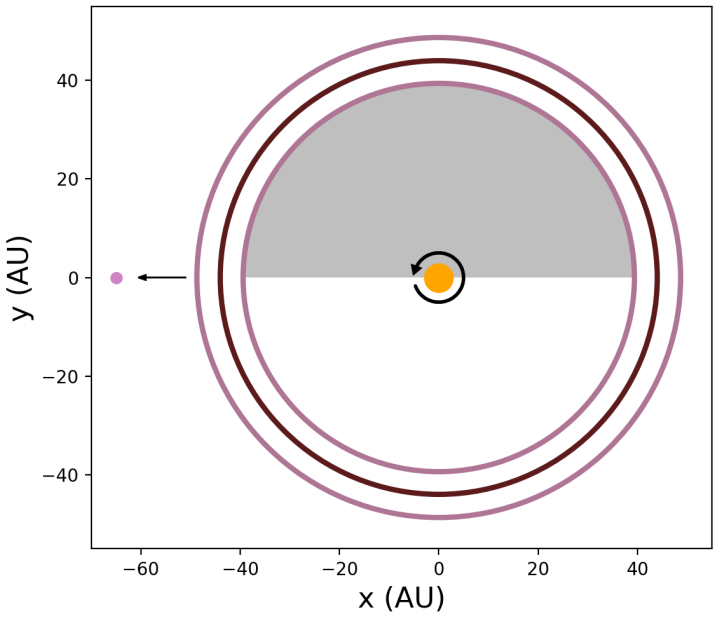

Figure 2.14: By observing Trojans only in the half of the sky shaded in gray, where $\Delta\phi$ from Planet Nine is positive, it is possible to differentiate perturbations from Planet Nine from those of the symmetric Kuiper belt, which always correspond to negative $\Delta\phi$. The direction of the Trojan orbit and of Planet Nine's position to the left are displayed for orientation, while the Kuiper belt is represented by 3 concentric rings.

detection significance rising to $\sigma_{\rm final} = \sigma_i\sqrt{N}$, where N is the number of Trojans whose orbits have been improved to the DR2-enabled precision. As an example, refinement of $N \sim 225$ Trojan orbits would generate $\sigma_{\rm final} = 5\sigma$ detection significance.

2.5 Additional Applications

The construction of the proposed cISP occultation network would have considerable potential not only in the search for Planet Nine, but also in a wide variety of other applications. We showed that regular stellar occultation measurements would deliver unprecedented constraints on the orbital ephemerides of Trojan asteroids. This concept can be extended to various other small-body populations within the solar system, including near-Earth objects,

main-belt asteroids, and TNOs, delivering unprecedented constraints on the distribution of minor body shapes and sizes within the solar system. Such applications can be utilized in a variety of ways.

In this section, we describe a selection of additional scientific use cases for the network. We focus on four sample applications: (1) synergies with planned NASA missions (§2.5.1), (2) investigating the origins of Jovian and Neptunian Trojans (§2.5.2), (3) using TNO occultations to develop an acceleration map of the outer solar system (§2.5.3), and (4) independently verifying asteroid diameters throughout the solar system (§2.5.4).

2.5.1 Synergies with Future NASA Missions

2.5.1.1 The Lucy Mission

As an added benefit of the search for Planet Nine, improved constraints on the orbital ephemerides of the Jovian Trojan population may assist NASA's *Lucy* Mission, which is planned for launch in October 2021 and will fly by seven asteroids – including six Jovian Trojans – from 2025-2033 (Levison et al., 2016; Levison et al., 2017). Improving the positional accuracy of the Jovian Trojans can potentially enhance the safety of the spacecraft, conserve fuel, and improve trajectory planning. Occultation measurements can also probe for rings and other opaque structures in the Trojan systems prior to spacecraft arrival.

Moreover, a combination of *in situ* range observations from the spacecraft and near-simultaneous Jovian Trojan positional measurements obtained via occultations would generate fully three-dimensional positions of extraordinary accuracy for the handful of Jovian Trojans studied in-depth by *Lucy*. These exquisite 3D positional measurements would be highly complementary to the larger statistical sample that we propose studying throughout this work.

2.5.1.2 The Psyche Mission

The *Psyche* Mission is planned for launch in 2022 and will arrive at Main Belt asteroid (16) Psyche in 2026, where it will orbit for 21 months (Oh et al., 2016; Lord et al., 2017). Equipped with an X-band telecommunications system and a multi-spectral imager, the mission spacecraft will map the gravity field of Psyche while also characterizing its topography (Hart et al., 2018).

By concurrently measuring occultations of Psyche as it is studied by the orbiting spacecraft, it will be possible to not only constrain the asteroid's position to exquisitely high precision (as for the Trojans studied by *Lucy*) but also to simultaneously obtain 2D projected spatial maps of the asteroid from multiple viewing points. Synergistic measurements from both an occultation network and the orbiting spacecraft will provide extremely detailed constraints on the asteroid's three-dimensional topography.

2.5.2 Origins of the Jovian & Neptunian Trojans

Trojan asteroids, introduced in the Jovian context in Section 2.4.1, are more generally defined as small bodies librating at the L4 and L5 Lagrange points of a planet-star system in a 1:1 mean motion resonance, first theorized by Lagrange (1873) as a stable solution to the restricted three-body problem. Wolf (1906) discovered the first known Trojan asteroid at Jupiter's L4 point, and the vast majority of Trojan asteroids discovered since then are also located in the Jovian system.

More recently, a population of Trojans was also found in the Neptunian system (Chiang et al., 2003; Sheppard & Trujillo, 2006, 2010a,b; Becker et al., 2008), with a comparable number of asteroids at each Lagrange point and similar dynamics in both populations. Previous works have explored possible capture mechanisms for the Jovian and Neptunian Trojans (e.g. Marzari et al., 2002; Morbidelli et al., 2005; Sheppard & Trujillo, 2006;

Lykawka et al., 2009), where most of the recent models incorporate capture from the same population that now comprises the modern-day Kuiper belt, as in the Nice model (Tsiganis et al., 2005). However, the exact capture mechanism of these two populations remains to be definitively established.

Sheppard & Trujillo (2006) and later Jewitt (2018) observed that the color distribution of the Jovian Trojans is statistically indistinguishable from that of the Neptunian Trojans – consistent with a shared origin for both populations. However, Jewitt (2018) found that this distribution does not match that of any dynamical sub-population within the Kuiper belt, where these sub-populations are defined in Peixinho et al. (2015). This discrepancy refutes the hypothesis that the two Trojan populations were captured as scattered KBOs, raising the question of where else the Trojans may have instead originated.

A robust comparison of asteroid size distributions would serve as another useful distinguishing factor, since the parent population should have a similar size distribution to that of the captured Trojans. The H magnitude distribution of the large Jovian Trojans, which is closely tied to the radius distribution, has been measured on several occasions and matches a steep power law with index α ($\alpha = 0.9 \pm 0.2$ (Jewitt et al., 2000), $\alpha = 1.0 \pm 0.2$ (Fraser et al., 2014), $\alpha = 0.91 \pm^{+0.19}_{-0.16}$ (Wong & Brown, 2015)); however, the paucity of bright Trojans limits the accuracy to which α can be determined. The translation of an H magnitude distribution to a radius distribution also requires assumptions about the albedos of the Trojans, which could be circumvented through direct size observations using stellar occultations. Currently, the ranges of asteroid sizes measured for the Trojan and Kuiper belt populations barely overlap – thus, precise constraints on both α and the size distribution of small KBOs would enable a more convincing direct comparison.

2.5.3 An Acceleration Map of the Outer Solar System

Another appeal of this network lies within its ability to pinpoint the positions of distant minor planets, such as TNOs, to high precision. While challenging, stellar occultation measurements of TNOs have been successfully demonstrated on several occasions (Ortiz et al., 2017; Müller et al., 2018; Sickafoose et al., 2019) and have become increasingly accessible with the advent of Gaia DR2. Deep sky surveys such as the Dark Energy Survey (Dark Energy Survey Collaboration, 2005) can also provide astrometric measurements of TNOs with sufficient accuracy to reliably predict upcoming occultation events (Banda-Huarca et al., 2019). By observing TNO occultations over time, it is possible to constrain the allowed parameter space for undiscovered minor planets in the outer solar system, leading to an effective "acceleration map." As new extreme TNOs with $a > 250$ AU are discovered, such a map will contextualize the distribution of known objects and constrain their abundance.

Furthermore, occultation measurements provide the added benefit of determining TNO shapes and sizes. Stellar occultation measurements are currently the only ground-based observational technique capable of constraining the geometry of TNOs to kilometer accuracy. However, as described in Appendix 2.11, the Fresnel diffraction limit becomes important for small TNOs and must be taken into consideration when planning TNO occultation measurements.

2.5.4 Asteroid Diameter Verification

Occultation measurements provide an independent test to verify previously obtained asteroid diameters, as well as tighter constraints to decrease the uncertainties in reported values. A combination of measurements from the WISE (Wright et al., 2010) and NEOWISE (Mainzer et al., 2011) infrared surveys has been modeled using the Near-Earth Asteroid

Thermal Model (NEATM; Harris (1998)), resulting in a database of measured diameters for ∼150,000 asteroids (Mainzer et al., 2016). The methods and underlying assumptions used by the NEOWISE team to obtain these asteroid diameters were recently challenged (Myhrvold, 2018a,b,c), with a corresponding rebuttal in Wright et al. (2018). Independent measurements obtained through separate methods can improve confidence in reported values obtained through thermal modeling, and this could serve as an external test of the accuracy of previously measured values.

2.6 Camera & Telescope Selection

The availability of funding will play a deciding role in the specific magnitude limits of this network, and we envision that financial support for the network would potentially leverage private-public partnerships. In this section, we discuss desired properties for the network and provide cost estimates based on several currently operating camera systems optimized for occultation observations. While we focus on hardware configurations and their associated expenses, we note that operating costs are nontrivial and would play a material role in the final project budget.

A large fraction of the total construction cost will be encompassed by camera expenses. We require a camera capable of observing at a high cadence with minimal dead time between frames in order to obtain the highest possible time resolution and corresponding spatial resolution. The field of view should ideally encompass at least $10'$ to include a sufficient number of comparison stars for calibration. To push our magnitude limit and timing cadence as low as possible, we benefit from a camera with the highest possible quantum efficiency, maximizing the percentage of incoming photons that are converted to detected electrons. For faint objects occulting faint stars, the dominant noise source will likely be read noise; as a result, an optimal detector will minimize this quantity.

For comparison, we consider the hardware setup used by the New Horizons team in preparation for the MU69 flyby. This configuration was sufficient to obtain five chords of varying lengths across MU69, all consistent with the bi-lobed structure observed during the flyby. To measure an occultation over a $V \sim 15$ star, the New Horizons team used a set of Sky-Watcher 16" truss-tube Dobsonian reflectors (∼$2000 each) paired with QHY174M-GPS CCD cameras (∼$1240 each with cooling).[4] These CCD cameras have 78% quantum efficiency and can run between 138-490 fps depending on the user's pixel binning. Two thousand sets of this telescope-CCD pairing would amount to $6.5 million. While this particular configuration is not endowed with robotic controls, switching the telescope model for an automated version with the same photon collecting area should provide similar results. Potential alternatives include 16-inch automated telescopes such as the Meade LX600-ACF 16" (f/8) Cassegrain telescope ($19,000 per unit, including tripod mount) or the Meade LX200-ACF 16" (f/10) computerized telescope ($13,500 per unit, without mount).

We also consider the possibility of adapting previous iterations of optimized occultation camera systems for our purposes. We focus on four recent systems: the Portable Occultation, Eclipse, and Transit System (POETS; Souza et al., 2006); the Portable High-Speed Occultation Telescope (PHOT; Young et al., 2011a); the Portable Instrument for Capturing Occultations (PICO; Lockhart et al., 2010) and the Astronomical Digital Video occultation observing and recording System (ADVS; Barry et al., 2015). We provide a brief overview of each of these CCD systems here, and we refer the reader to Table 3 in Barry et al. (2015) for a detailed comparison of their properties.

POETS was developed from 2005-2006 as a portable camera system small enough to fit into two standard carry-on luggage bags. With frame rate ∼35 frames per second (fps), high quantum efficiency (≥90% from 500-700nm), and low read noise (<1-49 e^- at 10 MHz),

[4] https://www.skyandtelescope.com/astronomy-news/solar-system/observers-track-new-horizons-next-target/

POETS has been used for a wide variety of observations ranging from exoplanet light curve observations (Winn et al., 2009; Adams et al., 2011) to studies of Pluto's atmosphere (Bosh et al., 2015) and total solar eclipse observations (Pasachoff, 2009). POETS uses a frame-transfer electron-multiplying CCD for which dead time between frames is nearly zero and limited only by the time required to shift accumulated charge to the storage region of the CCD (typically on the order of ms). Frame-transfer CCDs allow for high-cadence observations with high time resolution, making them suitable for occultation measurements. They are, however, relatively expensive, and thus we do not explore this option further.

The PHOT system has very similar specifications to those of POETS and has been used to study several stellar occultations of Pluto and Charon (Young et al., 2008, 2010, 2011b; Olkin et al., 2014). Like POETS, the PHOT system also uses an expensive frame-transfer CCD, so we do not discuss it in further depth here.

The PICO system, which is significantly less expensive than POETS and PHOT ($5,000 vs. $38,000 and $30,000, respectively) has read noise $\sim$13.6 e^- at 2.8 MHz and lower quantum efficiency peaking at 58%. It does not employ a frame-transfer CCD and has significant inter-frame dead time ranging from 0.65-0.76 seconds per frame, resulting in total observing cadence 1.63 fps for a sample observation of KBO 762 Pulcova (Lockhart et al., 2010). With average shadow speeds $\sim$17 km/s for Jovian Trojans and $\sim$25 km/s for KBOs, this cadence limits occultation detections to objects with minimum diameter 55 km and 82 km, respectively, where at least two in-occultation observations are required for an individual detection. Though this system would be sufficient to observe large bodies, we desire a shorter observing cadence to probe a large number of Jovian Trojans, many of which are significantly smaller than this diameter limit.

The final and most recently deployed system that we discuss here is the ADVS, which uses the GX-FW-28S5M-C Grasshopper Express CCD camera with frame rate 26 fps, quantum efficiency 68%, read noise 9.6 e^-, and inter-frame dead time $<$2 ms. The full system costs

$4,500, and, with frame rate 26 fps, it would be possible to probe objects with diameters down to $\sim$1.3 km for Jovian Trojans and $\sim$1.9 km for KBOs. While this performance is sufficient to observe Trojans down to our size limit, a camera with a higher frame rate would be necessary to reach the limits of the Gaia astrometric precision for $V = 15$ stars (225 fps or better to achieve 20 μas precision at 5.2 AU).

Beyond these previously implemented configurations, which all employ CCD systems, thermoelectrically cooled CMOS cameras provide another promising alternative for low-cost, low-read noise photometry at a high frame rate, with several versions currently available reaching $\geq$225 fps. Notably, two sets of a CMOS camera/11" telescope configuration ($16,000 for each set) were recently used to observe an occultation of a 1.3 km KBO, where both cameras successfully observed the event (Arimatsu et al., 2019). An in-depth exploration of the CMOS option would be warranted upon implementation of this network. However, we choose not to explore this option further here due to the substantial technological improvements that will likely be made prior to the time of network construction.

Because the cISP network adapts only already-existing communication towers, infrastructure expenses for this project would be minimized. In addition to the telescope/camera systems themselves, automated telescope covers would be necessary to protect and efficiently deploy all units. These expenses will constitute a small fraction of the budget by comparison with the driving costs of the telescope/camera systems employed throughout the grid. We estimate that, for bulk purchases, telescope housing and similar construction expenses will amount to $\sim$$1500 per telescope, or $3 million in total for all 2000 telescopes.

2.7 Network Multiplicity

We have presented this network under the assumption that one telescope would be placed at *every* location throughout the cISP network; here, we also consider the case in which

telescopes are placed not at all possible sites, but only at some subset. For the following reasons, we do not explore this possibility outside of the context of the proposed network site distribution:

1. Finding and setting up new sites that are not part of a pre-existing framework would substantially increase infrastructure expenses. As a result, the starting overhead cost of the network would be much higher when including sites that have not been pre-specified.

2. The quasi-random spread of telescopes across the country decreases the likelihood that the full network will ever be down at the same time due to poor weather. Thus, a more compact setup of telescopes is disfavored.

We acknowledge that an analogous network with a different set of sites and similar quasi-random spread of telescopes could achieve the same goals outlined here; however, an exploration of alternative network configurations is beyond the scope of this work.

A reduction in the number of network telescopes could substantially reduce the price of the network; in fact, it is likely that certain sites may not be amenable to astronomical observations due to light pollution. However, we argue that most, if not all, sites should be maintained within the network in order to maximize its potential for a range of scientific cases.

A smaller number of sites incorporated into the network would increase the average latitudinal gap between telescopes, meaning that occultations would more frequently slip through gaps in the network. This disadvantage would most strongly affect small ($\sim$km-size) asteroids, drastically reducing the total number of consistently observable occultation events. Individual asteroids would also become more difficult to regularly track due to this reduced efficiency.

Beyond increasing the likelihood that any *individual* occultation event is observable, a

large number of telescopes also increases the likelihood that a single occultation event may be observable at multiple sites. Observations of a single occultation event by multiple telescopes provide tighter constraints upon the size of an individual object while also improving confidence in low-SNR occultation measurements. Thus, auxiliary science cases, such as constraining the size distribution of asteroid populations, benefit substantially from the use of a large number of telescopes. Since we propose this network not only as a way to find Planet Nine, but also as a tool to better characterize the distribution of small bodies in the solar system, we choose to present it as a whole such that, upon implementation, it would be as useful as possible for a range of applications.

2.8 Conclusions

In this paper, we discussed the scientific merits of a national network of telescopes designed to observe occultations of solar system asteroids over background stars. We focused on an in-depth analysis of the network's capabilities as a probe for the proposed Planet Nine, where we use Jovian Trojan asteroids as trackers of the tidal differential acceleration imparted by Planet Nine. In summary, we find the following:

- Each Jovian Trojan asteroid occults a $V \leq 15$ Gaia star with a median rate of 7.42 occultations per Trojan per year, corresponding to $\sim$1.0$\times 10^5$ total observable events per year when accounting for the day/night cycle. These rates require both solar angle below -8° and airmass $z < 2$ while the Trojan is in the sky. In practice, some Trojans will occult background stars more frequently, while others will occult less frequently. This overall rate is sufficient to track many individual Trojans over a timescale of a few years in search for the gravitational signature of Planet Nine.

- We quantify the improvement in orbital element uncertainties over several occultations, showing that the precise positional constraints from occultation measurements

directly translate to exquisite constraints on the asteroid's orbital elements. We find that 5 occultations observed over 5 years constrain each orbital element to fractional uncertainty $\sigma/\mu \sim 1.5 \times 10^{-9}$ or better, while 5 observations evenly spaced over a full ~ 12-year Trojan orbit provide even tighter constraints.

- It is possible to distinguish the gravitational perturbing effects of an encircling, (nearly) symmetrically distributed Kuiper belt from that of Planet Nine by strategically observing Jovian Trojans within the half of their orbit in which they are moving towards the effectively stationary location of Planet Nine. This is due to the differing gravitational signature of a circular, approximately symmetric system – the Kuiper belt – by comparison to a point mass – Planet Nine.

We also outlined several further science cases for the network. In summary, we find that occultation measurements by the proposed network can (1) provide complementary observations to support and supplement the NASA *Lucy* and *Psyche* missions, (2) deliver direct, robust measurements to compare the Jovian and Neptunian Trojan size distributions as a test of solar system formation models, (3) permit the development of an outer solar system acceleration map through high-precision measurements of TNOs, and (4) provide an independent test of previously measured asteroid diameters.

Throughout this work, we focus on one configuration for a telescope grid in which all locations are tied to sites along the cISP low-latency communications network. Under the assumption that the latitudinal distribution of telescopes is not substantially altered, the same concept of a large-scale occultation network could, in theory, be implemented with deviations from these specifications without significant changes to the scientific case for the network. In particular, reducing the size of the largest latitudinal gaps present in the current telescope distribution – either by adding/adjusting telescope sites within the current configuration or by using an altogether separate configuration to implement the same concept – would improve the rate of successful occultation measurements for

primarily East-West occulting asteroids.

Similar projects may also consider leveraging the availability of university observatories and small telescopes worldwide. Though this would lead to a non-uniformity of telescope systematics, incorporating these facilities would expand the baseline of possible observations. In any of these forms, a large-scale occultation network would serve as a powerful and timely tool to utilize synergies between Gaia and LSST, demonstrating the remarkable potential of small telescopes in the upcoming era of astronomy.

2.9 Acknowledgements

We are grateful to Konstantin Batygin, Andy Szymkowiak, and Frank van den Bosch for helpful comments. We also thank the anonymous referee for constructive feedback that led to a substantial improvement of the manuscript. M.R. is supported by the National Science Foundation Graduate Research Fellowship Program under Grant Number DGE-1752134. This material is also based upon work supported by the National Aeronautics and Space Administration through the NASA Astrobiology Institute under Cooperative Agreement Notice NNH13ZDA017C issued through the Science Mission Directorate. We acknowledge support from the NASA Astrobiology Institute through a cooperative agreement between NASA Ames Research Center and Yale University. Simulations in this paper made use of the `REBOUND` code which can be downloaded freely at `http://github.com/hannorein/rebound`. This research also made use of the `numpy` (Oliphant, 2006; Walt et al., 2011), `matplotlib` (Hunter, 2007), `astropy` (Astropy Collaboration et al., 2013; Price-Whelan et al., 2018), `pandas` (McKinney et al., 2010), `emcee` (Foreman-Mackey et al., 2013), `corner` (Foreman-Mackey, 2016), and `Altair` (VanderPlas et al., 2018) Python packages.

2.10 Appendix: Assigning Jovian Trojan Diameters

Where diameters are not listed in the JPL Small-Body Database, we convert the reported absolute magnitudes (H) to diameters using Equation 2.11, derived from Equations 2.12 and 2.13 where F refers to flux, m is magnitude, I is intensity, and Ω is the solid angle over which I is collected.

$$D = \frac{2(1\,\mathrm{AU})}{\sqrt{p_v}} 10^{\frac{1}{5}(m_\odot - H)} \tag{2.11}$$

$$\frac{F_1}{F_2} = 10^{-\frac{2}{5}(m_1 - m_2)} \tag{2.12}$$

$$F = I\Omega \tag{2.13}$$

In Equation 2.11, $m_\odot = -26.74$ is the absolute magnitude of the Sun, and H and p_v are the absolute magnitude and geometric albedo of the small body, respectively. Albedo values have not been measured for the vast majority of Jovian Trojans; however, Fernández et al. (2003) observed a mean albedo of 0.056 ± 0.003 for 32 Jovian Trojans with diameter $D \geq 25$ km. Later, Fernández et al. (2009) reported a mean albedo 0.121 ± 0.003 for a separate sample of 44 small (5 km$< D < 24$ km) Jovian Trojans. Accordingly, we assign albedo $a = 0.056$ to all Jovian Trojans and, for those objects which return $D \leq 25$ km, we reassign albedo $a = 0.121$.

2.11 Appendix: Diffraction Effects

The Fresnel diffraction limit must be taken into consideration for observations of small-body populations employing stellar occultations. This limit quantifies the size scale of occulting objects below which diffraction effects become non-negligible. The Fresnel diffraction limit is given in Equation 2.14, where λ is the wavelength of observation and d is the distance from the telescope to the occulting object.

$$F = \sqrt{\lambda d/2} \tag{2.14}$$

Objects with radius $r \sim F$ display wavelength-dependent diffraction fringes that are irregular for objects with rough edges. At visible wavelengths ($\lambda = 700$ nm), the Fresnel scale for Jovian Trojans at distance $d = 5$ AU is $F = 0.5$ km, while $F = 1.5$ km for TNOs at a distance $d = 45$ AU. We anticipate that the Fresnel limit may need to be taken into account in the case of our smallest targets, particularly in the case of TNOs and other distant targets with relatively large F. Fresnel-Kirchoff diffraction theory can be applied in such cases to model objects with $r \sim F$.

2.12 Appendix: (Non)-Detectability of Dark Matter

When a massive body moves through a sea of less massive bodies, it experiences a loss of momentum and kinetic energy through a process called dynamical friction, which produces a trailing wake of concentrated density behind the massive body. Hernandez (2019) recently proposed that the sun's motion through the Milky Way should produce a wake of dark matter particles detectable through gravitational perturbations within the solar system. For estimated dark matter particle dispersion σ = 200 km/s and density $\rho_0 = 0.01 M_\odot$ pc^{-3} (Read, 2014), we use Equation 2 of Hernandez (2019) to find that the tidal acceleration of

the dark matter wake at 10 AU would be of order $a \sim 2 \times 10^{-19}$ cm/s^2, corresponding to a radial shift of roughly 1 cm over a 10-year timescale. This small shift is undetectable with our proposed network.

Part II

The Planet-Minor Planet Connection

Chapter 3

A Direct Shift-Stacking Search for Planet Nine with the Transiting Exoplanet Survey Satellite

Originally published as:

Rice, M. & Laughlin, G. 2020, The Planetary Science Journal, 1, 81

3.1 Abstract

We present results from a new pipeline custom-designed to search for faint, undiscovered solar system bodies using full-frame image data from the NASA Transiting Exoplanet Survey Satellite (TESS) mission. This pipeline removes the baseline flux of each pixel before aligning and co-adding frames along plausible orbital paths of interest. We first demonstrate the performance of the pipeline by recovering the signals of three trans-Neptunian objects – 90377 Sedna ($V = 20.64$), 2015 BP519 ($V = 21.81$), and 2007 TG422 ($V = 22.32$) – both through shift-stacking along their known sky-projected paths and through a blind recovery. We then apply this blind search procedure in a proof-of-concept survey of TESS Sectors 18 and 19, which extend through a portion of the galactic plane in the Northern Hemisphere. We search for dim objects at geocentric distances $d = 70 - 800$ au in a targeted search for Planet Nine and any previously unknown detached Kuiper belt objects that may shed light on the Planet Nine hypothesis. With no input orbital information, our present pipeline can reliably recover the signals of distant solar system bodies in the galactic plane with $V < 21$ and current distances $d \lesssim 150$ au, and we elaborate on paths forward to push these limits in future optimizations. The methods described in this paper will serve as a foundation for an all-sky shift-stacking survey of the distant solar system with TESS.

3.2 Introduction

The outer reaches of the solar system, at distances $d \geq 70$ au, remain largely unexplored. Individual objects with orbits exterior to Neptune's – known as trans-Neptunian objects, or TNOs – are challenging to detect; owing to the steep, $F \propto 1/d^4$ diminution of reflected flux, only the brightest such objects have been discovered. Indeed, at the time of writing, fewer than 100 detached Kuiper belt objects (KBOs), with perihelia $q \gtrsim 40$ au and no direct

interactions with the known solar system planets, have been discovered.

These objects, however, are of exceptional interest due to the unique window that they provide into the dynamical evolution of the outer solar system. The observed apsidal and nodal alignment of detached KBOs, combined with the unexpectedly high inclinations and eccentricities of several outer solar system small body populations, has led to the proposition that a ninth, super-Earth sized planet, commonly known as "Planet Nine", may reside in the distant solar system (Trujillo & Sheppard, 2014; Batygin & Brown, 2016; Malhotra et al., 2016). Dynamical simulations reveal that, in order to account for this observed alignment among KBOs, Planet Nine would likely possess a large semimajor axis ($a = 400 - 800$ au), a high inclination ($i = 15 - 25\,^{\circ}$), and a substantial eccentricity ($e = 0.2 - 0.5$; Batygin et al., 2019). The best-fit orbital solution for Planet Nine found by Batygin et al. (2019), with $a = 500$ au and aphelion distance $Q = 625$ au, corresponds to magnitude $V = 19.0 - 22.2$, where the exact value is determined by the planet's location along the orbit and its inferred albedo.

Several recent and ongoing searches have placed observational constraints on the properties of Planet Nine. Meisner et al. (2018) applied data from the Wide-Field Infrared Survey Explorer (WISE; Wright et al., 2010) to search for the proposed planet at high galactic latitudes, ruling out a bright planet ($W1 < 16.7$, where the $W1$ bandpass is centered at 3.4 μm) at 90% completeness over 3π radians on the sky. The Backyard Worlds: Planet 9 citizen science project described in Kuchner et al. (2017) has also used the WISE dataset in a more extended Planet Nine search, employing $W1$ along with an additional wavelength band ($W2$, centered at 4.6 μm) and resulting in the discovery of the brown dwarf WISEA J110125.95+540052.8. While it is not specifically designed to search for Planet-Nine-like signals, the Dark Energy Survey (DES; Dark Energy Survey Collaboration, 2005) covers 5000 square degrees in the southern sky and is sensitive to dim Planet Nine-like signals at optical and near-infared wavelengths. From the survey's first four years on-

sky, the DES collaboration reported the discovery of hundreds of TNOs, including some relevant to the Planet Nine hypothesis (Bernardinelli et al., 2020; Becker et al., 2018). Indirect, gravitational searches have also been pursued (Fienga et al., 2016; Holman & Payne, 2016a,b); however, these searches require a long time baseline of precise positional measurements across many gravitational probes to distinguish the effects of Planet Nine from those induced by the Kuiper belt (Rice & Laughlin, 2019a).

Planet Nine remains elusive. If it exists, the most promising places to look might now be the regions of the sky with the highest stellar density, where a slowly moving point source is particularly difficult to pinpoint and isolate. A bright and compelling signal lost in the noise of the galactic plane would not be unprecedented; for example, the binary brown dwarf system Luhman 16, only 2 pc from the Sun, remained undiscovered until 2013 as a consequence of its proximity to the galactic plane (Luhman, 2013).

We present results from a systematic search for objects in the outer solar system ($a = 70 - 800$ au) using a custom shift-stacking pipeline designed for use with full-frame images (FFIs) from the Transiting Exoplanet Survey Satellite (TESS; Ricker et al., 2015). The basic idea of the underlying shift-stacking technique – also called "digital tracking", "de-orbiting", or "pencil-beam surveys" in the literature – has been implemented in several preceding works to search for new solar system satellites (Holman et al., 2004; Kavelaars et al., 2004; Burkhart et al., 2016) and TNOs (Gladman et al., 1998, 2001; Bernstein et al., 2004), and a version of it was recently proposed for application to TESS by Holman et al. (2019). Our implementation includes a number of refinements that are specific to finding particularly distant solar system objects in the TESS dataset.

For the purposes of this study, we focus on Sectors 18 and 19, which lie directly along the galactic plane. Our focus on this region is motivated by two factors. First, based on the most recent parameter updates provided by Batygin et al. (2019), the most likely remaining parameter space for Planet Nine – and, specifically, the parameter space that is most poorly

constrained by other surveys such as Pan-STARRS (Kaiser et al., 2002) – lies in the vicinity of the galactic plane (see Figure 25 of Batygin et al. (2019)). If Planet Nine exists in the galactic plane, this would help to explain why it has not yet been discovered, since stellar contamination severely limits optical searches in this region of the sky. Thus, by focusing on Sectors 18 and 19, which encompass much of the galactic plane in the northern sky, we complete a targeted search aimed at the region in which Planet Nine is most likely to lie.

Second, a survey of Sectors 18 and 19 allows us to quantify the performance of our pipeline in a noisy region of the sky that has previously been difficult to study with single-frame exposures. We demonstrate that shift-stacking is a promising method to search for dim outer solar system objects using all-sky surveys, strengthened by its ability to recover sources that would otherwise be lost in the noise due to stellar crowding in single-frame exposures. While previous studies have employed the TESS dataset for solar system science applications (Pál et al., 2018; McNeill et al., 2019; Pál et al., 2020), we present results from the first survey designed to detect undiscovered solar system objects in the TESS dataset.

3.3 Data Overview

The TESS spacecraft includes 4 cameras each comprised of 4 CCDs with 2048 x 2048 pixels per CCD. Each pixel spans $21'' \times 21''$ for a combined, total field of view $24^\circ \times 96^\circ$, extending from the ecliptic pole towards the ecliptic plane at each spacecraft pointing. The survey's observing strategy covers most of the sky away from the ecliptic plane, extending down to ecliptic latitude $b \gtrsim 6^\circ$ and spending the most time observing the ecliptic poles (the "continuous viewing zone" that retains coverage as the pointings change). TESS is thus ideally suited to search for high-inclination solar system objects. Due to the TESS camera's point spread function (PSF), 50% of the flux from an object falls within 1 pix^2 of

the object's peak flux location while 90% falls within 4 pix^2.[1]

TESS observations are organized into sectors, each of which corresponds to a single spacecraft pointing. Each hemisphere is spanned by thirteen observing sectors, and TESS spends 27 consecutive days collecting data for each sector. These observations include 30-minute cadence full-frame images over the full field of view, as well as 2-minute cadence "postage stamp" images of individual bright stars.

We take advantage of TESS's extensive sky coverage by searching for dim, slow-moving objects in the calibrated Sector 18 and Sector 19 FFIs. The locations of these sectors in the sky relative to the ecliptic plane, the galactic plane, and the range of expected Planet Nine orbits are displayed in Figure 3.1. The sampled Planet Nine orbital elements were randomly selected from the ranges 400 au $< a < 800$ au, $15\,^\circ \leq i \leq 25\,^\circ$, and $0.2 \leq e \leq 0.5$ with randomly oriented orbital angles. We then used the `PyEphem` software package to translate the selected orbital elements to projected sky locations and to plot those locations over a span of 10,000 years (the limiting time span covered by `PyEphem`).

To quantify the performance of our pipeline, we also recover previously known objects using calibrated FFIs from Sector 5. The difference in flux density based on proximity to the galactic plane is apparent: the average pixel's flux in Camera 1 of Sector 19, which lies along the galactic plane, is a factor of $\sim$1.5 higher than in Camera 1 of Sector 5, which is at a $\sim 30\,^\circ$ separation from the galactic plane. We download all frames, which have been corrected for instrument/detector artifacts, directly from the Mikulski Archive for Space Telescopes (MAST).[2]

[1] https://heasarc.gsfc.nasa.gov/docs/tess/the-tess-space-telescope.html

[2] http://archive.stsci.edu/tess/bulk_downloads/bulk_downloads_ffi-tp-lc-dv.html

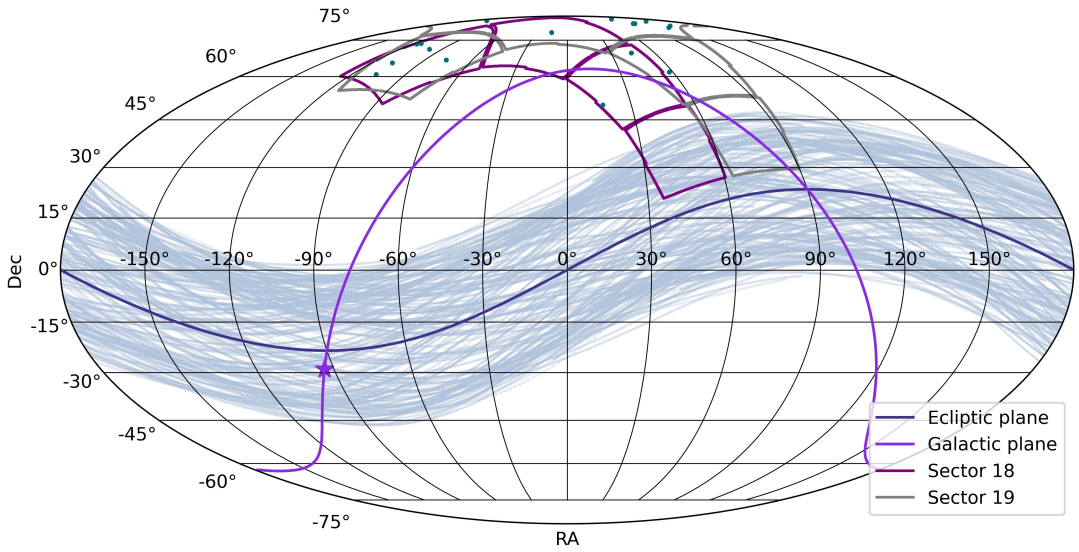

Figure 3.1: Mollweide projection of the TESS field of view in celestial coordinates, with the ecliptic plane shown in navy and the galactic plane in purple. A purple star denotes the location of the galactic center. The regions of the sky observed in Sectors 18 and 19 are delineated in red and gray, respectively. We sample and plot 150 random Planet Nine orbits in light blue to show the overlap between the possible Planet Nine orbits and the region observed by TESS. We target the region of space in which the galactic plane intersects with possible Planet Nine orbits. The candidate outer solar system objects presented in this study are included in green for reference.

3.4 Methods

To search for TNOs in the TESS dataset, we developed a linear shift-stacking pipeline that checks a pre-defined set of potential orbit paths for objects of interest. This pipeline includes three primary components: (1) an algorithm that subtracts the flux baseline on a per-pixel basis, (2) a shift-stacking algorithm that co-adds all baseline-subtracted images along potential TNO paths in search of strong local maxima, and (3) an automated algorithm for extracting candidates. By summing many small deviations from the flux baseline over a large number of exposures, we recover dim objects far below the detection limit of an individual frame.

We divided each TESS CCD into 256×256 pixel cutout regions in a grid that overlaps itself by 128 pixels in either direction. With this partition, each region of the sky interior to the edge of the CCD is present in four different shift-stacking frames. This degeneracy improves the likelihood that any given frame will enclose the full path of an outer solar system object. It also decreases the likelihood that we will miss any given object during our vetting procedure.

We ran our pipeline on the Grace cluster at the Yale Center for Research Computing, using one compute node and one CPU per 256×256 pixel cutout region. The full pipeline described in this section takes 1-3 hours (wall time) to run with the polynomial baseline subtraction, while this time is increased to 7-8 hours with the PCA baseline subtraction. We processed each cutout frame independently; as a result, we were able to run a full sector at a time, processing all cutout regions in parallel.

3.4.1 Baseline Subtraction Algorithms

Our baseline subtraction procedure includes both a pre-processing pixel masking stage (Section 3.4.1.1) and two independent methods for removing the flux baseline, where

both methods fit the baseline on a pixel-by-pixel basis. We use the polynomial baseline subtraction method described in Section 3.4.1.2 to return our key results, and we use the Principal Component Analysis (PCA) baseline subtraction method described in Section 3.4.1.3 as a consistency check. By probing parameter space with two independent data processing procedures, we ensure that only the most robust candidates remain in our final search results.

3.4.1.1 Pre-Fit Pixel Masking

We first use a predetermined mask, set individually for each sector during a preliminary testing phase, to remove a fraction of frames displaying large systematic signals from the full time series. These masks eliminate flux discontinuities in the pixel time series that generally occur at the beginning, middle, and/or end of each TESS sector. In Sectors 18 and 19, the removal of discontinuities at the center of each time series leaves two roughly equal-length $\sim 6 - 7$ day light curves for each pixel, separated by a $\sim 5 - 10$ day gap.

The two time-series masks – one for Sector 18 and another for Sector 19 – were selected using an iterative trial-and-error process with our polynomial baseline subtraction method (Section 3.4.1.2). For each sector, we obtain a mask that removes the fewest possible time frames while still providing a consistently well-performing fit to each of the two light curve segments. We examine the residuals of our polynomial fits in search of asymmetric systematics and alter our masks accordingly.

After this first pass at removing discontinuities, we also remove the 10% of time series frames where the median flux gradient across all pixels has the largest magnitude. In this way, we discard systematics corresponding to rapid flux changes affecting the full cutout region, which introduce additional scatter and degrade the quality of our fits. This allows us to more accurately determine the baseline background flux present in all pixels.

We then determine the maximum flux of each remaining pixel time series and mask out the

10% of pixels that reach the highest flux values. Removal of the brightest pixels eliminates sequences with the highest Poisson noise. A retained pixel will strengthen a detected signal only if the flux observed from the solar system object of interest is greater than the Poisson shot noise from the flux counts within that pixel.

The bright eliminated pixels are typically associated with stars and bright asteroids in the frame that would reduce the signal-to-noise ratio of any detections. If an object passes over one of these masked pixels, it makes no contribution to the signal determined across the full shift-stack. As a result, if a TNO lies directly on top of masked pixels during the majority of a TESS sector, it will likely not be detected by our algorithm. To be detectable, a TNO must cross over masked pixels for a sufficiently small fraction of a TESS sector – defined such that the total summed flux over the unmasked frames produces a $\geq 5\sigma$ signal above zero flux in our aggregate frame.

3.4.1.2 Polynomial Baseline Subtraction

Our first baseline subtraction method uses a low-order polynomial to fit and subsequently subtract off the baseline of each pixel, with flux as a function of time $F(t)$ given by

$$F(t) = k_0 + k_1 t + k_2 t^2 + ... + k_n t^n. \tag{3.1}$$

Here, k_n are constants with values determined in the fitting process. We fit the two halves of the light curve separately and consider polynomials with degree $n_p = 1 - 5$ for each, calculating the reduced χ^2 value,

$$\chi^2 = \frac{1}{(n_t - n_p)} \sum_{i=1}^{i=n_t} \Big(\frac{(F_{\mathrm{obs},i} - F_{\mathrm{fit},i})^2}{\sigma_i^2} \Big) \tag{3.2}$$

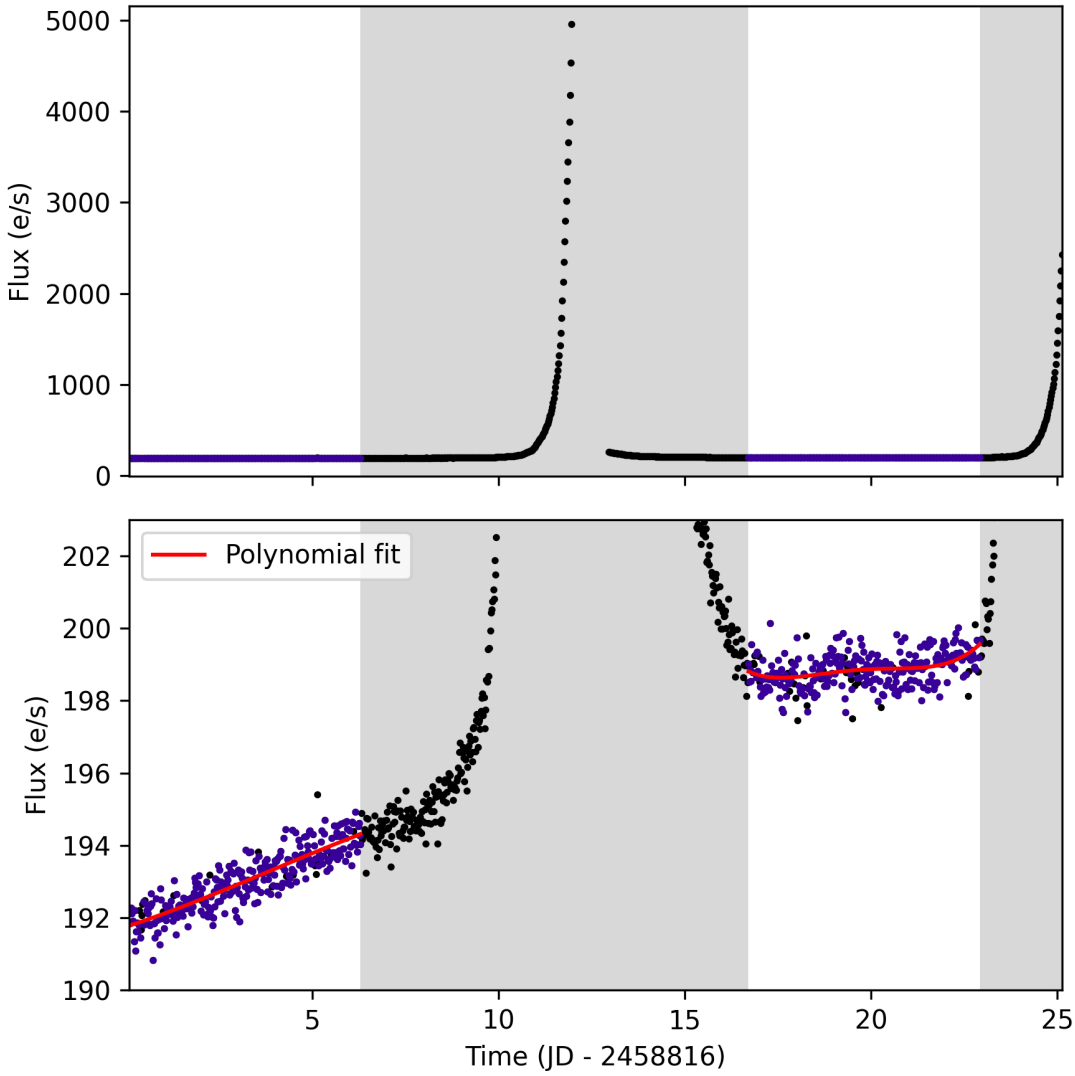

Figure 3.2: Sample polynomial fits to the light curve of a representative Sector 19 pixel. The top panel shows the full flux range covered by the pixel's time series, while the bottom panel zooms in to the regions that we fit after discarding the sloping systematics dominant at the middle and end of the time series. The gray shaded regions are excluded from our analysis, and times included in the baseline fit and subsequent shift-stack are shown in purple. 540 frames are retained after masking. Each light curve component is fit separately with a polynomial shown in red.

for each fit. The total number of frames in the time series is given by n_t, whereas $F_{\mathrm{obs},i}$ and $F_{\mathrm{fit},i}$ are the observed and polynomial fit flux values, respectively, at frame i, and σ_i is the uncertainty reported for the fitted pixel at frame i in the full-frame image FITS file. We iterate over n_p values for each pixel, keeping only the fit that results in the lowest χ^2 value from Equation 3.2. We independently apply this fitting procedure to both light curve segments associated with each pixel. An example fit for a Sector 19 light curve is shown in Figure 3.2.

This baseline subtraction procedure makes no attempt to preserve information regarding the baseline flux of each pixel relative to its neighbors. As a result, if any pixels in the frame still include substantial stellar flux after the removal of bright pixels described in Section 3.4.1.1, the baseline flux of the star(s) in the pixel should also be removed by this process. This helps to prevent spurious "signals" that appear from stacking frames along paths that cross over the locations of stars that have not been fully masked. The polynomial baseline subtraction method assumes a smooth flux profile that can be well-captured by a polynomial function, so it performs more poorly for pixels with substantial variability – for example, those containing a transient source with strong flux variability over a time scale of $\sim$days.

3.4.1.3 PCA Baseline Subtraction

In our second baseline subtraction method, we utilize the `RegressionCorrector` class[3] of the `lightkurve` Python package to estimate the baseline flux profile of each pixel. This method reduces the properties of N regressors, or surrounding pixels, into their constituent components using Principal Component Analysis. Reduction to these principal components removes long-term variability and stochastic noise from the estimated baseline. The PCA components are then combined with the `RegressionCorrector`

[3] See https://docs.lightkurve.org/api/lightkurve.correctors.RegressionCorrector.html for documentation.

to determine the best-fitting baseline correction using linear regression.

For each pixel subtraction, we use the 2000 nearest pixels as regressors. We exclude from this analysis all pixels that lie within 5 pixels of the pixel for which the baseline is being determined. In this way, we ensure that light from an object of interest is not incorporated into its baseline subtraction, reducing the risk of self-subtraction. We use three PCA components; from initial testing, we find that additional components add to the algorithm's computation time without a substantial improvement in performance.

The PCA method described here determines the baseline flux of a pixel based on its neighbors in the frame, meaning that residual long-term signals that remain after our initial masking are not removed by this baseline subtraction. While this reduces the likelihood of self-subtraction for slowly-moving objects such as Planet Nine, it also increases the rate of false positives due to an incomplete subtraction of stellar signals. For this reason, we require that all candidates are first detected using the polynomial baseline subtraction method before verifying their recoverability with the PCA method.

3.4.2 Shift-Stacking

3.4.2.1 Description of the Algorithm

After completing the baseline subtraction process, we input the reduced images into our shift-stacking algorithm to recover TNO candidates of interest. Our pipeline operates in two modes: one designed to recover known objects along a known path, and another designed to implement a blind search for unknown objects.

When we have prior knowledge of a TNO's path, it is possible to recover the object's signal by simply co-adding all baseline-subtracted frames along the known path using the process outlined in Steps $1-3$ of Figure 3.3. In this mode of operation, our pipeline collects and sums the small, systematic flux increase from the TNO over the full time series, resulting

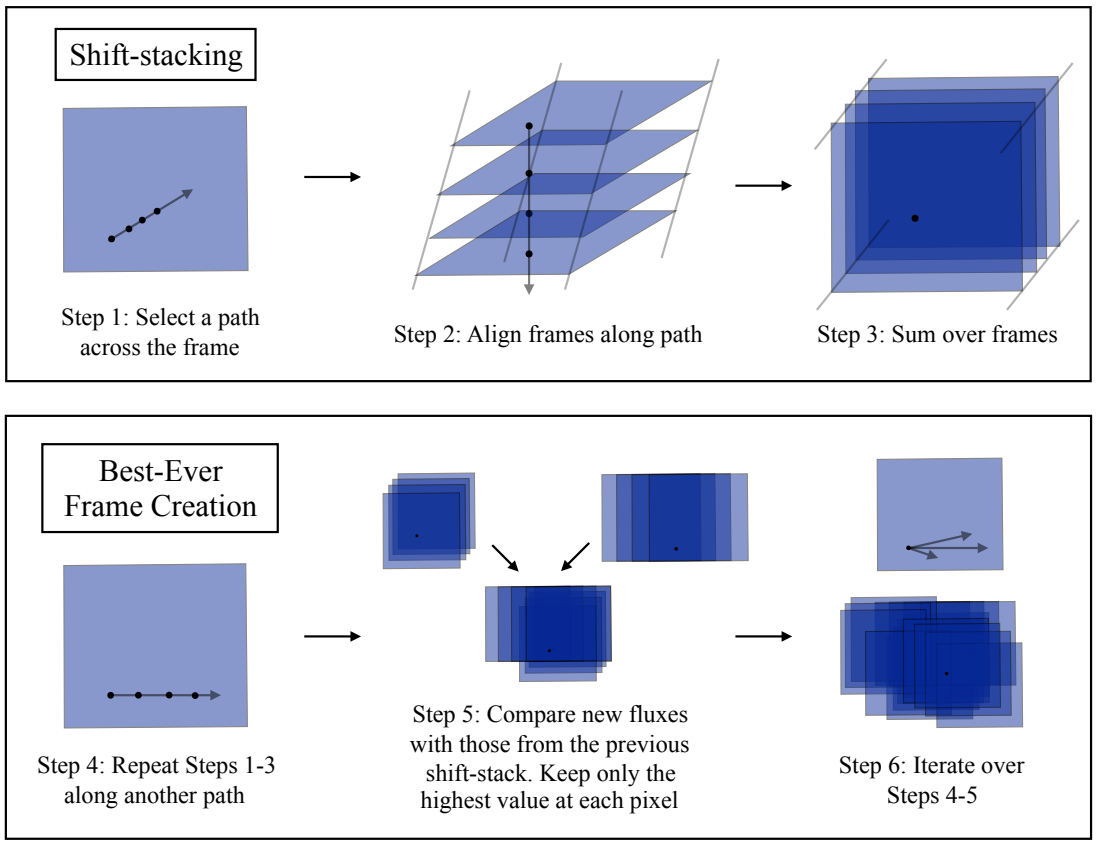

Figure 3.3: Overview of the shift-stacking algorithms included within our pipeline. The top row (Steps $1-3$) describes the general process of shift-stacking and can be applied to recover known objects along their projected orbits. The bottom row (Steps $4-6$) illustrates the additional steps implemented to create the "best-ever" frames used in our blind search.

in a point-like signal. By contrast, fluctuations in the surrounding pixel values are produced by Poisson noise; they should, as a result, sum to values centered on zero that adhere to $\sqrt{N}$ statistics. We demonstrate in Section 3.5.2 the results of this method as applied to several known objects in the TESS dataset.

The discovery of objects with *unknown* locations and orbital parameters is a more challenging task. To find new candidate objects in trans-Neptunian space, we shift and stack our baseline-subtracted images across all plausible orbital trajectories to create an aggregate "best-ever" frame using the process described in Steps $4-6$ of Figure 3.3. These best-ever frames collate the results from all shift-stack paths of interest in a region of the sky into a

single, information-dense aggregate frame.

We initiate the creation of a best-ever frame by placing the first baseline-subtracted image of our time series in the center of an empty image – a 2D array of zeroes. As illustrated in Figure 3.3, we then shift and co-add all images in the time series along a path of interest. Because our baseline-subtracted images record only deviations from the flux baseline, each constituent pixel stack, in the absence of complex systematics or objects in the frame, sums to a small number close to zero.

We then compare our post-shift-stacking image to the previous image and save only the highest flux value attained by each pixel (Step 5 in Figure 3.3). For the first shift-stack path, all positive values are saved, since our original image consists of only zeroes. We repeat this process along all possible paths; after many iterations, the zeroes of the original image are replaced with positive values. The emergent best-ever frame tracks systematic positive deviations from the baseline flux along a path. We stress that these frames are not sky images, but, rather, a graphic assessment of all the shift-stack paths under consideration within a single patch of sky.

We make some simplifying assumptions to minimize the computational expense of this search. First, we assume a constant speed across the frame for each TNO and do not incorporate deviations from a linear path between the start and end pixel. This approximation should hold for all bound TNOs of interest, since their orbital motion generates only relatively small deviations from a primary projected path determined by Earth's parallax. We also consider only integer total pixel shifts, and we shift and stack frames only at integer pixel intervals. While our pipeline includes the option to interpolate to sub-pixel shifts, we found that this alteration significantly slows down the pipeline without substantially changing the results. Sub-pixel shifts may be useful in future iterations of this work to maximize the recovered flux from a source; this could be especially useful near the edges of our current detection limits.

Without interpolation, the shift-stacking component of our pipeline consists of only quick array operations set up in a grid, with $(n_x + 1) \cdot (n_y + 1)$ operations for n_x $x-$shifts and n_y $y-$shifts. For example, for $4 < \Delta x < 47$ pixel shifts in the $x-$direction and $-8 < \Delta y < 8$ pixel shifts in the $y-$direction, used in our blind Sector 18 search, $n_x = 43$ and $n_y = 16$ for a total of 748 array operations.

The range of orbits considered in a given search determines the number of shifts necessary to include in this grid. In the parameter space explored by this work, where Earth's motion dominates objects' motion along the detector, the range of distances spanned by the population of interest at the observing epoch is the primary determining factor for appropriate values of n_x and n_y. This distance, together with the projected magnitude of an object, is also our primary observable for the distant solar system objects studied in this survey, which subtend short orbital arcs spanning consecutive nights (see e.g. Bernstein & Khushalani (2000)). We refer the reader to Section 3.5.8 of this manuscript for a description of the relationship between an object's orbital elements and its projected path on the TESS detector, which can be used to deduce the $x-$ and $y-$shifts of interest.

3.4.2.2 Distinguishing between systematic and true signals

The excess flux from an astronomical object is maximized by summing along the closest possible path to that of the object's true orbit. Similar paths that only partially align with the correct orbit produce an excess in flux, but that excess is weaker than that from the shift-stack that most closely fits the object's track. Thus, a properly recovered object should appear as a point source with a surrounding "cloud" of radially decreasing brightness, where the central pixel has the peak flux and the "cloud" corresponds to similar orbit paths that overlap with, but are slightly misaligned with, the true path.

Imperfections in the baseline subtraction method can also lead to artificial flux excesses. If the determined baseline for a given pixel does not closely follow that pixel's true profile, a

bias is introduced such that the baseline-subtracted light curve no longer follows Poisson statistics about a baseline of zero. The bias is then propagated forwards such that the associated pixel stacks no longer sum to values centered around zero, producing systematic features in our best-ever frames.

Bright asteroids outside of our search limits serve as a major source of systematics in the shift-stacked frames. Asteroids moving across the frame can produce strong spikes in individual pixel light curves, leading to a peaked flux profile that is captured by neither a polynomial fit nor a PCA fit using the surrounding pixels. Our initial pixel masking process removes only gradients that are uniformly present across the entire frame, rather than those attained by individual pixels.

Fortunately, these systematics typically manifest as extended structures in the best-ever frames, without the tapering "cloud"-like effect of the true objects (see Section 3.5.2 for specific examples). Thus, differentiating between systematics and true objects becomes a challenging task only for very dim objects or very slow-moving objects (see Section 3.5.7). We apply both baseline subtraction methods, which each produce different systematic patterns, to mitigate this problem.

3.4.3 Automated Candidate Extraction

After creating our best-ever frames (Steps 4-6 in Figure 3.3), we then identify all local maxima in these frames with summed flux $\geq 3\sigma$ above zero, where the standard deviation is computed across the entire frame. We accomplish this by first applying a maximum and minimum filter to the best-ever image, each with a neighborhood size of 5 pixels. We then find locations where the difference between the maximum and minimum filtered images is at least 3σ. To quickly identify the strongest signals, we distinguish between sources recovered at $3-5\sigma$ and at significance higher than 5σ. Finally, we calculate the weighted center of flux of each identified local maximum to more precisely pinpoint each candidate

location in the frame.

By applying this procedure to the best-ever frames, rather than to each individual shift-stacked frame, we simultaneously extract all of the strongest signals in a given frame. Our significance thresholds are set relative to zero, meaning that they are more representative of an object's significance in a single shift-stack frame (obtained from Steps 1-3 in Figure 3.3) than its significance relative to other shift-stacking combinations of neighboring pixels. However, we note that the standard deviation of the best-ever frames may substantially differ from that of an individual shift-stack frame.

In its current form, our automated source extraction algorithm does not distinguish between point-like sources and more extended sources. As a result, all sources must be further examined to verify whether they resemble compact signals rather than elongated systematics. Regardless, the automated algorithm serves as a useful tool to quickly identify possible sources of interest and to guide the eye. Future developments of this pipeline will replace this module with an alternative computer vision algorithm to efficiently distinguish between systematics and true signals in a more fully automated manner.

3.4.4 Full Pipeline Workflow - Blind Candidate Search

Combining our three pipeline components, we iteratively run the full pipeline on each frame cutout across each camera and CCD of a sector. After initial pixel masking – both in the temporal and spatial dimensions – we apply our polynomial baseline subtraction to individual cutout regions. We shift-stack the reduced images, then use the results to select promising sources with our automated candidate extraction algorithm. We vet these candidates by eye and select only those that resemble point-like signals. Then, we re-reduce the cutout regions with identified promising sources using the more computationally expensive PCA baseline subtraction. We repeat the shift-stack process and the automated candidate extraction, then cross-match to find which candidates were re-recovered.

3.5 Results

After developing our pipeline, we demonstrated its performance by recovering the signals of three known outer solar system objects. We then applied the same framework to blindly search for new candidate objects in TESS sectors 18 and 19, using injection tests to verify the performance of our pipeline. Finally, we developed a formalism to interpret the orbits of objects recovered in the TESS frames with shift-stacking.

3.5.1 Search Limits

Because outer solar system bodies have slow orbital velocities relative to that of Earth, their movement across the TESS CCDs is dominated by Earth's parallactic motion. All of the TESS cameras are aligned with the ecliptic plane in which Earth orbits, meaning that, over a 27-day observing sector, slowly-orbiting outer solar system objects primarily move in one direction – horizontally across the TESS frame (in the x-direction) – with little vertical (y-direction) motion. As a result, we set the x-pixel shift limits of our searches based on the expected parallactic movement of an object on a circular orbit at the distance of interest (see Section 3.5.8). We also allow for shifts of up to ± 8 pixels ($\pm 168''$) in the y-direction across the masked temporal baseline to account for orbital motion on high-inclination orbits.

The TESS spacecraft itself is not stationary relative to the Earth; it follows a highly eccentric geocentric orbit – characterized by perigee and apogee at $17 R_\oplus$ and $59 R_\oplus$, respectively – with a 13.7-day period in a 2:1 resonance with the Earth's moon (Ricker et al., 2015). For an object at $d = 35$ au, the difference between perigee and apogee can induce a positional shift of up to $19''$ (< 1 pixel) on the TESS detector. While this sub-pixel shift is too small to affect the results presented here, which focus on the distant solar system ($d \geq 35$ au), the TESS spacecraft orbit should be taken into account in studies focusing on more nearby solar system populations – particularly objects interior to $d = 32$ au, where the TESS spacecraft

motion can induce a shift exceeding 1 pixel.

3.5.2 Recovery of Known Objects

We begin by testing the pipeline's ability to blindly recover the known outer solar system objects listed in Table 3.1. We show that distant TNOs with magnitudes down to $V \sim 22$ are readily recoverable and distinguishable from systematic effects in our Sector 5 best-ever frames, and we provide three examples of known TNOs – 90377 Sedna (Brown et al., 2004), 2015 BP519 (Becker et al., 2018), and 2007 TG422 – run through our pipeline using its two modes of operation: with input orbital information and with our blind search method. To encompass the orbits of each of these objects, our blind searches in this section span pixel shifts corresponding to circular orbits at distances between 35 and 800 au. Our results are summarized in Figure 3.4, and we describe each individual recovery below.

Table 3.1: Blind recovery results for the three known objects shown in Figure 3.4. Values are reported at the last unmasked time in Sector 5, at t =2458461.19 JD (December 8, 2018) for all three frames. Nominal values were extracted from JPL Horizons at the same epoch, and radii were computed using $p_V = 0.32, 0.08$ and 0.04 for Sedna, 2015 BP519, and 2007 TG422, respectively. We do not include nominal shift-stack paths for these objects because, while the projected path of each object is known, nonlinearities in these paths imply that the "best" recovery may not be a straight line from the start to end location of the object during this time series. The cutout origin is reported in pixels, referenced from the first data pixel of the FFI. 2015 BP519 also crosses through Sector 4, Camera 3, CCD 1; for a direct comparison with the other two objects, we elect to include only its Sector 5 track in our analysis. Because 2007 TG422 was not recovered in the corresponding best-ever frame, we include only its nominal and known-path recovery values here for reference.

Name	(S, Cam, CCD)	Cutout Origin	Type	RA (deg)	Dec (deg)	$(\Delta x, \Delta y)$	V	d (au)	r (km)	N_{frames}	Significance
90377 Sedna	(5, 1, 4)	(1450, 1024)	nominal	56.8960	7.6094	-	20.64	86.2	500	-	-
			PCA, known	-	-	-	-	-	-	733	15.3σ
			poly, blind	57.1268	7.6637	(40, -2)	20.81	93.3	539	733	11.6σ
			PCA, blind	57.1268	7.6637	(40, -2)	20.43	93.3	644	733	8.7σ
2015 BP519	(5, 3, 2)	(900, 70)	nominal	66.8552	-33.6904	-	21.81	54.4	299	-	-
			PCA, known	-	-	-	-	-	-	733	14.3σ
			poly, blind	67.3436	-33.3435	(68, 6)	22.25	54.8	95.9	733	3.2σ
			PCA, blind	67.1341	-33.6660	(65, -3)	21.55	57.5	146	733	5.5σ
2007 TG422	(5, 1, 3)	(1367, 1260)	nominal	69.9593	8.7433	-	22.32	36.8	168	-	-
			PCA, known	-	-	-	-	-	-	736	5.8σ

3.5.2.1 90377 Sedna

One particularly interesting test case is that of Sedna, a detached Kuiper belt object described in the first row of Table 3.1. Sedna has a high eccentricity ($e = 0.86$) and large semimajor axis ($a = 484$ au), and its current proximity to perihelion in its orbit ($q = 76.1$ au) makes it one of the brightest detached KBOs, at magnitude $V = 20.64$ during the time of the TESS Sector 5 observations. Sedna's location in its orbit at the time of observation is shown in the top right panel of Figure 3.4, with the ecliptic and galactic planes included for reference. Sedna is a prime test object for our algorithm, since it is a distant object with a correspondingly short path length across the TESS cameras (42 pixels over the 22-day baseline of Sector 5 data kept after initial masking) which lies below the TESS single-frame detection limit.

We illustrate our known path recovery of Sedna in the top left frame of Figure 3.4, where we recovered Sedna's signal at 15.3σ. We define the significance of the known-path recoveries relative to the standard deviation of the full frame after masking out the 8×8 pixel region surrounding the recovered object.

This recovery was obtained by summing 733 PCA baseline-subtracted FFIs along Sedna's known path. Because we complete the shift-stacking process without accounting for Sedna's sub-pixel location, the primary flux peak is shared among a few neighboring pixels. This is expected, since a moving object spends the most time towards the edges of a pixel, resulting in substantial flux spillover into neighboring pixels. Furthermore, because of the TESS PSF, only 50% of the flux from an object falls within 1 pix^2 of the object's peak flux location.

The second and third columns of Figure 3.4 show our pipeline's blind recovery results obtained using the polynomial and PCA baseline subtractions, respectively. We emphasize that these two frames are both best-ever images obtained with no prior information related to Sedna's location or expected properties. Despite the absence of this input information, we

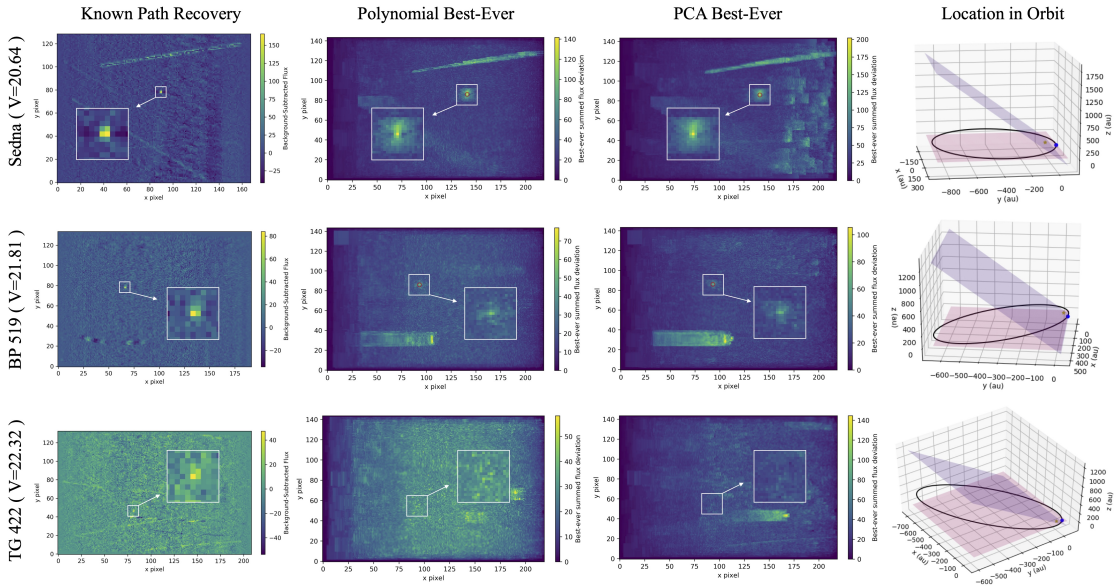

Figure 3.4: Dim solar system objects Sedna ($V = 20.64$), 2015 BP519 ($V = 21.81$), and 2007 TG422 ($V = 22.32$) recovered through our shift-stacking pipeline. The first column displays the PCA baseline-subtracted recovery of each object along its known projected path on the sky, with varying image dimensions set by the path length covered by the object. The second and third columns show best-ever frames of the same objects, which compile the results of a blind search for any objects with x-pixel shifts corresponding to $d = 35 - 800$ au and with y-pixel shifts $-8 < \Delta y < +8$ across the single-sector baseline of TESS in a 128×128 pixel frame encompassing the known objects. No prior orbital information about the objects is incorporated in these blind recoveries. The fourth column shows the location of each object in its orbit at the time of recovery, with the ecliptic plane shaded in red and the galactic plane shaded in blue for reference.

recover Sedna at high significance in both images: at 11.6σ in the polynomial recovery and at 8.7σ in the PCA recovery. In the best-ever images, the standard deviation used to set the recovery significance is determined after masking out the 20×20 pixel region surrounding the recovered object, since the recovered signals are substantially more extended than in the known-path recoveries.

3.5.2.2 2015 BP519

2015 BP519 is a high-eccentricity ($e = 0.923$) extreme trans-Neptunian object with semimajor axis $a = 454$ au and perihelion $q = 35.2$ au. For assumed albedo $p_V = 0.08$ corresponding to a typical red TNO (Stansberry et al., 2008), 2015 BP519 has radius $r = 299$ km. The object's known-path recovery, best-ever recoveries, and location along its orbit are portrayed in the second row of Figure 3.4. Each of the stacked 2015 BP519 frames in Figure 3.4 aggregates the signals of 733 TESS FFIs after initial masking.

2015 BP519 is readily recoverable through shift-stacking due to its current location near perihelion and its relatively large radius, though its fainter magnitude results in a weaker recovery than that of Sedna. The recovery of 2015 BP519 along its known path is only marginally weaker than the corresponding recovery for Sedna, with a 14.3σ peak. The best-ever frames, on the other hand, show substantially weaker detections than those of Sedna, likely owing to the higher magnitude of 2015 BP519. While 2015 BP519 still produces a clear signal recovered at high significance, its weaker recoveries suggest that it is approaching the magnitude limit of our blind search.

3.5.2.3 2007 TG422

2007 TG422 is the dimmest of the three sample TNOs that we recover, and it clearly demonstrates the power of shift-stacking to recover even very dim ($V \sim 22.3$) signals with confidence. Though it is the nearest of the three objects detailed in Table 3.1, at a distance

$d = 36.8$ au, 2007 TG422 is also significantly smaller than Sedna and 2015 BP519, with radius $r = 168$ km assuming $p_V = 0.04$ – appropriate for a neutral-colored TNO such as 2007 TG422 (Stansberry et al., 2008). With eccentricity $e = 0.931$, semimajor axis $a = 512$ au, and perihelion $q = 35.6$ au, 2007 TG422 is currently observable due to its proximity to perihelion, as shown in the bottom right panel of Figure 3.4.

Our 5.8σ known-path recovery of 2007 TG422, aggregated over 736 frames, corresponds to a contributed flux of only 0.0645 e/s, well below the single-frame detection limit of TESS. As a result, the known path recovery of 2007 TG422 produces a point source signal just marginally brighter than the background. In both of the two best-ever frames, no local maximum is found at the expected location of 2007 TG422. This indicates that the TNO's signal is not strong enough to be recovered by our blind search, placing a clear bound on our magnitude limit.

3.5.2.4 Systematics and Uncertainties

In addition to the signals of interest, each of the frames in Figure 3.4 also includes systematic features. Systematics persist due to imperfections in the baseline subtraction process, leading to residual effects that are asymmetric about the flux baseline.

Each recovery of Sedna in Figure 3.4 includes a long, diagonal streak towards the top of the frame. Figure 3.5 provides intuition for the origin of this systematic feature, which directly corresponds to a bright asteroid passing over the frame (see the frames spanning 2018-11-18 to 2018-11-21). Though the central pixels along the asteroid's path were masked due to their high maximum flux, the surrounding regions remain unmasked. This results in a high transient flux within pixels coincident with the asteroid's path, leading to a high summed flux when those pixels are included in a shift-stack.

The path of a second bright asteroid (see 2018-12-01 to 2018-12-04 in Figure 5) is also evident as a diagonal streak in the lower region of Sedna's known path recovery, while the

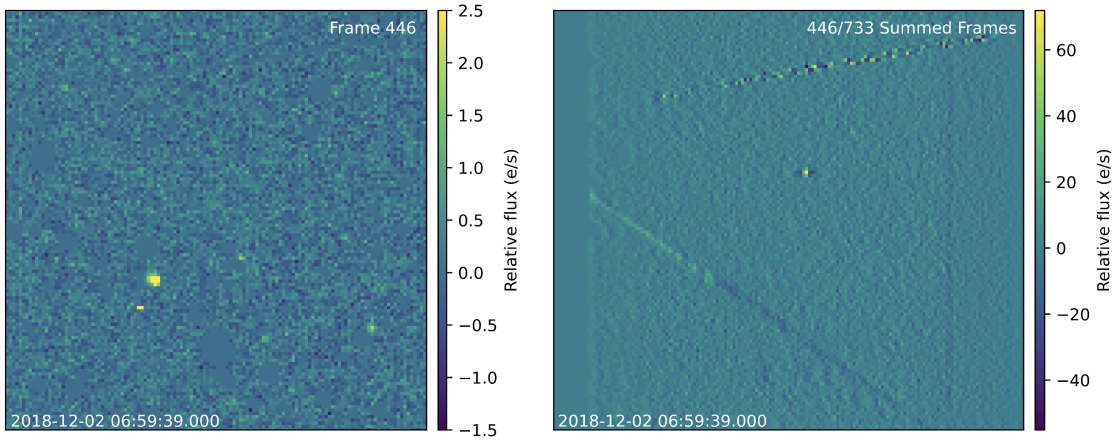

Figure 3.5: Animation of all 733 polynomial baseline-subtracted frames used to recover Sedna in Figure 3.4. The paths of two bright asteroids passing through the frame from 2018-11-18 to 2018-11-21 and from 2018-12-01 to 2018-12-04 produce diagonal systematics in the corresponding Sedna best-ever frame.

corresponding feature is more diffuse in the best-ever frames. The diffusion of the second asteroid's signal is likely due to the larger y-component of its movement, meaning that the blind shift-stack range of interest does not capture movement directly along the asteroid's orbit.

Systematics produced by foreground asteroids should be most common close to the ecliptic plane of the solar system, corresponding to TESS Camera 1 (where Sedna is located). While we made no direct effort to discard these signals in this work, the automated removal of such signals can be accomplished through cross-matching with existing asteroid catalogs to further clean the TESS dataset prior to larger-scale surveys.

In the PCA best-ever frame, numerous block-like systematics are present to the right of Sedna, while these systematics are absent in the polynomial subtraction frame. These features likely result from an incomplete removal of residual stellar fluxes in the PCA baseline subtraction, which does not incorporate the flux baseline of the pixel at hand. We find from this example and others in our blind search that the PCA best-ever frames tend to be noisier and thus more prone to false positive signals than the polynomial-subtracted

frames. However, as we describe in Section 3.5.7, the PCA frames are also more sensitive to dim, slow-moving signals and thus contribute great value to the Planet Nine search.

Horizontally extended systematic signals are present in the 2015 BP519 and 2007 TG422 best-ever frames, as well. These systematics are not directly attributable to asteroids, and they appear with higher frequency near the galactic plane. This suggests that they may be associated with stars that have not been fully masked from the frame.

The maximum fluxes attained in each of the PCA best-ever images are notably higher than those of their polynomial best-ever counterparts. This is likely because the polynomial subtraction method is more prone to self-subtraction, since the baseline profile of each pixel is determined solely from that pixel's time series flux, rather than surrounding "background" pixels in the frame. The known path recovery frame peaks at a higher flux than the polynomial best-ever frame, meaning that the path determined from our best-ever frame is similar to, but not quite the same as, Sedna's actual path. Deviations between these two maximum fluxes stem from nonlinearities in the object's real path across the TESS frame, which are ignored in the best-ever image creation.

While our PCA best-ever frames can be used to circumvent the issue of self-subtraction, uncertainties related to the extent of these path nonlinearities are inherent to our linear search method. As a result, our magnitude estimates are systematically lower than the real object magnitudes, and this effect is more apparent for more nearby objects with stronger nonlinearities. The path uncertainty and thus dilution of signal reduces the range of magnitudes recoverable with this method; however, it also suggests that recovered objects should be more amenable to follow-up observations, since their actual magnitudes may be slightly elevated over those suggested by our best-ever frames.

3.5.3 New Candidates from a Blind Search

3.5.3.1 Candidate Extraction

After verifying our ability to recover known objects, we then conducted a blind search for previously unknown outer solar system objects. We completed this search using all FFIs from Sectors 18 and 19. Due to the presence of the bright star Polaris ($I = 1.22$; Ducati, 2002) in Camera 3, CCD 3 of Sector 19, a column of pixels across the CCD was saturated, leading to an overestimated smear correction affecting multiple columns of pixels.[4] These columns produced strong systematics in our pipeline; as a result, we discarded the affected 256×256 frames, which constitute a small fraction of our total sample.

We searched all paths with x-shift bounds set to match $a = 70 - 800$ au for circular orbits: around $4 \leq \Delta x < 47$ pixels after the initial masking step has reduced the full temporal baseline to $\sim$19 days in Sector 18, and around $5 \leq \Delta x < 58$ pixels for the $\sim$ 23-day baseline in Sector 19. As in Section 3.5.2, we also allowed for y-shifts of up to ± 8 pixels across the masked baseline to account for small path deviations from solely parallax-dominated motion. After initial masking, our final image stacks included 574 frames in our Sector 18 analysis and 540 frames in our Sector 19 analysis. This is roughly 25% fewer frames than were used in the Sector 5 recoveries described in Section 3.5.2, meaning that our magnitude limit will accordingly be about 0.3 mag higher before accounting for the increased stellar density in the galactic plane.

Our automated extraction process returned a large number of $\geq 3\sigma$ local maxima (typically between a few and a few tens per 256×256 frame) from the polynomial baseline-subtracted best-ever frames. However, upon visual inspection, we found that most of these local maxima could be quickly discarded as systematics. We carefully inspected all frames,

[4] https://archive.stsci.edu/missions/tess/doc/tess_drn/tess_sector_19_drn26_v02.pdf

taking note of those containing potentially promising candidate sources, and re-ran the promising frames using the more computationally intensive PCA baseline subtraction method.

We cross-compared frames to determine whether the point-like features present in the polynomial-subtracted best-ever frames were recovered with the PCA subtraction. In many cases, the second baseline subtraction revealed that these features more closely resembled systematics when processed in a different manner. By comparing our results from both reductions, we narrowed down our set of candidates to those listed in Table 3.2.

3.5.3.2 Physical Property Estimates

For each candidate, we estimated several physical properties – sky coordinates, distance, radius, and magnitude. We extracted each property using both baseline subtraction methods, resulting in two separate estimates of each parameter in Table 3.2. The discrepancy between results from each reduction method provides a sense of the parameter uncertainty for individual candidates. Below, we detail our methods for determining each of these properties.

Sky Coordinates We directly extracted the coordinates of candidates at each time based on their pixel locations in the TESS frames. Because we did not implement our shift-stacking method with sub-pixel precision, an uncertainty floor is set by the size of the pixels ($21'' \times 21''$). Our true uncertainty is higher – on the order of a few pixels – because only most, but not necessarily all, of the true object path is required to match the simplified shift-stacked path for our algorithm to return a recovery.

Distance The distance to each object was determined from the shift-stack path length covered by the object over the TESS observing baseline, where we assumed that the

$N_{\rm cand}$	(S, Cam, CCD)	Cutout Origin	Type	RA (deg)	Dec (deg)	$(\Delta x, \Delta y)$	V	d (au)	r (km)	$N_{\rm frames}$	Epoch (JD)	Significance
1	(18, 2, 1)	(256, 384)	poly	43.9497	69.3189	(18, 1)	21.11	167.7	1517	573	2458810.25	9.29σ
			PCA	43.9497	69.3189	(18, 1)	21.15	167.7	1487	573	2458810.25	3.98σ
2	(18, 2, 3)	(0, 1664)	poly	16.6250	49.8912	(24, 4)	20.80	123.5	948	573	2458810.25	6.69σ
			PCA	16.6250	49.8912	(24, 4)	20.56	123.5	1060	573	2458810.25	4.31σ
3	(18, 3, 3)	(1280, 1664)	poly	344.4852	78.6428	(36, -3)	21.32	83.8	343	574	2458810.25	6.09σ
			PCA	344.5048	78.6538	(34, -3)	20.92	88.7	462	574	2458810.25	5.70σ
4	(18, 4, 2)	(1664, 1664)	poly	244.9781	73.4310	(32, 3)	22.05	95.6	320	574	2458810.25	6.96σ
			PCA	244.9781	73.4310	(32, 3)	21.70	95.6	375	574	2458810.25	5.24σ
		(1664, 1536)	poly	244.9781	73.4310	(32, 3)	22.04	95.5	321	574	2458810.25	7.20σ
			PCA	244.9781	73.4310	(32, 3)	21.69	95.5	376	574	2458810.25	5.21σ
5	(18, 4, 2)	(768, 1664)	poly	260.9884	70.8163	(32, 0)	22.16	96.3	309	574	2458810.25	5.21σ
			PCA	260.9884	70.8163	(32, 0)	21.43	96.3	375	574	2458810.25	3.91σ
		(768, 1792)	poly	260.9884	70.8163	(32, 0)	22.16	96.3	310	574	2458810.25	5.81σ
			PCA	260.9884	70.8163	(32, 0)	21.74	96.3	431	574	2458810.25	3.44σ
6	(18, 4, 3)	(1152, 1280)	poly	252.6156	65.2268	(38, 4)	22.06	80.7	227	574	2458810.25	6.03σ
			PCA	252.5973	65.2182	(37, 3)	21.4	83.1	319	574	2458810.25	3.91σ
7	(19, 2, 3)	(896, 1664)	poly	57.8788	61.6236	(36, 3)	20.74	106.7	727	539	2458838.92	5.73σ
			PCA	57.8788	61.6236	(36, 3)	20.20	106.7	933	539	2458838.92	3.95σ
8	(19, 3, 2)	(1664, 1024)	poly	122.2438	81.4212	(36, 0)	22.02	107.1	407	539	2458838.92	5.16σ
			PCA	122.2698	81.4253	(36, 3)	21.66	106.7	477	539	2458838.92	4.55σ
9	(19, 3, 2)	(896, 1536)	poly	105.0573	86.9216	(19, -2)	21.77	201.7	1616	539	2458838.92	7.07σ
			PCA	105.1514	86.9321	(19, 0)	21.62	202.8	1758	539	2458838.92	5.66σ
10	(19, 3, 2)	(1024, 1024)	poly	98.0576	83.7762	(42, 3)	21.89	91.5	316	539	2458838.92	5.74σ
			PCA	98.0063	83.7598	(42, 0)	21.48	91.8	384	539	2458838.92	5.02σ
11	(19, 3, 2)	(896, 1024)	poly	99.9219	83.7321	(23, 6)	21.92	162.1	976	539	2458838.92	5.76σ
			PCA	99.9219	83.7321	(23, 6)	21.59	162.1	1137	539	2458838.92	3.34σ
12	(19, 3, 2)	(1664, 1408)	poly	132.6612	82.2808	(30, 6)	21.85	126.0	608	539	2458838.92	6.06σ
			PCA	132.6612	82.2808	(30, 6)	21.65	126.0	667	539	2458838.92	4.53σ
13	(19, 3, 2)	(1408, 1664)	poly	140.7700	85.2651	(39, 5)	21.96	98.0	350	540	2458838.92	5.80σ
			PCA	140.7700	85.2651	(39, 5)	21.71	98.0	393	540	2458838.92	3.14σ
14	(19, 3, 3)	(640, 1280)	poly	206.4893	85.7287	(47, -3)	22.23	81.8	216	539	2458838.92	5.46σ
			PCA	206.3024	85.7390	(47, 0)	21.81	82.0	263	539	2458838.92	5.45σ
15	(19, 4, 1)	(1152, 1792)	poly	283.4931	66.3750	(32, 3)	21.91	119.9	538	540	2458838.92	5.19σ
			PCA	283.4931	66.3750	(32, 3)	21.40	119.9	679	540	2458838.92	3.33σ
16	(19, 4, 2)	(1408, 896)	poly	242.2779	73.2150	(42, 1)	22.04	91.7	295	539	2458838.92	6.17σ
			PCA	242.2779	73.2150	(42, 1)	21.62	91.7	359	539	2458838.92	5.22σ
17	(19, 4, 3)	(896, 1024)	poly	253.6466	60.7441	(45, -2)	22.20	85.6	240	539	2458838.92	5.33σ
			PCA	253.6466	60.7441	(45, -2)	21.85	85.6	282	539	2458838.92	3.04σ

Table 3.2: Candidates recovered in best-ever frames obtained with both baseline subtraction algorithms. We report values recovered from both subtraction methods. Coordinates are reported at the last unmasked time in the sector, and the reported distances (d) refer to the predicted distance between the candidate object and the TESS spacecraft at the epoch of detection. For objects recovered in two separate stacks, four entries are included in the table, with results from the second stack provided as the third and fourth rows. Significances are reported as the deviation above zero flux recovered in our automated candidate extraction, where the standard deviation is calculated across the full best-ever frame.

movement of an object across the sky was dominated by Earth's parallactic motion. Thus, we inferred the distance to an object using its extracted pixel shift together with Earth's known orbital velocity (see Section 3.5.8 for further details of this calculation). We approximated that contributions to the pixel shift from the object's orbital motion were negligible.

Radius To calculate estimated radii r, we used the scaling relation $F \propto d^{-4}$ for reflected light, where F is flux at the TESS CCD and d is the distance to the object from Earth. Then, the number of counts N_c collected in a pixel is related to the distance d and radius r of an outer solar system object using Equation 3.3.

$$N_c \propto \frac{r^2}{d^4} \tag{3.3}$$

Combining this scaling relation with Sedna's known properties and recovered flux, we estimated the radii of all recovered and candidate objects in our pipeline. We used the peak flux obtained for Sedna in its PCA known-path recovery for this scaling, noting that the uncertainty in our radius estimate is directly tied to the uncertainty in an object's true contributed flux. Uncertainties in this contributed flux are introduced by self-subtraction and discrepancies between the object's recovered, linear path and its true path on the sky. By using this scaling, we implicitly adopt the geometric albedo of Sedna, $p_V = 0.32$ (Pál et al., 2012).

Flux Calibration We scaled all extracted signals with the flux of Sedna in order to deduce the magnitude of each object, as well. This standard scaling relation is given by Equation 3.4.

$$\frac{F_1}{F_2} = 10^{(V_2 - V_1)/2.5} \tag{3.4}$$

We assume that the visual (V) magnitude scaling is roughly equivalent to that of the TESS passband, which most closely resembles the I band in the Johnson-Cousins UBVRI system. This assumption would hold in the case of a perfectly spherical, Lambertian surface; however, deviations from sphericity and variations in albedo across objects introduce additional uncertainties to our scaling, and, as a result, to our extracted magnitude estimates. By scaling with Sedna, we again implicitly assume that the albedos of candidate objects are similar to that of Sedna ($p_V = 0.32$). For objects with a true albedo lower than that of Sedna, this means that our predicted sizes may be underestimated.

The results of this scaling are shown in Figure 3.6, with the known magnitudes and extracted flux values of 2015 BP519 and 2007 TG422 included for reference. We use the known path recoveries to determine the per-frame flux contribution of each object, dividing the peak summed flux by the total number of frames contributing to the sum. This extrapolation leads to an overestimated magnitude for 2015 BP519 and 2007 TG422 given their recovered fluxes, suggesting that real, recovered objects may be slightly brighter than our algorithm reports.

3.5.4 Cross-Check Against Known Objects

After identifying our candidates, we checked all known distant solar system objects to verify whether any should have been observable in our survey. The utility of this exercise is twofold. First, it allows us to determine whether our candidates coincide with previously discovered objects and to verify whether our candidates are truly "new" potential objects. Second, it provides information regarding whether our algorithm failed to recover any known objects that should have been detectable.

We considered all objects in the `distant_extended.dat` file downloaded from the International Astronomical Union Minor Planet Center[5] on February 8th, 2020. This

[5] https://www.minorplanetcenter.net/data

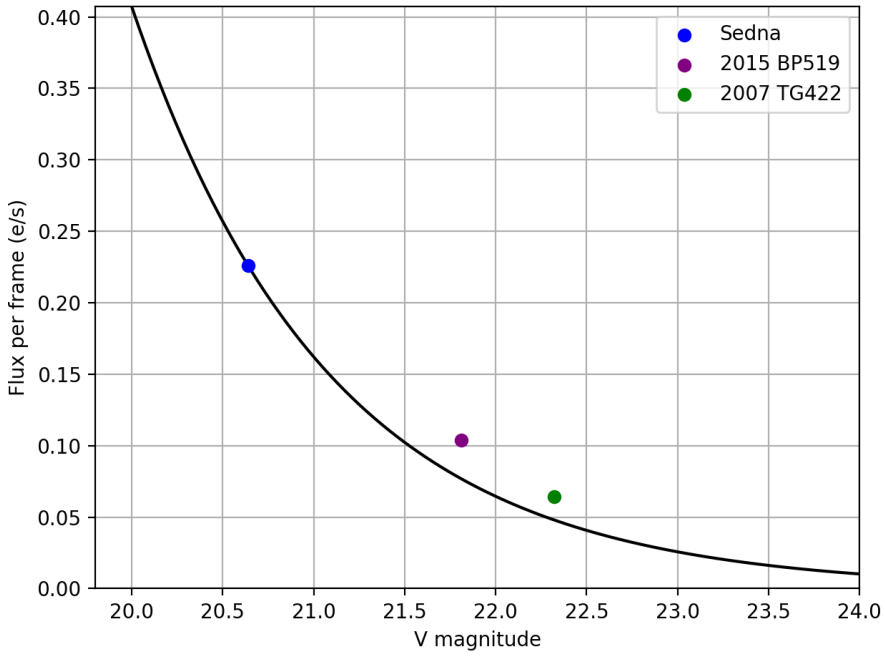

Figure 3.6: Flux calibration used to estimate the magnitude of all recovered signals through scaling with Sedna. The recovered fluxes of 2015 BP519 and 2007 TG422 are included for reference. The elevated flux per frame of these two objects over the predicted values suggests that true objects may be slightly brighter than our algorithm reports.

file includes all known TNOs, Centaurs, and scattered disk objects in the Minor Planet Center catalog. We first extracted the sky position of each object as viewed by the TESS spacecraft during the midpoints of Sectors 18 and 19 using the `jplhorizons` module of `astroquery` (Ginsburg et al., 2019). Afterwards, we applied the `TESSCut` module of `astroquery` to determine whether those coordinates fell within the TESS field of view during Sector 18 and/or 19.

Once we identified a list of objects within the field of view, we then determined whether any were consistent with the limits of our shift-stack search. Our search includes objects moving at a sky-plane projected speed consistent with that of a body on a circular orbit between $d = 70 - 800$ au. More specifically, this means that we searched for objects whose projected speed corresponds to a certain range of pixel shifts: for Sector 18, $4 \leq \Delta x \leq 47$ pixels over a $\sim$19-day temporal baseline, and, for Sector 19, $5 \leq \Delta x \leq 56$ pixels over a $\sim$23-day baseline. In each case, we also allowed for $-8 \leq \Delta y \leq 8$ pixels of movement in the y-direction over the full baseline. The number of pixels traversed by an object at known orbital velocity and distance is described in Section 3.5.8. We determined the total Δx and Δy shifts expected for each object and confirmed that no known objects lie within our search parameter space; instead, they all produce larger $x-$shifts corresponding to smaller geocentric distances than are covered by our search.

3.5.5 Time Constraints for Follow-Up

For the distant solar system objects that we focus on in this work, with short orbital arcs across a single TESS sector, we are sensitive to only the distance and magnitude of recovered objects during the time of the TESS observations. We gain insufficient information from the shift-stacks to place meaningful constraints on the orbital elements of candidates. As a result, candidates must be followed up relatively quickly after the TESS observing sector from which they were extracted.

The sky-plane position of an object on a circular orbit at $d = 80$ au, the distance of our most nearby candidate, would change by up to 30 from orbital motion over the span of one year. The most distant candidate at $d = 200$ au would shift in position by up to 8. The primary direction of motion can be approximated using the assumption that these objects are on prograde orbits; however, without more detailed orbital information, the candidates would be rapidly lost in the absence of follow-up observations. It may be possible to connect orbital arcs from the current TESS observations with data from the upcoming TESS extended mission, which would allow for a more refined orbit and an extended time available for follow-up observations.

3.5.6 Expected Yield

In total, there are three known distant TNOs – Sedna, Eris, and Gonggong – with 70 au $\leq d \leq$ 100 au and $V \lesssim 22$ au. If we assume that these objects are evenly distributed across all TESS sectors, and that any objects in this parameter space lying outside of the galactic plane would have already been detected, we expect that there is a roughly 33% chance that one of our candidates in that range is real. This approximation comes from the prior that we have searched 2 of 8 sectors crossing through the galactic plane at relatively low inclinations, while roughly 18 sectors lie outside of the galactic plane.

An exact false positive rate is difficult to estimate with this method because our candidate list comprises signals that are not only outliers in flux, but that also appear as point sources in the results from both baseline subtraction methods. While our automated candidate extraction rapidly selects local maxima, it does not distinguish between point sources and more extended flux maxima corresponding to systematics. This final step is conducted through a visual assessment in which the vast majority of flux maxima are dismissed as likely systematics.

Many, if not most, of the high signal significances reported in Table 3.2 are likely the result

of unmodeled systematic errors. Most best-ever frames have several sources detected at $\geq 5\sigma$ significance due to systematic noise. This is because our significance is defined as the deviation above zero attained for a given shift-stack, rather than the deviation above all sources across all possible shift-stacks. In this way, we avoid discarding candidates due to the presence of a strong systematic signal elsewhere in the frame. Despite the large number of flux maxima that are initially recovered, only a small number of these sources pass our visual examination test.

Even with this final vetting step, we anticipate a high false positive rate due to the expected rarity of true objects in this parameter space. For the relatively small sample of objects presented in this work, we propose that the best method to verify the false positive rate would be to follow up the candidates directly with observations.

Future work will extend this search to a more optimized, all-sky survey that incorporates neural networks to rapidly and accurately identify true signals (Rice et al. in prep). By using these neural networks to fully automate the candidate identification process, it will be possible to more rigorously constrain the expected false positive rate for larger-scale surveys.

3.5.7 Injection Recovery

We also developed an injection recovery algorithm to study the recoverability of objects with varying magnitudes and path lengths across frames with differing systematics. Each injection consists of a 13×13 pixel two-dimensional Gaussian, described by

$$g(x, y) = Ae^{-\left(\frac{(x-x_0)^2}{2\sigma_x^2} + \frac{(y-y_0)^2}{2\sigma_y^2}\right)} \tag{3.5}$$

We select σ, which determines the width of the Gaussian, to match the PSF of TESS, where 50% of light from a point source falls within 1 pix^2. We use a symmetric PSF for which

$\sigma = \sigma_x = \sigma_y$. The amplitude of the Gaussian injection, A, is scaled to match the flux peak expected for an object of a given magnitude based on the extrapolation in Figure 3.6, multiplied by a factor of 3.5 to account for the fact that the peak recovered flux is shared by roughly 3.5 neighboring pixels. We verified that this scaling provides correctly-calibrated injections by reproducing roughly the expected peak fluxes of the three TNOs described in Section 3.5.2.

In our injection tests, we added these signals into the data frames prior to all data pre-processing, then ran our full pipeline with the injected objects included. Afterwards, we applied the same automated source extraction algorithm that we had used in our images with no injections for a direct comparison showing whether these sources, if present in each frame, would have been detected.

3.5.7.1 Injection Completeness Tests

To quantify our recovery rates, we injected grids of 24 identical signals into one 256×256 pixel cutout frame from each TESS camera. Signals may be more or less recoverable in different regions of a given best-ever frame due to spatially nonuniform systematics. As a result, we injected many signals into each frame and used the fraction of recovered sources in each frame as an indicator of our recovery rates.

We injected the signal grids into the cutout region with origin (1024, 1024) towards the center of CCD 1 from each camera. Because the on-sky stellar density varies with proximity to the galactic plane, we report results for each camera separately.

We varied the magnitudes and path lengths of our injected signals, then determined whether each was recovered by our automated candidate extraction algorithm in the corresponding best-ever frames. For simplicity, all injections in this section move only in the $x-$direction, with zero $y-$shift. We conducted these injection tests across the parameter space of signals explored in this work, testing both the polynomial and PCA recovery methods for

comparison. Results from both baseline subtraction methods are provided in Figure 3.7. To better understand our survey's sensitivity to Planet Nine, we also conducted injection tests using a finer grid of path lengths spanning the Planet Nine parameter space, with results in Figure 3.8.

In a small number of cases, a local maximum had already existed at the location of the injection. In these ambiguous cases, our algorithm accepts the signal as a recovery; this occurs at most 1-2 times in a given camera and therefore cannot artificially inflate our fractional recovery above 10%. Similarly, lower magnitudes and higher $x-$shifts do not unilaterally lead to higher recovery fractions because, at certain magnitude ranges, the injected signals blend with systematics in the frame and shift the locations of recovered local maxima.

Figure 3.7 demonstrates that we reliably recover objects with $V < 21$ and $x-$shift > 30 pixels ($d \lesssim 150$ au). At shorter path lengths (larger distances) and dimmer fluxes, injected signals are recovered with decreasing consistency. The galactic plane crosses through Camera 1 of Sector 19 and Camera 2 of Sector 18, contributing to a reduced fraction of recoveries in those regions.

The more finely gridded Figure 3.8 shows that, while our shift-stacking search grid covers much of the parameter space spanned by plausible Planet Nine orbits, the survey presented here cannot rule out the majority of Planet Nine parameter space. This is, in large part, because we require that any accepted candidates are recovered by *both* baseline subtraction methods. Figure 3.8 demonstrates that the PCA baseline subtraction method consistently produces more reliable recoveries for objects with $d > 190$ au – corresponding to $\Delta x \leq 15$ pixels in Sector 18, or $\Delta x \leq 19$ pixels in Sector 19 – and that it is capable of reliably recovering distant ($d \leq 800$ au) objects down to $V \sim 20$ with a single sector of data alone.

Distant objects remain in the same pixel for a large fraction of the full time series such that

they easily blend into the baseline flux. By including only surrounding pixels well outside the central pixel's PSF in the PCA subtraction, we circumvent the polynomial baseline subtraction's tendency to shift the full baseline fit slightly upward in the presence of a very slow-moving object. The PCA method shows great promise to narrow down the parameter space for even objects out to ∼800 au ($\Delta x \sim 5$ pixels) in future optimizations of this technique.

3.5.7.2 Uncertainty Characterization

We also characterized our uncertainties in the properties of recovered objects using the gridded injection tests. At each magnitude and $x-$shift, we calculated the mean offset of all 24 gridded injections relative to the known values. Then, we determined the standard deviation across the full parameter space displayed in Figure 3.7.

We found that the magnitude uncertainties retrieved with the polynomial and PCA baseline subtraction methods were $\sigma_V = 0.7$ and $\sigma_V = 1.1$, respectively. The path length uncertainties were characterized by $\sigma_{\Delta x} = 13$ pixels and $\sigma_{\Delta y} = 2$ pixels for the polynomial subtraction, and $\sigma_{\Delta x} = 8$ pixels and $\sigma_{\Delta y} = 1$ pixels for the PCA subtraction. These large path length uncertainties, particularly in the $x-$direction of motion, preclude meaningful recoveries of orbital information from our candidates, and they propagate to correspondingly large fractional uncertainties of up to $\sim 70\%$ in our estimated candidate distances and radii. The PCA subtraction recovers notably more accurate candidate properties than the polynomial method due to its lack of self-subtraction.

3.5.8 Interpretation of Orbits

Once we have extracted candidates from our best-ever frames, we consider the orbital constraints that can be placed on these objects. In our current framework, these constraints are limited by a few factors – to keep computation times manageable, our shift-stacking

algorithm does not consider sub-pixel shifts, and we assume a straight-line path for the distant solar system objects of interest. For such distant objects, a simplified orbital arc spanning just a few weeks at best does not yield meaningful constraints on the TNOs' true orbits.

Nevertheless, we can set up the formalism to understand what types of projected tracks correspond to plausible orbits and which correspond instead to either physically impossible or highly implausible tracks. The formalism described here demonstrates how a set of orbital elements can be translated to $x-$ and $y-$ TESS pixel shifts by determining the angular sky-plane path associated with the orbit. Our translation to pixel shifts is TESS-specific; however, we note that Bernstein & Khushalani (2000) has also outlined a general framework to approximate orbital properties in the short-arc limit, and we refer the interested reader to this study for details.

We use standard coordinate transformations to translate the orbital elements into Cartesian space. We complete two transformations: first, a transformation to the heliocentric ecliptic reference frame, and then to a barycentric frame that aligns with the TESS FFIs for ease of interpretation. The first of these transformations is given by Equation 3.6, where we deviate from the standard labels for each direction in order to maintain consistency with the 2D xy directions that we use to describe locations in the TESS FFIs. Throughout this section, we use the capitalized XYZ notation to describe Cartesian distances, while we continue to refer to TESS pixel shifts using the lower-case xy notation.

$$\begin{pmatrix} Z_h \\ X_h \\ Y_h \end{pmatrix} = r \begin{pmatrix} \cos\Omega\cos(\omega+f) - \sin\Omega\sin(\omega+f)\cos i \\ \sin\Omega\cos(\omega+f) + \cos\Omega\sin(\omega+f)\cos i \\ \sin(\omega+f)\sin i \end{pmatrix} \tag{3.6}$$

In this heliocentric ecliptic reference frame, Z_h is defined to be perpendicular to the

ecliptic plane, while X_h points in the direction of the vernal equinox and Y_h is mutually perpendicular. We then switch to a custom geocentric frame with Equation 3.7 using a translational transformation to Earth's location, as well as one final rotational transformation about the Y_h axis using the angle ϕ between the vernal equinox and the Sun-Earth vector at the central time of a sector.

$$\begin{pmatrix} Z \\ X \\ Y \end{pmatrix} = \begin{pmatrix} Z_h \cos\phi - X_h \sin\phi \\ Z_h \sin\phi + X_h \cos\phi \\ Y_h \end{pmatrix} + \begin{pmatrix} d_{ES} \cos\phi \\ d_{ES} \sin\phi \\ 0 \end{pmatrix} \tag{3.7}$$

Here, d_{ES} is the length of the Sun-Earth vector. While this coordinate system is centered on the Earth rather than the TESS spacecraft itself, the difference between the two is negligible for our purposes. In this final coordinate frame, Z points in the radial direction perpendicular to the TESS FFI frames. The X and Y directions are aligned with the TESS CCDs, where X is parallel to the ecliptic plane in the projected frames (consistent with our notation for the x-pixel shift direction), while Y is perpendicular to the ecliptic plane (in the y-shift direction).

Using these steps, we can obtain the Cartesian distances traveled by each object in 3D space along a specific orbital arc. Then, we can translate those distances into a projected sky-plane track as observed from Earth. For a known orbit, this tells us exactly where an object is in the TESS frame at each time.

The y-component of a TNO's track across the detector should include no contribution from Earth's parallax, since the TESS frames are aligned with the ecliptic plane (by definition, the plane of Earth's orbit). Thus, any observed y-shift of a TNO in the TESS CCDs is due to orbital motion. Large y-shifts must correspond to objects on orbits with high inclination and with locations on their orbits that bring them relatively close to Earth – either through

a high-eccentricity orbit where the detected object is near periastron, or through a small semimajor axis.

The projected angular movement of an object in the y-direction across TESS's CCDs is provided by θ_Y, which is described by

$$\theta_Y = \frac{\Delta Y}{Z}\,. \tag{3.8}$$

Here, Z is the mean line-of-sight distance between the start and end of the orbital arc. This radial distance should be much smaller than the radial movement ΔZ of the object across the TESS single-sector time baseline ($Z << \Delta Z$; $\Delta t = t_2 - t_1 \sim 27$ days for one sector). The movement of the object over Δt in the $Y-$direction, defined as perpendicular to the ecliptic plane, is given by ΔY. We assume that $\Delta Y << Z$ such that $\tan\theta_Y \sim \sin\theta_Y \sim \theta_Y$. The resulting θ_Y from Equation 3.8 provides the object's vertical projected movement in units of radians.

Given that TESS's pixels span $21'' \times 21''$, we directly convert these angular sky movements to projected pixel shifts on the TESS detector using Equation 3.9.

$$N = \theta \times \frac{206265''}{1\ \mathrm{rad}} \times \frac{1\ \mathrm{pix}}{21''} \tag{3.9}$$

The projected horizontal angular movement of an object across the detector, θ_X, can also be extracted from the object's known orbit. In this direction, parallax contributes directly to the path length spanned by an object. Earth moves in its orbit at $v_\oplus = 29.78$ km/s, and its velocity in the ecliptic plane relative to that of a candidate object v_c contributes to the candidate's θ_X path length. For simplicity, we approximate that the velocity vectors of Earth and the observed object are roughly parallel at the time of observation, meaning that the net horizontal angular shift θ_X from both parallax and orbital motion is

$$\theta_X = \frac{\Delta X + (v_\oplus - v_c)\Delta t}{Z}. \tag{3.10}$$

Then, we can again apply Equation 3.9 to convert this angular shift to a pixel shift in the TESS frames, providing the x- pixel shift for a given orbital arc.

This translation between orbital elements and pixel shifts will be useful to forward model orbit bundles consistent with a given orbital track. While we do not have strong enough constraints on the outer solar system objects in this study to warrant such an analysis, this type of forward modeling will be applicable to extensions of this survey tracking shorter-period orbits, where a larger fraction of the orbit is spanned by the TESS observations. In the shorter-period regime, shift-stacking may serve as a useful tool to refine the orbits of known objects.

3.6 Discussion

3.6.1 Expected Yield of an All-Sky Survey

There is a marked difference between the systematics present in, for example, Camera 1 of Sector 19 – located directly along the galactic plane – and Camera 4 of Sector 19, at the northern ecliptic pole with $\gtrsim 30\,^\circ$ separation from the galactic plane. As expected, we find that our algorithm returns a much cleaner baseline subtraction for frames far from the galactic plane than those pointed towards the plane and neighboring regions. This is due to stellar crowding in the vicinity of the galactic plane, which has stymied previous optical searches in this region. Roughly 30% of all TESS frames lie either directly along or adjacent to the galactic plane.

The three objects presented in Holman et al. (2019) – Sedna, 2015 BP519, and 2015 BM518 – each also have $\gtrsim 30\,^\circ$ separation from the galactic plane. This suggests that

the yield estimates in Payne et al. (2019) may be optimistic, since they are based upon these three objects that are located in particularly "clean" regions of the sky. However, our algorithm recovers both Sedna and 2015 BP519 at significantly higher SNR than the recoveries presented in Holman et al. (2019), suggesting that the magnitude limits in these "clean" regions may be even lower than previously estimated. An all-sky shift-stacking search using the TESS FFIs should have varying sensitivity as a function of proximity to the galactic plane.

We stress that even within Sectors 18 and 19, our search is not exhaustive. As demonstrated by our detection recovery grids in Section 3.5.7, the two baseline subtraction methods applied throughout this work are not equally sensitive to distant TNOs. We ultimately find that the PCA baseline subtraction method, because of its robustness against self-subtraction, is more capable of narrowing down the Planet Nine parameter space and discovering new distant TNOs than the polynomial baseline subtraction method. While we required in this work that our candidates were recovered by both methods, this may not be necessary in future work. Extensions of this project may instead consider using only one baseline subtraction method, incorporating data from the ongoing TESS extended mission that is in progress at the time of publication for an additional check.

Hardware limitations also place some minor limits on the completeness of our search. A gap is left between each TESS camera, and the four CCDs in each camera are separated by 2 mm each. As described in Section 3.5.3, some frames in CCD 3 of Camera 3 also show substantial systematic effects due to saturation from the bright star Polaris, dramatically reducing our sensitivity in the small subset of our frames containing that column. However, while these regions of negligible sensitivity exist within our search, the likelihood is slim that a rare object within our magnitude limit falls within these regions.

3.6.2 Future Directions

The analysis presented here focuses on only two TESS sectors – a small fraction of the full TESS dataset's vast sky coverage. This work can be naturally extended to a larger-scale survey by incorporating convolutional neural networks (CNNs) into the existing pipeline to streamline the process of extracting promising candidates from the pre-processed best-ever images. These CNNs can be trained on injected Gaussian signals corresponding to dim objects at a known magnitude and distance, which will simultaneously provide a rigorous test of our algorithm's performance in various regions of the sky. Automating the candidate vetting process allows for a similar search on a much larger scale, incorporating the full TESS FFI dataset to survey almost the entire sky. Future work will explore this in greater detail (Rice et al., in prep).

Beyond its originally planned 2-year survey, TESS has been approved for an extended mission that will again survey nearly the full sky, but at a 10-minute cadence rather than a 30-minute cadence. When combined with the original survey, the longer temporal baseline provided by the extended mission can be leveraged to more efficiently search for slow-moving objects in the very distant solar system. While enabling higher signal-to-noise detections and accordingly lower magnitude limits for all TNOs, this additional data will be especially useful in searches for objects as distant as the predicted Planet Nine.

Objects located hundreds of au from the Earth are particularly prone to self-subtraction, since they remain in an individual pixel for a large fraction of the full time series and therefore easily blend into the baseline flux. The longer temporal baseline afforded by the TESS extended mission will help to push the detection limits of very distant TNOs by increasing their total path lengths in the TESS dataset. Furthermore, the two independent sets of observations may be analyzed separately, and the combined results can be used to reject false positive signals. This would make it possible to use only the PCA baseline

subtraction method, which is sensitive to a wider range of objects than the polynomial method, throughout the analysis rather than requiring two separate methods.

3.6.3 Additional Applications

3.6.3.1 Small Bodies Interior to 70 au

While this work focuses specifically on the distant solar system, the same algorithm can also be applied to study solar system objects interior to 70 au. Payne et al. (2019) identifies Centaurs, as well as high-inclination Kuiper belt objects (including those interior to 70 au) as some of the most promising objects to search for with TESS shift-stacking. Other high-inclination objects interior to the Kuiper belt would be possible to study, as well; however, for more nearby objects, orbital motion comprises a larger component of the targets' total sky-plane motion. This means that the approximation that Earth's parallax dominates the objects' motion no longer holds. As a result, surveys looking for objects interior to our current search limit will require additional planning to search along specified orbits, rather than along straight paths across the detector.

Examining the full range of possible nonlinear orbits is more computationally expensive than a linear path search. Nevertheless, the problem is still tractable. Burkhart et al. (2016) found that, using the methods outlined in Parker & Kavelaars (2010), only $\sim$35 total paths were required to conduct an exhaustive search for satellites of Haumea lying 10,000-350,000 km from the primary. This is because only sufficiently distinct sky tracks need to be searched in order to recover sources along all possible orbits within a set error tolerance. Beyond this threshold, similar tracks with substantial overlap provide diminishing returns.

3.6.3.2 Interstellar Objects

Given that the TESS field of view probes the high-inclination solar system, shift-stacking with TESS FFIs may be a powerful tool to search for interstellar objects (ISOs) and place

limits on their occurrence rate. The two ISOs confirmed so far – 1I/'Oumuamua and 2I/Borisov – peaked in flux at $V \sim 20$ and $V \sim 15$, respectively, both comfortably within our magnitude limits for a single-sector shift-stack search (Meech et al., 2017; Guzik et al., 2020). Using the TESS continuous viewing zones at the ecliptic poles, these magnitude limits could be pushed even lower. With 13 sectors of data included in a single analysis, our magnitude limit would be extended to $V \sim 24 - 25$ for an object remaining in the field for the full duration of these observations. The discovery of both 1I/'Oumuamua and 2I/Borisov in relatively rapid succession suggests an abundance of ISOs passing through the solar neighborhood (Rice & Laughlin, 2019b), and a deep study with TESS could place stringent constraints on the occurrence rate and size distribution of these rare objects.

Each of our pipeline components has been designed with flexibility and interchangeability in mind, meaning that our algorithm can be adapted for application to different datasets with relative ease. The *Kepler* dataset (Borucki et al., 2010) may be a particularly powerful probe of ISO occurrence rates given its extended temporal baseline and its pointing towards the galactic apex – the direction from which interstellar objects are most likely to approach the solar system. Although the initial treatment of systematics would differ for the *Kepler* dataset, the remainder of our pipeline could be easily applied to this new set of images given an adjusted, *Kepler*-specific baseline subtraction module.

Beyond a statistical study of ISOs, a shift-stacking survey with TESS FFIs may also be capable of finding individual ISOs for further follow-up. While TESS data is not publicly released until months after observing, it may be possible to discover individual objects that have been observed by TESS on their way towards perihelion. Because individual ISOs are typically observable only for a short span of time, this survey would require a rapid turnaround time for follow-up observations once candidates have been identified. This may be possible with a fully automated pipeline that incorporates CNNs.

3.6.3.3 Directly Imaged Planets

Males et al. (2013) have demonstrated that orbital motion will be significant over the integration times needed to directly image habitable-zone planets in extrasolar systems, suggesting that "de-orbiting" prospective planets over many trial orbits, while leading to an increase in false alarms, will also be necessary for survey completeness. By shift-stacking along theoretical orbits, it is possible to recover an object's signal at its nominal value without leakage from orbital motion. Thus, particularly in the era of large telescopes, shift-stacking may also provide an important tool to fully utilize the information content of direct imaging surveys.

3.7 Conclusions

We have developed a novel pipeline that is custom-designed to search for outer solar system objects by shift-stacking FFIs from the TESS dataset. In this paper, we highlighted the performance of this pipeline by recovering three known TNOs down to $V \sim 22$. We then applied the pipeline to two sectors of TESS data – Sectors 18 and 19 – located along the northern galactic plane in a targeted search for Planet Nine and other extreme trans-Neptunian objects. From this search, we extracted a list of promising candidates that can be easily checked with optical follow-up observations using meter-class telescopes.

This work serves as a proof-of-concept that develops the foundation for larger-scale applications of a similar technique. The existing framework applied in this work can reliably blindly recover the signals of distant solar system bodies in the galactic plane with $V < 21$ and current distances $d \lesssim 150$ au. These limits are currently set by the methodology applied – that is, the requirement that candidates are recovered using two separate methods, one of which is less sensitive than the other – rather than the dataset itself, indicating that the physical boundaries of our detection limits have not yet been met. The sensitivity of our

survey also improves with distance from the galactic plane. Future optimizations of this framework, including the incorporation of neural networks and of additional data from the TESS extended mission, will further push the boundaries of these detection limits and the range of solar system bodies that can be discovered and studied with TESS.

3.8 Acknowledgements

We thank Matt Payne and Darin Ragozzine for discussions that have helped to refine the ideas explored in this work. We also thank the anonymous referees for thoughtful comments and suggestions that substantially improved this manuscript. M.R. is supported by the National Science Foundation Graduate Research Fellowship Program under Grant Number DGE-1752134. This material is based upon work supported by the National Aeronautics and Space Administration through the NASA Astrobiology Institute under Cooperative Agreement Notice NNH13ZDA017C issued through the Science Mission Directorate. We acknowledge support from the NASA Astrobiology Institute through a cooperative agreement between NASA Ames Research Center and Yale University. This research has made use of data and/or services provided by the International Astronomical Union's Minor Planet Center. We thank the Yale Center for Research Computing for use of the research computing infrastructure. This project was developed in part at the Expanding the Science of TESS meeting, which took place in 2020 February at the University of Sydney.

`numpy` (Oliphant, 2006; Walt et al., 2011; Harris et al., 2020), `matplotlib` (Hunter, 2007), `lightkurve` (Lightkurve Collaboration et al., 2018), `astroquery` (Ginsburg et al., 2019), `PyEphem` (Rhodes, 2011), `astropy` (Astropy Collaboration et al., 2013; Price-Whelan et al., 2018), `scipy` (Virtanen et al., 2020)

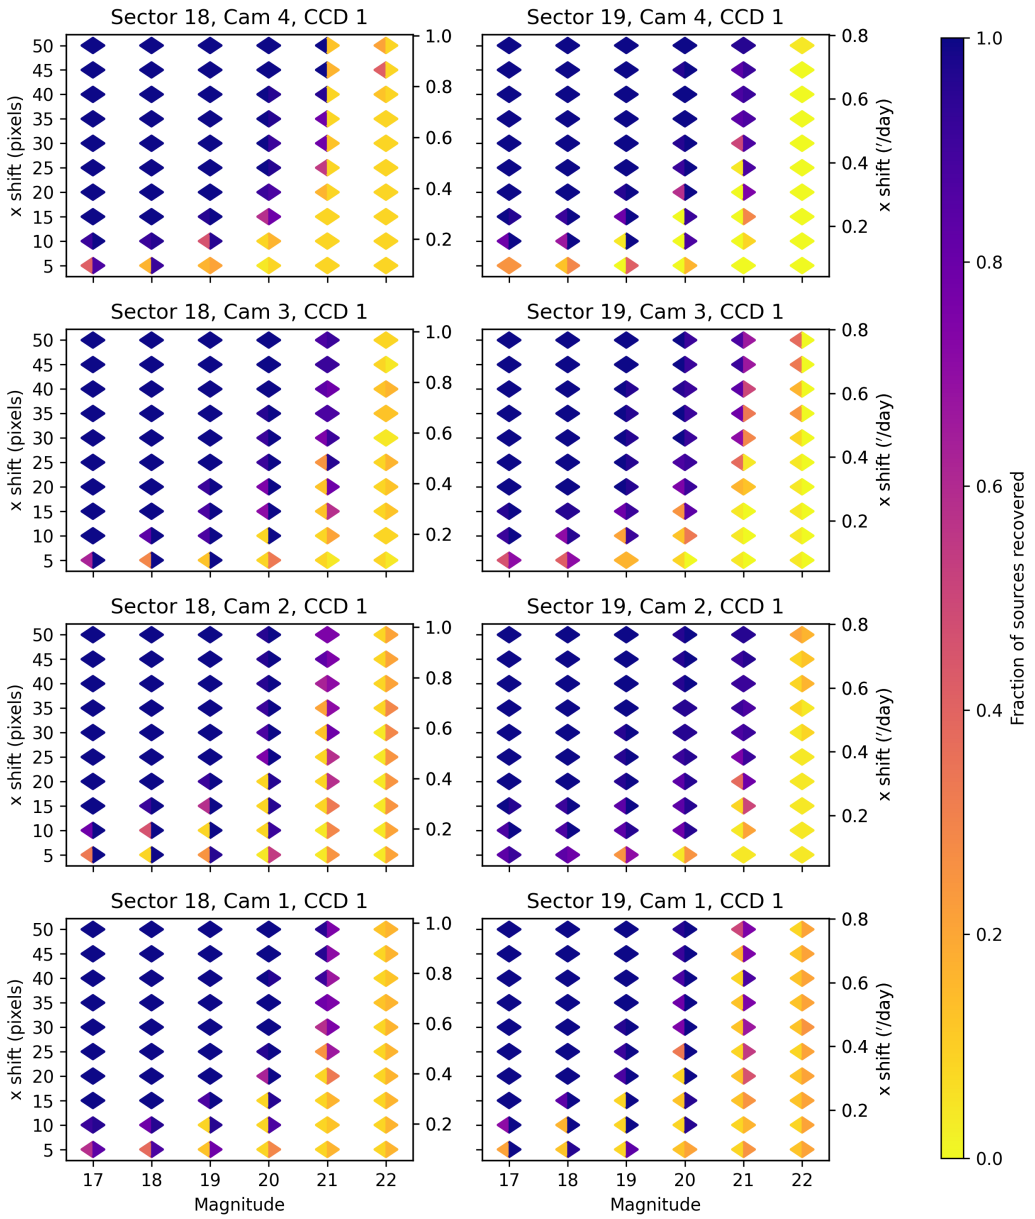

Figure 3.7: Injection recovery completeness as a function of path length and magnitude for the parameter space explored by our blind search. At each grid point, the left half of the diamond provides the recovery fraction using the polynomial baseline subtraction, while the right half provides the recovery fraction using the PCA baseline subtraction. Each pixel shift directly corresponds to an angular movement provided as an alternative y-axis on the right. For reference, 0.1/day corresponds to an object at $d = 550$ au, and 1/day corresponds to $d = 50$ au. This angular movement is primarily due to the Earth's parallactic motion, rather than the objects' movement within their own orbits.

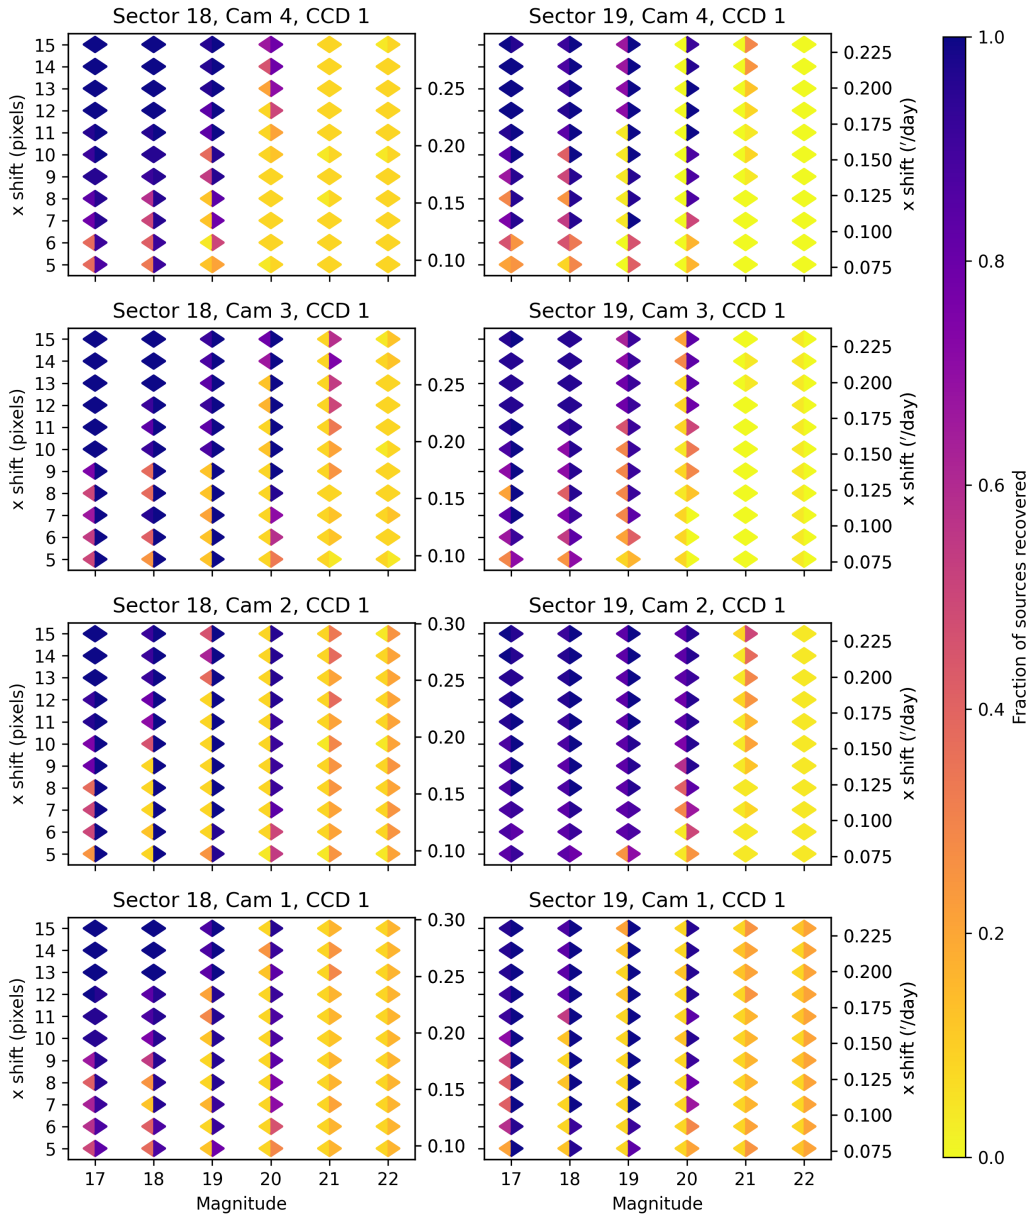

Figure 3.8: Injection recovery completeness as a function of path length and magnitude for Planet Nine-like orbits. At each grid point, the left half of the diamond provides the recovery fraction using the polynomial baseline subtraction, while the right half provides the recovery fraction using the PCA baseline subtraction. The PCA method is substantially more sensitive to distant, slow-moving objects.

Chapter 4

Hidden Planets: Implications from 'Oumuamua and DSHARP

Originally published as:

Rice, M. & Laughlin, G. 2019, The Astrophysical Journal Letters, 884, L22

4.1 Abstract

The discovery of 'Oumuamua (1I/2017 U1), the first interstellar interloper, suggests an abundance of free-floating small bodies whose ejection into galactic space cannot be explained by the current population of confirmed exoplanets. Shortly after 'Oumuamua's discovery, observational results from the DSHARP survey illustrated the near-ubiquity of ring/gap substructures within protoplanetary disks, strongly suggesting the existence of a vast population of as-yet undetected wide-separation planets that are capable of efficiently ejecting debris from their environments. These planets have $a \gtrsim 5$ au and masses of order Neptune's or larger, and they may accompany $\sim$50% of newly formed stars (Zhang et al., 2018). We combine the DSHARP results with statistical constraints from current time-domain surveys to quantify the population of detectable icy planetesimals ejected by disk-embedded giant planets through gravity assists. Assessment of the expected statistical distribution of interstellar objects is critical to accurately plan for and interpret future detections. We show that the number density of interstellar objects implied by 'Oumuamua is consistent with 'Oumuamua itself having originated as an icy planetesimal ejected from a DSHARP-type system via gravity assists, with the caveat that 'Oumuamua's lack of observed outgassing remains in strong tension with a cometary origin. Under this interpretation, 'Oumuamua's detection points towards a large number of long-period giant planets in extrasolar systems, supporting the hypothesis that the observed gaps in protoplanetary disks are carved by planets. In the case that 'Oumuamua is an ejected cometary planetesimal, we conclude that LSST should detect up to a few interstellar objects per year of 'Oumuamua's size or larger and over 100 yr^{-1} for objects with $r > 1\,\mathrm{m}$.

4.2 Introduction

In October 2017, 1I/2017 U1, now 'Oumuamua, was identified by the Panoramic Survey Telescope and Rapid Response System (Chambers et al., 2016, Pan-STARRS) as the first interstellar interloper observed traversing the solar system (Meech et al., 2017). 'Oumuamua's measured eccentricity $e = 1.1956 \pm 0.0006$ places its trajectory firmly in the regime of hyperbolic orbits. It was on its way out of the solar system with heliocentric distance $d_{\rm H} = 1.22$ au upon its discovery, having passed perihelion with $d_{\rm H} = 0.25$ au at closest approach (Meech et al., 2017). There is no clear consensus about the exact nature of 'Oumuamua, due in part to the short window of observability after its discovery and the corresponding sparsity of data, as well as the seemingly conflicting lines of evidence that have since emerged (Sekanina, 2019).

Following the detection of 'Oumuamua, several authors estimated occurrence rates for analogous interstellar objects (hereafter, ISOs) free-floating in galactic space, with all estimates falling within the range $n \approx 0.004 - 0.24\,{\rm au}^{-3}$ deduced by Portegies Zwart et al. (2018)). Assuming a cylindrical galaxy with 10^{11} stars, $R = 3 \times 10^4\,{\rm pc}$ and $H = 10^3\,{\rm pc}$ as in Laughlin & Batygin (2017), even the most conservative of these estimates, $n \approx 0.004\,{\rm au}^{-3}$, results in at minimum $0.36 M_\oplus$ of free-floating material per star, while the highest estimate $n \approx 0.24\,{\rm au}^{-3}$ implies over $20 M_\oplus$ per star. Throughout this work, we adopt the fiducial estimate $n = 0.2\,{\rm au}^{-3}$ from Do et al. (2018), which carefully incorporates the volume probed by Pan-STARRS to provide the most robust estimate to date.

'Oumuamua's detection appears to require a high density of material ejected into interstellar space, though both the ejection mechanism and the origins of the body remain under debate (e.g. Bannister et al., 2019). No coma, nor any cometary molecular emission bands, was detected in association with 'Oumuamua (Meech et al., 2017; Trilling et al., 2018). Micheli et al. (2018), however, found that 'Oumuamua accelerated out of the solar system at a rate

that could not be explained by gravitational forces alone, suggesting a cometary nature with acceleration induced by out-gassing.

Icy material is ejected from stellar systems during multiple evolutionary stages. Early on, close interactions between stars forming in open clusters can liberate icy planetesimals from their circumstellar disks (Hands et al., 2019). Later, after giant planets have formed in the system, debris is ejected through dynamical interactions with these planets (e.g. Barclay et al., 2017; Raymond et al., 2018). Much of a star's enveloping Oort cloud is ultimately shed during the post-main-sequence stage of a star's lifetime (e.g. Veras et al., 2011; Do et al., 2018; Torres et al., 2019), releasing further volatile-rich material.

In this paper, we quantify the rate of icy planetesimal ejection induced by gravity assists from circumstellar planets, with the aim of setting expectations for future observations from the Large Synoptic Survey Telescope (LSST; Abell et al., 2009). This exercise is critical to accurately assess the arrival prospects for future interstellar objects, each of which provides a uniquely valuable window into the nature of other planetary systems.

To accomplish this, we study planetesimal ejection rates from the long-period giant planet population suggested by the Disk Substructures at High Angular Resolution Project (DSHARP; Andrews et al., 2018), which surveyed 20 nearby protoplanetary disks at high resolution using the Atacama Large Millimeter Array (ALMA). Considering the case in which 'Oumuamua stems from the population of ejected volatile-rich planetesimals, we show that gravity assists by long-period giant planets are capable of reproducing the number density of interstellar asteroids implied by the detection of 'Oumuamua. We use our results to predict the detection rate by LSST for ISOs produced through this channel.

Moro-Martín (2018) previously used multi-component power law models to explore a wide range of possible ISO size distributions consistent with 'Oumuamua's appearance. Our estimates build on this previous work in several ways:

- Moro-Martín (2018) adopted relatively low giant planet occurrence rates ($f_{\rm pl}$ = 0.2 for A-K2 stars and $f_{\rm pl}$ = 0.03 for K2-M stars) based on radial velocity measurements probing planets with semi-major axes $a \lesssim 3$ au. The DSHARP survey, however, suggests an abundance of giant planets ($f_{\rm pl} \sim 0.5$) at wider separation $a \gtrsim 5$ au. We focus on these long-period planets, which are substantially more effective planetesimal ejectors than their shorter-period counterparts.

- Changes in the boundaries of a size distribution ($r_{\rm min}$ and $r_{\rm max}$) can dramatically alter the resulting fit. In Moro-Martín (2018), these boundaries were set as constant values $r_{\rm min} = 1000$ km and $r_{\rm min} = 1\,\mu$m, motivated by solar system models. We include no assumptions for the value of $r_{\rm min}$ and leave $r_{\rm max}$ as a free parameter within our models. We thus avoid assumptions about the degree of similarity between the size distributions of solar system KBOs and free-floating ISOs, which may originate from a range of collisional histories that do not necessarily resemble that of the Kuiper belt.

- Lastly, Moro-Martín (2018) assumes that 100% of solid circumstellar material around single stars and wide binaries is ejected. We instead perform N-body simulations of several representative systems from the DSHARP survey to deduce the expected mass of ejected material.

We present our simulation setup in Section 4.3 and subsequent results in Section 4.4, including our final range of possible power law solutions. We then discuss implications of our findings for future detections in Section 4.5 and sources of uncertainty in Section 4.6 prior to concluding.

4.3 Methods

4.3.1 Planets as Interstellar Comet Ejectors

Ejection of planetesimals is a natural outcome of close encounters with Jupiter, as observed in N-body simulations of early solar system evolution (e.g. O'Brien et al., 2006; Levison et al., 2009). Yet, not all planets are capable of efficiently ejecting material from their circumstellar systems. To readily expel material through gravity assists, a planet must have Safronov number $\Theta \gtrsim 1$, with Θ given by

$$\Theta = \frac{v_{esc}^2}{v_{orb}^2} = \frac{M_p a}{R_p M_*} \,. \tag{4.1}$$

Here, M_p, R_p, and a are the mass, radius, and semimajor axis of a planet, respectively, while M_* is the mass of the host star. Notably, hot Jupiters and super Earths cannot efficiently eject material from their systems. Long-period ($a \gtrsim 5\,\mathrm{au}$) planets of Neptune's mass or greater are the most effective ejectors; these planets, however, lie in a region of parameter space that is heavily disfavored by the detection biases of the transit and radial velocity (RV) methods. Direct imaging results show that occurrence rates for giant planets with orbital separation $10 - 100$ au are low, at $9^{+5}_{-4}\%$ between $5 - 13 M_{\mathrm{Jup}}$ (Nielsen et al., 2019), indicating that the primary planetary ejectors may be less massive than the current direct imaging detection limits and/or lie within $\sim 10\,\mathrm{au}$ of the host star.

Recently, a candidate population of such planets was inferred by Zhang et al. (2018) using results from DSHARP. Disk substructures were near-universally found in the DSHARP sample, where axisymmetric gaps and rings are most common. The DSHARP sample is biased towards relatively large disks around massive stars; however, gaps and rings at comparable radii have also been identified in samples more representative of the average protoplanetary disk population (e.g. Long et al., 2018), suggesting that they may be prevalent

among the underlying population. Although these substructures could be induced by a number of mechanisms, mounting evidence for planetary companions in similar systems (e.g. PDS 70; Haffert et al., 2019) favors the hypothesis that the substructures in the DSHARP sample are caused by planets. Recent kinematic detections also point towards the presence of planets at radial locations coincident with observed gaps in protoplanetary disks (Pinte et al., 2019), further strengthening the case that some of the observed substructures are indeed induced by planets.

Consequently, the DSHARP sample suggests a $\sim 50\%$ occurrence rate for giant planets with masses between $\sim M_{\rm Nep}$ and a few $M_{\rm Jup}$ orbiting their host stars at separation $a \gtrsim 5$ au (Zhang et al., 2018). This abundance of long-period giant planets is roughly in agreement with previous results from Bryan et al. (2016), which combined Keck RV measurements with NIRC2 adaptive optics imaging to obtain an occurrence rate of $52 \pm 5\%$ for planets with $M = 1 - 20 M_{\rm Jup}$ and $a = 5 - 20$ au. An abundance of Neptune-mass planets is also supported by microlensing results, which find a peak in planet occurrence rates at planet-to-star mass ratio $q \sim 10^{-4}$, corresponding to roughly $20 M_{\oplus}$, or 1.2 $M_{\rm Nep}$, for typical host star mass $\sim 0.6 M_{\odot}$ (Suzuki et al., 2016).

4.3.2 Simulation Setup

To determine how this planet population connects to the population of interstellar asteroids, we complete a detailed assessment of mass ejection rates in three DSHARP disk systems with clear radial gaps – HD 143006, AS 209, and HD 163296 – using the `REBOUND` orbital integrator (Rein & Liu, 2012). We adopt host star masses and radial planet positions from Table 3 of Zhang et al. (2018). For each predicted planet, Zhang et al. (2018) reports several masses derived with varying assumptions for the disk dust size distribution and α viscosity parameter. Because these model parameters are not well-constrained, we randomly assign the mass of each planet from the range $[M_{\rm min}, M_{\rm max}]$, where $M_{\rm min}$ and $M_{\rm max}$ are

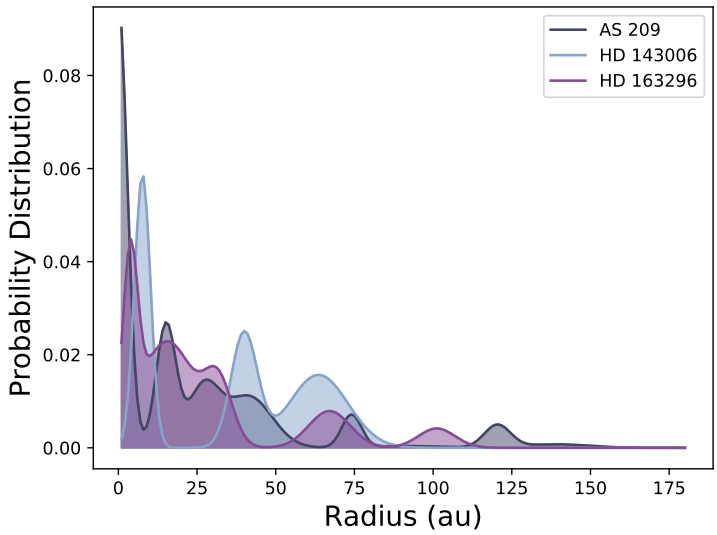

Figure 4.1: Initial distribution of dust particle semimajor axes for each disk in our sample. AS 209 is shown in gray, HD 143006 in blue, and HD 163296 in purple.

the minimum and maximum derived masses, respectively. In accordance with results from Zhang et al. (2018), the HD 143006 and AS 209 systems each include 2 planets, while the HD 163296 system includes 3 planets.

We focus on the millimeter-sized dust population in our simulations, since dust masses in this size regime are relatively well-constrained by ALMA. We estimate the total mm dust mass $M_{\rm dust}$ of each disk using the relation

$$\log(M_{\rm dust}/M_{\oplus}) = 1.3 \times \log(M_*/M_{\odot}) + 1.1 \tag{4.2}$$

derived in Pascucci et al. (2016), where M_* is the mass of the host star. This derivation is based on the 887 μm flux measured from observations of the Chamaeleon I star-forming region, with bandwidth spanning $867 - 917\,\mu$m. Thus, the dust masses in our simulations represent only the subset of grains with radii in this size range.

To reproduce the dust mass distribution in each system, we initialize 3500 equal-mass test

particles radially distributed using the semimajor axis probability distributions depicted in Figure 4.1. These distributions are based on radially symmetric best-fitting models consisting of several superposed Gaussians for each disk, with parameters for AS 209, HD 163296, and HD 143006 given in Guzmán et al. (2018), Isella et al. (2018), and Pérez et al. (2018), respectively. Each planetesimal is initialized with orbital elements

$$\Omega = \omega = 0$$

$$f \in [-\pi, +\pi]$$

$$i \in [0.0005, 0.005]$$

$$e \in [0.04, 0.06]$$

where all angles are in radians. We initialize each of our three disk-star-planet systems in three iterations each with a different random seed, resulting in a total of nine simulations with varying planet masses, orbital elements, and initial dust particle placements. We integrate each of these systems for one week on the *Grace* supercomputing cluster at Yale University.

4.4 Results

To encapsulate the typical behavior of all systems, we find the average mm mass $m_{\rm ej}$ ejected by each star as a function of time t for all nine simulations, with our results displayed in Figure 4.2. The resulting curve, representing the total mass of all particles with a formally positive energy relative to the host star, is well-fit by exponential function

$$m_{\rm ej} = a \log_{10} t + b, \tag{4.3}$$

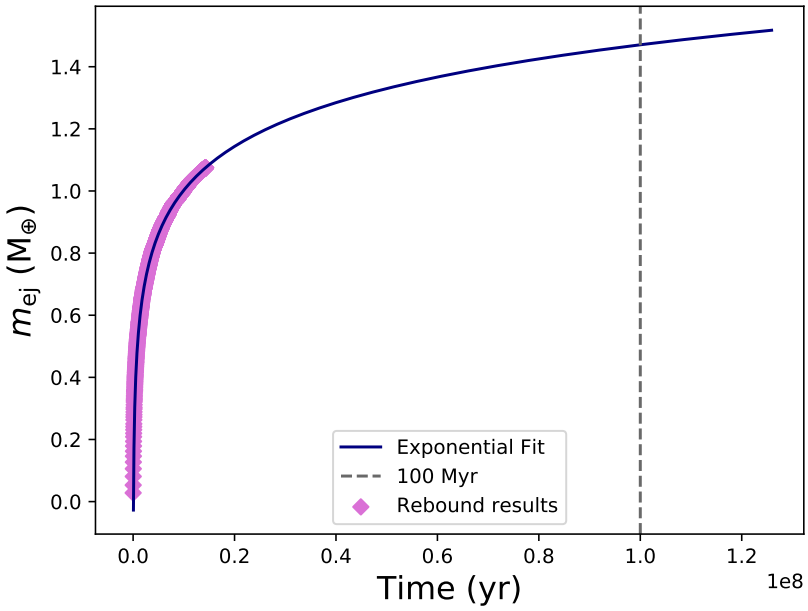

Figure 4.2: Average mass ejected as a function of time for all disks in our sample.

with $a = 0.468$ and $b = -2.274$. Using this profile, we extend our results to $t = 10^8\,\mathrm{yr}$, at which point $m_{\mathrm{ej}} = 1.47 M_\oplus$. We take this as the representative mm-sized dust mass ejected by each giant-planet-hosting star throughout its main-sequence lifetime, selecting a relatively late time t to compensate for mass ejected prior to the gas-clearing phase in each system. Adopting a 50% occurrence rate for this population of planets, we approximate that the average rate of mm-sized planetesimal ejection is $0.74\,M_\oplus\,\mathrm{star}^{-1}$.

In order to connect the mass ejected per star to an ISO spatial number density n, we must first determine the background stellar number density n_*. We adopt the local midplane stellar mass density $\rho_* = 0.043 \pm 0.004\,M_\odot\,\mathrm{pc}^{-3}$ found in McKee et al. (2015), then estimate the average stellar mass in the solar neighborhood using the sample of 75 known stars (including brown dwarfs) within 5 pc of the Sun. For stars without measured masses, we apply the piecewise mass-luminosity relation from Eker et al. (2015). We obtain an average stellar mass $M \sim 0.35 M_\odot$, which translates to stellar number density $n_* \approx 0.12\ \mathrm{stars\ pc}^{-3}$.

Combining n_* with our average planetesimal mass ejection rate per star, we find projected ISO mass density $\rho_{\rm ISO} \sim 0.09\,M_\oplus\,{\rm pc}^{-3}$ free-floating in space. This means that, for number density distribution dn/dr with boundaries $r_{\rm min,obs} = 867\,\mu$m and $r_{\rm max,obs} = 917\,\mu$m,

$$\int_{r_{\rm min,obs}}^{r_{\rm max,obs}} \frac{dn}{dr} m(r) dr = 0.09\,M_\oplus\,{\rm pc}^{-3}. \tag{4.4}$$

We use $m(r) = \frac{4}{3}\pi r^3 \rho$ with $\rho = 1\,{\rm g\,cm}^{-3}$ and fit dn/dr with a power law radius distribution of the form

$$\frac{dn}{dr} = C r^{-q}, \tag{4.5}$$

where q is the power law index and C is constant. We interpret previous estimates of n obtained from the detection rate implied by 'Oumuamua as $n(r \gtrsim r_{\rm O})$ where $r_{\rm O}$ is 'Oumuamua's effective radius, since larger ISOs would also be detectable when passing through the solar system at 'Oumuamua's distance. Equivalently,

$$\int_{r_{\rm O}}^{r_{\rm max}} \frac{dn}{dr} dr = n. \tag{4.6}$$

Combining and integrating Equations 4.4, 4.5, and 4.6, we parameterize our solution space with the function

$$f = \frac{4}{3}\pi\rho n \Big(\frac{1}{0.09 M_\oplus\,{\rm pc}^{-3}}\Big)\Big(\frac{1-q}{4-q}\Big)\Big[\,\frac{r_{\rm max,obs}^{4-q} - r_{\rm min,obs}^{4-q}}{r_{\rm max}^{1-q} - r_{\rm O}^{1-q}}\,\Big],$$

which is equal to unity when Equations 4.4, 4.5, and 4.6 are simultaneously satisfied. Using f, we solve for power law indices q over a range of $r_{\rm max}$, with $r_{\rm O} = 55\,{\rm m}$ from Jewitt et al. (2017).

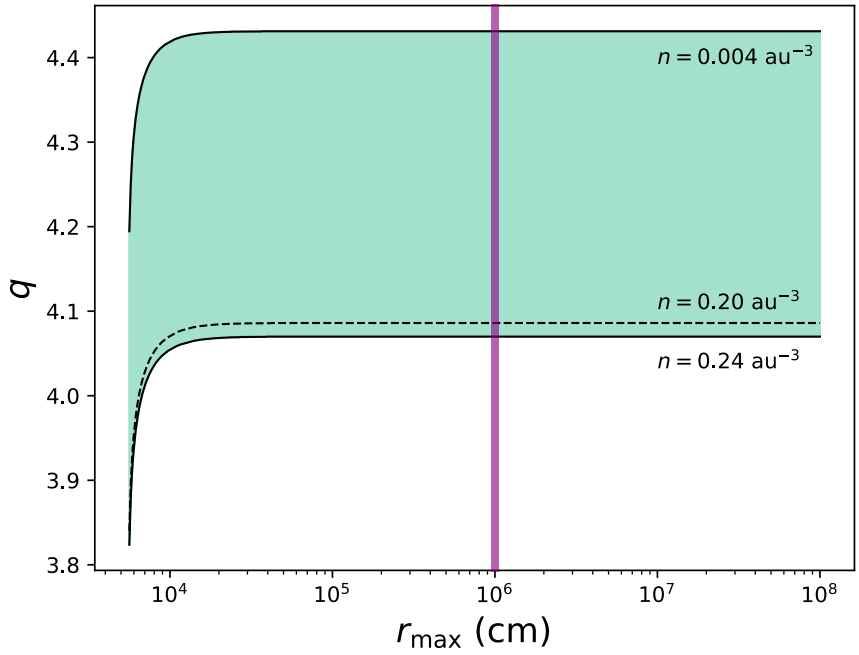

Figure 4.3: Solutions to the power law index q as a function of $r_{\max}$ for the full range of n estimates given in Portegies Zwart et al. (2018). The Do et al. (2018) estimate $n = 0.2$ au^{-3} is indicated with a dotted line. The purple slice of solutions at $r_{\max} = 10^6$ cm displays the cross-section corresponding to Figure 4.4.

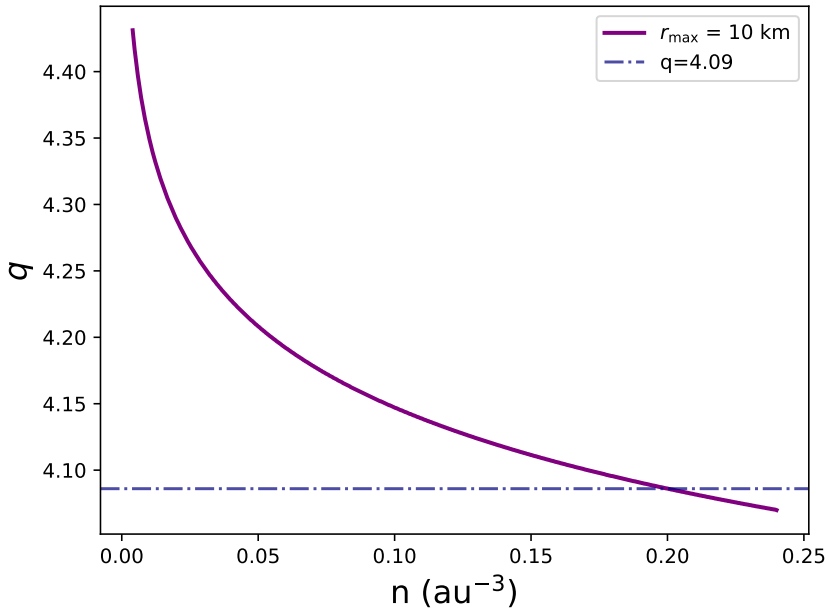

Figure 4.4: Power law solutions q shown as a function of n. Here, we set $r_{\rm max} = 10\,{\rm km}$, noting that q varies little with changes in $r_{\rm max}$ (see Figure 4.3). The dot-dashed blue line denotes the solution where $n = 0.2\,{\rm au}^{-3}$.

The resulting solutions are displayed in Figure 4.3, indicating that, for a given n, our power law solutions are robust to changes in $r_{\rm max}$ for an ISO distribution in which $r_{\rm max} \not\approx r_O$. We take a vertical slice through Figure 4.3 at $r_{\rm max} = 10^6\,{\rm cm}$ to display q as a function of n in Figure 4.4. Ultimately, we find that high power law indices $q > 4$ are required to fit both the n values implied by 'Oumuamua and our simulated mass ejection rates. Adopting $q = 4.09$, this distribution corresponds to a total mass of ejected ISO material $\sim 24 M_\oplus$ per DSHARP-type system for ISOs in the size range $10^{-3}\,{\rm cm} \leq r \leq 10^6\,{\rm cm}$.

Another potential constraint for this distribution is the number density $n_{\rm CNEOS} = 10^{6^{+0.75}_{-1.5}}\,{\rm au}^{-3}$ of $r_{\rm CNEOS} \gtrsim 0.45$ m interstellar objects implied by the candidate interstellar meteor identified from the Center for Near-Earth Object Studies (CNEOS) bolide catalog in Siraj & Loeb (2019). We note that Devillepoix et al. (2018) reported that the United States government sensors used to discover this object are generally unreliable for orbit calculations. We

include the candidate CNEOS meteor for completeness and primarily focus on determining whether the object is consistent with our analysis up to this point. If the CNEOS meteor originates from the same underlying distribution as 'Oumuamua, it would require

$$\int_{r_{\rm CNEOS}}^{r_{\rm max}} \frac{dn}{dr} dr = n_{\rm CNEOS}. \tag{4.7}$$

Adopting $q = 4.09$, corresponding to $n = 0.2\,{\rm au}^{-3}$ from Do et al. (2018), we obtain number density $n = 5.63 \times 10^5\,{\rm au}^{-3}$ for objects with $r_{\rm CNEOS} \leq r \leq r_{\rm max}$. This value is firmly within the error bars for $n_{\rm CNEOS}$ from Siraj & Loeb (2019), indicating that our results are consistent with those implied by the candidate CNEOS meteor.

4.5 Observability with LSST

Although the first ISO was not confirmed until 2017, anticipative estimates for the detectability of such objects have been made on several earlier occasions. To account for the prior non-detection of ISOs, early assessments predicted an interstellar comet number density $n \sim 10^{13}$ pc^{-3} for objects with $r \geq 1$ km (McGlynn & Chapman, 1989). Later, Jewitt (2003) projected that, if all stars eject 10^{13} comets with $r \geq 1\,{\rm km}$, approximately 0.3 such interstellar interlopers should pass within 5 au of the sun – the approximate detection limit of Pan-STARRS – each year.

Moro-Martín et al. (2009) followed up this study by calculating the expected detection rate by LSST of inactive extrasolar comets passing through the solar system at distances greater than 5 au. They found that the probability of LSST detecting an inactive interloper during its 10-year lifetime is $\sim 0.01\% - 1\%$, and they ultimately concluded that LSST will likely not observe even one such object. Cook et al. (2016) then updated prior estimates by numerically simulating the detectability of interstellar asteroids within 5 au of the sun,

concluding that an optimistic estimate would result in 1 interstellar object detected during LSST's lifetime. Since 'Oumuamua's discovery with Pan-STARRS, Trilling et al. (2017) used a scaling argument to project that the detection rate of analogous interlopers will increase from 0.2 yr^{-1} with Pan-STARRS to 1 yr^{-1} with the advent of LSST.

Leveraging the power law radius distributions obtained in Section 4.4, we independently estimate the expected detection rate by LSST for protostellar disk-ejected ISOs. Adopting single-frame magnitude limit $m_{\mathrm{LSST}} \sim 24$ (LSST Science Collaboration et al., 2009) and cometary geometric albedo $p_V = 0.04$ (Lamy et al., 2004), we calculate the smallest ISO radius observable with LSST at a given distance d, given in au, using

$$r_{\mathrm{min,LSST}} = \frac{1\,\mathrm{au}}{\sqrt{p_v}}\, 10^{\frac{1}{5}(m_\odot - H)} \tag{4.8}$$

$$H = m_{\mathrm{LSST}} - 5\log_{10} d, \tag{4.9}$$

where $m_\odot = -26.74$ is the magnitude of the Sun and H is the absolute magnitude of the ISO. Our calculation implicitly assumes that all ISOs are observed at opposition; however, in practice the solar angle and frequency of sky coverage will also play an important role in the final ISO detection rate, meaning that our rates should be treated as upper limits.

At each of 10^5 evenly spaced circular annuli around the Earth, from $0 - 40$ au, we sum over the number density profile from $r_{\mathrm{min,LSST}}$ to r_{max} to obtain the total observable ISO number density n_{ISO} within that annulus. Combining n_{ISO} with the surface area of each annulus σ, we calculate the total detection rate of ISOs as

$$R = \sum n_{\mathrm{ISO}} \sigma v, \tag{4.10}$$

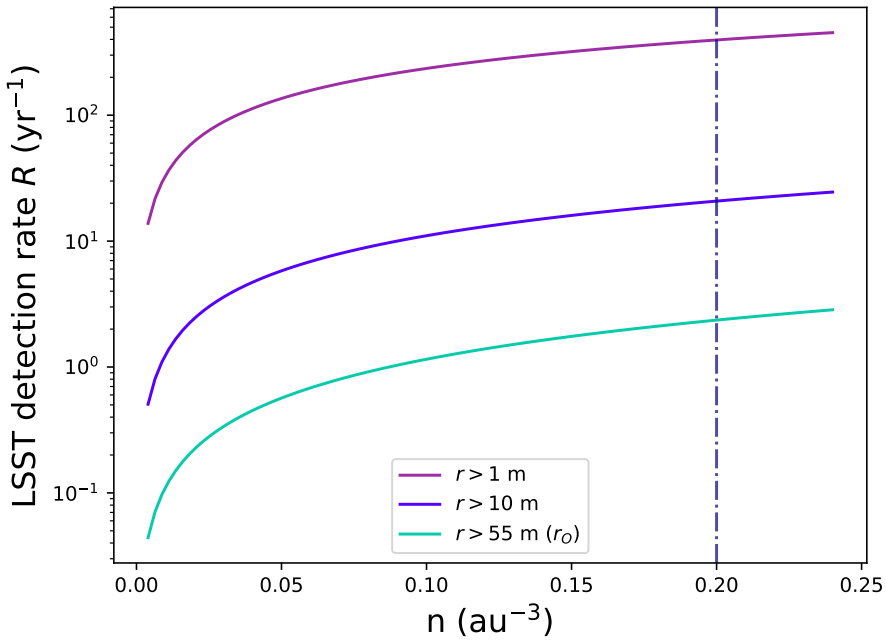

Figure 4.5: LSST detection rates as a function of n for minimum object radii 1 m, 10 m, and $r_O = 55\,\mathrm{m}$.

summing over all annuli. We take 'Oumuamua's velocity at infinity, 26 km/s (Mamajek, 2017), as our representative v. Completing this calculation with reference minimum radius thresholds $r > 1\,\mathrm{m}$, $r > 10\,\mathrm{m}$, and $r > 55\,\mathrm{m}$, we obtain the results outlined in Figure 4.5.

We find that, if 'Oumuamua is an ejected cometary planetesimal drawn from an isotropic power law distribution, LSST will find several ISOs each year of 'Oumuamua's size or larger, as well as up to hundreds of ISOs per year with $r > 1\,\mathrm{m}$.

4.6 Sources of Uncertainty

Zhu et al. (2019) recently found that protoplanetary disk masses may be heavily underestimated due to the assumption that the disks are optically thin. As a result, in the case that these disks are optically thick, our mm mass ejection rates may be underestimated, leading to an underestimate in the power law slope.

However, it may instead be the case that our disk masses are overestimated. Our three sample disks are located around relatively massive stars ($0.83M_{\odot}$, $1.78M_{\odot}$, and $2.04M_{\odot}$ for AS 209, HD 143006, and HD 163296, respectively), meaning that treating these systems as 'typical' may instead lead to an overestimate of the total mass of ejected material. This would accordingly correspond to an overestimate of the LSST detection rate, particularly on the lower end of the ISO size distribution.

Furthermore, the number density of ISOs, n, may be locally enhanced at the present time due to the Sun's current proximity to the Galactic midplane: $z_{\odot} \sim 17 \pm 5$ pc (Karim & Mamajek, 2016), as compared to its maximum height above the midplane $49 - 93$ pc (Bahcall & Bahcall, 1985). If the Sun is currently passing through a low-dispersion, 'cold' population of ISOs, the apparent number density as observed from the solar system may be higher than the bulk average density throughout the Galaxy. Because we do not yet have strong constraints on the scale height of ISOs, we cannot rule out this possibility, which suggests that the true n may be lower than most current estimates. This could result in a steeper power law size distribution and a correspondingly lower LSST detection rate for large ISOs.

It is also quite possible that the distribution of interstellar comets does not follow a simple power law. Moro-Martín (2018), after sampling a wide range of possible models, finds that 'Oumuamua is likely not representative of an isotropic background population. Though our present model is already overconstrained and therefore does not warrant the inclusion of additional free parameters, we acknowledge that a power law fit may not fully capture the true characteristics of the ejected ISO population. We may be overestimating the number density of large ISOs if, for example, the collisional evolution of solid material in extrasolar systems typically grinds the vast majority of larger bodies into dust prior to ejection.

Lastly, the detection rates reported in Section 4.5 may be modestly increased by gravitational focusing and comet brightening from passage near the sun, each of which enhances

observability of the ISO population.

4.7 Conclusions

The discovery of the first interstellar interloper, 'Oumuamua, has already provided exceptional insight to inform our understanding of planetary systems while simultaneously presenting new puzzles. In our work, we have reconciled simulations of the observed DSHARP planet-disk systems with the detection of 'Oumuamua to constrain the range of possible size distributions for interstellar objects ejected through interactions with circumstellar giant planets. We conclude that the population of long-period giant planets suggested by the DSHARP sample is capable of ejecting the population of free-floating planetesimals implied by 'Oumuamua's detection. Furthermore, 'Oumuamua is consistent with an origin within a population of ISOs following a single power law radius distribution, resulting in an anticipated LSST detection rate ranging from just a few detections per year for 'Oumuamua-sized ISOs to over $100\,\mathrm{yr}^{-1}$ for ISOs with $r > 1$ m. Future observations of interstellar objects hold tremendous potential to answer long-standing questions about not only the range of processes taking place in extrasolar systems, but also the population statistics of long-period giant exoplanets, the diversity of small-body populations throughout the Galaxy, and the evolutionary path of the solar system itself.

4.8 Acknowledgements

We thank the anonymous referee for helpful comments that substantially improved the manuscript. We also thank Matt Holman for insightful input as the manuscript was being finalized. M.R. is supported by the National Science Foundation Graduate Research Fellowship Program under Grant Number DGE-1752134. This material is also based upon work supported by the National Aeronautics and Space Administration through the

NASA Astrobiology Institute under Cooperative Agreement Notice NNH13ZDA017C issued through the Science Mission Directorate. We acknowledge support from the NASA Astrobiology Institute through a cooperative agreement between NASA Ames Research Center and Yale University. Simulations in this paper made use of the `REBOUND` code which can be downloaded freely at `http://github.com/hannorein/rebound`. This research also made use of the `numpy` (Oliphant, 2006), `matplotlib` (Hunter, 2007), and `seaborn` (Waskom et al., 2014) Python packages.

Part III

The Star-Planet Connection

Chapter 5

Stellar Characterization of Keck/HIRES Spectra with The Cannon

Originally published as:

Rice, M. & Brewer, J. 2020, The Astrophysical Journal, 898, 119

5.1 Abstract

To accurately interpret the observed properties of exoplanets, it is necessary to first obtain a detailed understanding of host star properties. However, physical models that analyze stellar properties on a per-star basis can become computationally intractable for sufficiently large samples. Furthermore, these models are limited by the wavelength coverage of available spectra. We combine previously derived spectral properties from the Spectroscopic Properties of Cool Stars (SPOCS) catalog (Brewer et al., 2016) with generative modeling using *The Cannon* to produce a model capable of deriving stellar parameters ($\log g$, $T_{\rm eff}$, and $v \sin i$) and 15 elemental abundances (C, N, O, Na, Mg, Al, Si, Ca, Ti, V, Cr, Mn, Fe, Ni, and Y) for stellar spectra observed with Keck Observatory's High Resolution Echelle Spectrometer (HIRES). We demonstrate the high accuracy and precision of our model, which takes just $\sim$3 seconds to classify each star, through cross-validation with pre-labeled spectra from the SPOCS sample. Our trained model, which takes continuum-normalized template spectra as its inputs, is publicly available at `https://github.com/malenarice/keckspec`. Finally, we interpolate our spectra and employ the same modeling scheme to recover labels for 477 stars using archival stellar spectra obtained prior to Keck's 2004 detector upgrade, demonstrating that our interpolated model can successfully predict stellar labels for different spectrographs that have (1) sufficiently similar systematics and (2) a wavelength range that substantially overlaps with that of the post-2004 HIRES spectra.

5.2 Introduction

High-precision spectroscopic stellar characterization is critical to understand the host environments of planetary systems. The diversity of observed exoplanetary systems suggests a wide range of system properties influenced by a correspondingly wide range of formation

environments. Furthermore, several of the most prominent current methods to study exoplanets rely upon indirect measurements, probing how planets gravitationally perturb their host stars (radial velocity measurements; e.g. Lovis & Fischer, 2010; Butler et al., 2017; Cumming, 2004)) and/or alter the time-series photometry of their host star (transit and phase curve measurements; e.g. Haswell, 2010; Morello et al., 2014; Esteves et al., 2013; Cowan et al., 2013; Borucki et al., 2009; Southworth, 2008). To appropriately disentangle the properties of planets from their host stars' signals, and to interpret the relationship between these planets and their formation environments, it is necessary to robustly determine the properties of the host stars within these systems (Brewer et al., 2018).

Several existing catalogs report the derived properties of stars based on different spectroscopic surveys (e.g. Hinkel et al., 2014, and references therein). One example, the Spectroscopic Properties of Cool Stars catalog (SPOCS; Valenti & Fischer, 2005), analyzed nearly 2,000 Keck HIRES spectra of over 1,000 F, G, and K dwarfs that were obtained as part of the Keck, Lick, and AAT planet search programs (Cumming et al., 1999; Fischer et al., 1999; Butler et al., 2003). The precise stellar parameters and 5 elemental abundances (Fe, Si, Ti, Na, Ni) obtained in this survey demonstrated for the first time the positive correlation between the frequency of close-in giant planets and host star metallicity (Fischer & Valenti, 2005), providing a landmark constraint towards a more cohesive understanding of planet formation (e.g. Robinson et al., 2006).

The bulk of the spectra analyzed in Fischer & Valenti (2005) were obtained with the Keck HIRES spectrograph (Vogt et al., 1994) prior to a detector upgrade that took place in August 2004. The newer three-chip detector, once installed, proved advantageous: it allowed for more extensive spectral analyses, including higher-precision gravity measurements (Brewer et al., 2015) and abundances for 15 elements (Brewer et al., 2016; Brewer & Fischer, 2018) obtained using the stellar modeling program Spectroscopy Made Easy (SME; Valenti &

Piskunov (1996)). These improved parameters enabled the measurement of more precise masses and radii for observed stars and their accompanying planets (Brewer et al., 2016).

SME incorporates empirical atomic and molecular line data to develop physically motivated synthetic spectra, which can then be compared to stellar data to fit for parameters of interest. However, the computational expense of SME becomes prohibitively high for large stellar samples, since each individual star typically takes $\sim$14 hours to model in SME. Furthermore, the analysis techniques used to develop uniform catalogs in Brewer et al. (2016) and Brewer & Fischer (2018) rely on the extended wavelength coverage of the newer three-chip detector and, as a result, cannot be applied to the older (pre-2004) spectra.

Keck HIRES includes an iodine cell used to extract the radial velocity signals of planets orbiting stars (Butler et al., 1996). For each observed star, a 'template' spectrum of only the star is obtained without the iodine cell in place. To measure the reflex motion of the star, the same star is then observed through the iodine cell, imprinting a rest-frame iodine spectrum onto the stellar spectrum. Planets produce a radial velocity shift in the star, manifest in the observed spectra as a slight offset of the stellar spectrum relative to the iodine lines. The template spectrum of the star at different radial velocity offsets is convolved with a reference spectrum of the iodine cell and an instrumental profile in order to determine the observed radial velocity. The iodine-free template spectra can also be used to deduce properties of the observed stars.

Not all stars observed before 2004 had another template spectrum taken afterwards, and many stars were dropped early on from radial velocity surveys if, after a few observations, they were found to have a root-mean-square (RMS) scatter below the precision of the spectrograph used at the time ($\sim$3 m/s). Although those stars were deemed "planetless" based on the absence of high-amplitude signals, much-improved spectrographs in the modern era currently reach precision an order of magnitude lower than these prior surveys (e.g. the EXtreme PREcision Spectrograph, EXPRES (Jurgenson et al., 2016; Petersburg et al.,

2020); the Echelle SPectrograph for Rocky Exoplanets Search and Stable Spectroscopic Observations, ESPRESSO (Pepe et al., 2013); and the upcoming NN-explore Exoplanet Investigations with Doppler spectroscopy spectrograph, NEID (Schwab et al., 2016)), meaning that many of these stars are again targets of interest in current planet searches.

Fortunately, new spectral analysis techniques can be used to address both problems described above: the computational expense of synthetic spectral models, as well as the dependence of these models upon a specific wavelength coverage. *The Cannon* (Ness et al., 2015; Casey et al., 2016) is a supervised learning algorithm that determines stellar labels by identifying correlations on a pixel-by-pixel basis. By "learning" the properties of a uniformly classified dataset of stars, *The Cannon* can efficiently and accurately transfer these learned correlations to a new set of stars spanning a similar parameter space.

Previous studies have applied *The Cannon* to obtain stellar parameters and abundances using spectra from the Apache Point Observatory Galactic Evolution Experiment (APOGEE) as part of the Sloan Digital Sky Survey (SDSS; e.g. Ness et al., 2016; Abolfathi et al., 2018; Holtzman et al., 2018); from the Galactic Archaeology with HERMES (GALAH) survey (e.g. Buder et al., 2018; Kos et al., 2018); from the RAdial Velocity Experiment (RAVE; Casey et al., 2017), and from the Large Sky Area Multi-Object Fiber Spectroscopic Telescope (LAMOST; Ho et al., 2017), among others. Behmard et al. (2019) also completed an analysis of 141 cool stars observed with Keck HIRES, focusing on the subset of K and M stars with $T_{\rm eff} < 5200$ K, to estimate precise effective temperatures ($T_{\rm eff}$), stellar radii (R_*), and metallicities ([Fe/H]). We complete a more extended study here, applying the full Brewer et al. (2016) SPOCS catalog to develop a new model predicting 18 stellar labels for dwarfs and subgiants spanning $T_{\rm eff} = 4700 - 6674$ K.

In this paper, we first train *The Cannon* using the Brewer et al. (2016) SPOCS catalog to produce a model which rapidly and reliably retrieves 18 precisely determined stellar labels ($\log g$, $T_{\rm eff}$, $v \sin i$, and 15 elemental abundances: C, N, O, Na, Mg, Al, Si, Ca, Ti, V, Cr,

Mn, Fe, Ni, and Y) for input post-2004 Keck HIRES stellar spectra. This model is made available as a tool for public use and is applicable to all current and future Keck HIRES spectra taken since the 2004 detector upgrade. While it covers a considerably smaller temperature range than `SpecMatch-Emp` (Yee et al., 2017), an empirical grid-based code designed to characterize HIRES spectra with $T_{\rm eff} \approx 3000 - 7000$ K, our model returns 18 precisely determined stellar labels in comparison with 3 returned by `Specmatch-Emp` ($T_{\rm eff}$, R_*, and [Fe/H]).

We then interpolate the SPOCS spectra to the pre-2004 detector's wavelength range and again apply *The Cannon* to develop a homogeneously analyzed catalog of these 18 stellar labels for 477 archival Keck stars, using overlapping stars observed in both the pre- and post-2004 samples as our test set. We demonstrate that, given a reliably labeled training set, *The Cannon* can be used to efficiently obtain high-precision stellar parameters from large-scale spectroscopic surveys, with a combined speed and accuracy unattainable for more time-intensive, single-object stellar classification methods.

5.3 Methods: The Cannon

A detailed overview of the methods applied in *The Cannon* can be found in Ness et al. (2015) and Casey et al. (2016). We briefly review these methods here, and we refer the reader to these articles for a more in-depth description of the code. In this work, we apply the version of *The Cannon* described in Casey et al. (2016).

In short, *The Cannon* develops a generative spectral model, described by coefficient vectors θ_j corresponding to each modeled parameter at each pixel j, to characterize the relationship between flux and label values in each pixel. This model is trained using an input dataset with known labels spanning the same parameter space as the stars that are characterized. The coefficient vector of each pixel is determined by minimizing the summed log likelihood

function

$$\theta_j, s_j^2 \leftarrow \underset{\theta, s}{\text{argmin}} \Big[\sum_{n=0}^{N-1} \ln p(y_{jn}|\theta, l_n, s^2) \Big], \tag{5.1}$$

where the pixel's log likelihood is given by

$$\ln p(y_{jn}|\theta, l_n, s^2) = \frac{[y_{jn} - \theta \cdot f(l_n)]^2}{s^2 + \sigma_{jn}^2} + \ln(s^2 + \sigma_{jn}^2). \tag{5.2}$$

Here, l_n is a vector containing the star's labels (in our case, 18 labels per star), y_{jn} is the flux at a given pixel j and stellar spectrum n, and $f(l_n)$ is a vectorizing function that determines the form of our model (in our case, a quadratic polynomial with all cross-terms included). The noise is characterized by $s^2 + \sigma_{jn}^2$, where σ_{jn} encapsulates the reported instrumental and Poisson uncertainty, while$s$ provides the intrinsic scatter of the model.

Ultimately, this produces a generative model that describes the probability density function of flux at each wavelength as a function of stellar labels. Once the coefficient vectors have been obtained through supervised learning with a set of pre-labeled input stars, the generative model can be applied to a new set of spectra to transfer stellar labels based on the trained model's vectorizing function and coefficients. This step is accomplished by optimizing Equation 5.3, which sums over all j pixels in the spectrum, to find the label vector l for each test star m.

$$l_m \leftarrow \underset{l}{\text{argmin}} \Big[\sum_{j=0}^{J-1} \ln p(y_{jm}|\theta_j, l, s_j^2) \Big], \tag{5.3}$$

In this way, *The Cannon* "learns" the characteristics of the stars that it is trained on in order to efficiently transfer labels to a new set of stars with similar properties. Ness et al. (2015) demonstrated that *The Cannon* provides robust results for low signal-to-noise spectra when

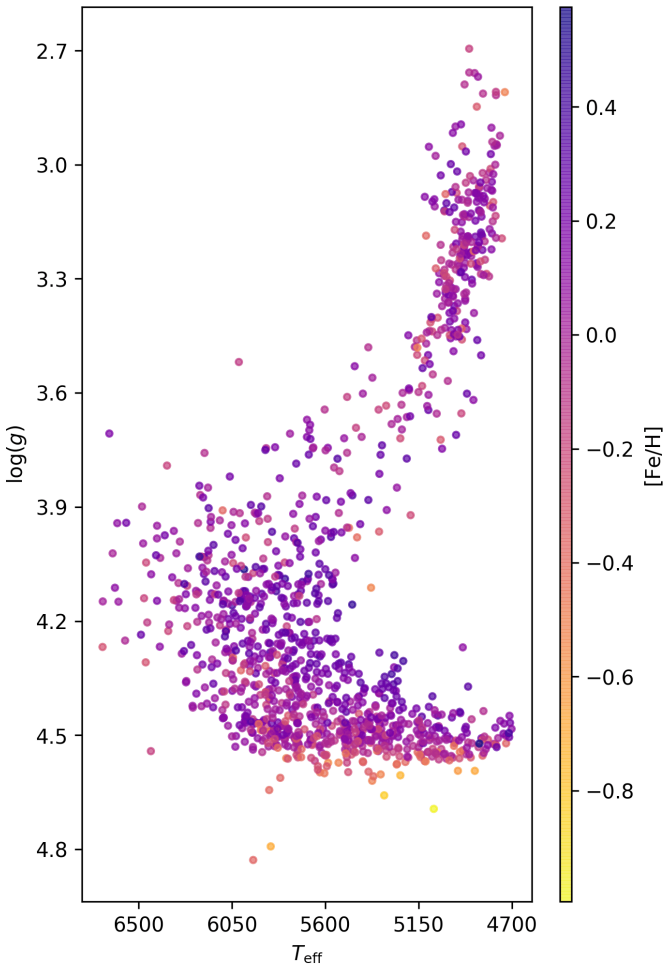

Figure 5.1: Distribution of our final sample of 1202 stars, colored by metallicity.

trained upon higher resolution spectra.

While the training and validation step of this label transfer process can be time-intensive, a well-characterized model, once trained, can be easily saved and applied to new datasets for rapid characterization. We employ this property of *The Cannon* to develop a new open-source code applicable to current and future Keck HIRES stellar spectra in Section 5.5.

5.4 Data Selection & Processing

Throughout this work, we use the SPOCS dataset for training and model validation testing. Our pre-labeled dataset includes ∼3800 HIRES spectra of ∼1600 objects, with labels from Brewer et al. (2016) obtained using the SME software combined with atomic and molecular line data from the Vienna Atomic Line Database 3 (VALD-3; see Brewer et al. (2016) for a comprehensive set of SPOCS line list contributors). All stars were observed with HIRES in the red configuration, with the iodine cell out and $R \sim 70,000$. From this original sample, we removed all spectra with one or more of the following properties:

- Labeled with 'NGC' (deep sky objects; not individual stars)
- Flagged as 'bad'
- Signal-to-noise ratio (SNR) < 100

Together, these cuts reduced our pre-labeled sample to 1202 stars with 2018 spectra, where the distribution of stellar parameters in our final sample is displayed in Figure 5.1. To optimize our inputs, we selected only the highest-SNR spectrum from each star, resulting in 1202 total spectra.

An automated version of the data reduction pipeline used to reduce the SPOCS data set is available for public use online (Butler et al., 1996; Howard et al., 2010).[1] Because the Keck HIRES instrument is an echelle spectrograph, each of our sample spectra contains 16 separate echelle orders. Each echelle order initially included a blaze function convolved with the instrumental response function, leading to an underlying continuum upon which the spectral features of the observed star were imprinted. To deduce the shape of the continuum, each spectrum was individually fit with iterative polynomials using the algorithm described in Valenti & Fischer (2005). These continuum fits were then divided out of the

[1] https://caltech-ipac.github.io/hiresprv/index.html

corresponding spectra to obtain a set of continuum-normalized spectra with baseline flux set to unity.

Our selected spectra were initially slightly shifted relative to each other in wavelength space due to the varying line-of-sight velocities of stars within our sample, leading to slight shifts in the wavelength solutions. To account for this effect, we interpolated all spectra to the same wavelength grid for a one-to-one comparison of each pixel across spectra. This grid was determined by finding the maximum wavelength range spanned by all spectra in our sample and keeping the total number of data points in each spectrum the same, and we carried out this process independently for each echelle order. Each final spectrum includes 16 echelle orders each with 4021 pixels, resulting in a total of 64,336 data points per star.

5.5 Developing a Model: Current and Future Keck HIRES Spectra

Keck's current HIRES spectrograph has been in use since 2004 to search for and study extrasolar planets. Paired with the 10-meter Keck I telescope, HIRES is a powerful tool to probe dim stars, such as many *Kepler* planet hosts, that are prohibitively faint for study with other telescopes.

Many stars that were not part of the original SPOCS catalog have been and continue to be observed with HIRES. Thus, a reliable model to extract stellar properties from HIRES spectra is crucial. We describe here our methods in developing a new, open-source model that rapidly delivers 18 stellar labels, including $T_{\rm eff}$, $v \sin i$, log$g$, and 15 elemental abundances (C, N, O, Na, Mg, Al, Si, Ca, Ti, V, Cr, Mn, Fe, Ni, and Y).

5.5.1 Model Selection Framework

Throughout our model testing phase, we used three different 80%/20% train/test splits to check our model performance. These splits were randomly selected at the beginning of our testing, and we used the same divisions at each progressive test step for a direct comparison between models. We ran three tests at each step to verify that any observed differences in performance were due to generalizable changes in the model performance, rather than stochastic variations in the selected test/train samples.

After training our model on the 80% training set, we benchmarked its performance by (1) checking the model's ability to recover the input training set labels and (2) cross-validating using our 20% test set with known "true" labels. The first of these benchmarks was used only to verify that the model was performing correctly, while we report all results based on the second benchmark, which provides an independent check for our model performance.

While exploring various configurations to optimize our model performance, we ran tests on individual echelle orders as well as the combined 16 orders. Our reasons for this were twofold.

First, testing with a limited wavelength range is far less computationally expensive than with the full range, and, as a result, we were able to complete a more extensive analysis in the single-echelle-order case. This informed our more computationally-intensive tests that included all 16 echelle orders, allowing us to more quickly optimize our models and to consider a wider range of possible adjustments.

Second, some orders contain particularly important spectroscopic lines – for example, the gravity-sensitive magnesium Ib triplet at 5183, 5172, and 5167 Å and the forbidden oxygen line at 6300 Å – and should therefore perform particularly well to extract associated parameters. Beyond producing a model useful for the characterization of Keck HIRES

spectra, we were interested to determine (1) from which wavelength ranges *The Cannon* obtained the most useful information and (2) with what precision a smaller wavelength range with more concentrated information could determine our stellar parameters of interest. Thus, we optimized both a single-echelle-order model and a model including all 16 echelle orders. We report our results in both cases but make only the all-orders model publicly available due to its improved performance over the single-order model.

Our metric for model performance is a χ^2 test in which we minimize the function

$$\chi^2 = \frac{1}{N}\sum_{i=1}^{N}\frac{(x_i - E_i)^2}{|E_i|}. \tag{5.4}$$

Here, N is the number of values in the sample, and E_i and x_i are the expected and predicted value, respectively, for the parameter at each step in the summation. By dividing the mean squared error by the expected value of each parameter, we normalize our performance metric to avoid biases from the unequal scales of each label. Our adopted χ^2 thus measures a modified percent deviation from the expected value of each label. We determine the average χ^2 across our three models after implementing each new adjustment and compare that value with the previous best-performing χ^2 to determine which adjustments to implement in our final model.

5.5.2 Outlier Removal

From initial testing, we discovered that there existed several spectra with outlier stellar labels within our training/testing sample. As a result, we ran tests in which we removed outliers using several different thresholds in search of a threshold that improved the accuracy of our model without removing enough spectra to degrade the model's performance. Where q_1 and q_3 represent the first and third quartile of the data, respectively, the interquartile range (IQR) is given by

$$\mathrm{IQR} = |q_3 - q_1|. \tag{5.5}$$

We define an outlier as a value that falls more than a factor $(x_O \cdot \mathrm{IQR})$ below q_1 or above q_3, where x_O determines the stringency of our requirement for a data point to be classified as an outlier. We tested three different values $x_O = 1.5$, 3, and 10 (resulting in 59, 10, and 1 outlier stars removed from the sample, respectively), as well as the case in which no outliers are removed, to determine the optimal x_O value.

We found that, in the all-orders case, we obtained the lowest χ^2 with $x_O = 10$. Our best-performing single-order model was order 10, spanning wavelength range $5355 - 5445$ Å, with no outliers removed. In Table 5.1, we compare the total number of atomic lines returned by VALD-3 in the best- and worst-performing single-order wavelength ranges. While the number of atomic lines is not a perfect metric to compare the information content of different wavelength orders due to the varying strength of atomic lines, as well as the presence of molecular lines, a zeroth-order comparison between these two wavelength orders reveals that our best-performing wavelength order contains multiple known atomic lines associated with each element. In contrast, this is not the case for our worst-performing wavelength order, which contains no known Na lines and thus may perform particularly poorly in returning the [Na/H] label.

In general, each individual order provided systematically better results with no outliers removed than with $x_O = 3$ or $x_O = 10$, although the case with $x_O = 1.5$ provided similar results. We chose to move forward in testing with the top-performing order as a representative wavelength range that performs well on its own, noting that the stochastic variation in performance due to the random train/test split is larger than the margin of improvement obtained from using this order rather than the second-best-performing order.

Element	$5355 - 5445$ Å	$6312 - 6418$ Å
Fe	1414	1253
C	69	53
N	12	32
O	46	31
Na	3	0
Mg	8	8
Al	12	24
Si	56	72
Ca	134	340
Ti	224	208
V	380	320
Cr	654	683
Mn	385	395
Ni	486	582
Y	81	64

Table 5.1: Number of atomic lines for each analyzed element in our best-performing single wavelength order ($5355 - 5445$ Å) compared with our worst-performing single wavelength order ($6312 - 6418$ Å).

5.5.3 Tuning the Model

Next, we consider a range of possible model adjustments to determine an optimal configuration for our final model. To cover a breadth of model configurations, we use the single, best-performing individual wavelength order found in Section 5.5.2 for initial testing purposes (order 10; $5355 - 5445$ Å, with no outliers removed). Once we have run these tests on an individual order, we use the results to inform further testing with our full wavelength range.

In each subsection, we test different approaches to continuum normalization (§5.5.3.1), telluric contamination (§5.5.3.2), label censoring (§5.5.3.3), and regularization (§5.5.3.4). We run three configurations of every test setup, each with a different randomly drawn train/test split, to ensure that our results are generalizable across samples spanning a similar range of labels.

We caution that the hyperparameters selected to optimize our best-performing single-order model do not necessarily translate to the best possible model when applied to all of our echelle orders together. Furthermore, we progressively build upon our model adjustments, accepting or rejecting changes to our base model in a set order. An exhaustive search for the single best-performing model would include all possible permutations of these model adjustments and would test all of these models with all echelle orders included. However, the computational expense of this exercise would be prohibitive with potentially diminishing returns. As a result, we operate under the assumptions that (1) a model that performs well for a single echelle order will also perform well with all orders included and (2) altering the order in which we apply model adjustments would not result in substantial improvements in our best-performing model.

5.5.3.1 Data-Driven Continuum Renormalization

The Cannon accepts continuum-normalized spectra with flux baseline set to unity as its inputs, and, as a result, the manner with which the continuum normalization is applied also affects the performance of *The Cannon*. With the goal of improving our continuum-fitting procedure while reducing the signal-to-noise dependence of our model, we tested the effects of applying data-driven continuum renormalization methods when optimizing our training setup.

We completed this process by finding the "true" continuum pixels in a data-driven manner. These pixels, each with a corresponding wavelength value, act as part of the continuum in that they (1) vary minimally with changes in the stellar label values and (2) return flux values close to unity in the spectral model's baseline spectrum, defined as the zeroth-order coefficient vector returned by *The Cannon*.

We first trained *The Cannon* using our initial continuum normalization, and we used this model to identify the pixels that varied the least with our four dominant stellar labels: log$g$, [Fe/H], $T_{\rm eff}$, and $v \sin i$. We selected some percentage $N\%$ of pixels with coefficients closest to zero for each label, then determined the set of overlapping pixels across all four labels. Finally, we applied a cut removing all pixels from this set that lay outside 1.5% of unity in the spectral model's baseline spectrum. To explore several possible model configurations, we used four different thresholds for pixel selection: N = 50, 60, 70 and 80. For these thresholds, the final percentage of pixels identified as "true" continuum pixels ranged from $\sim$3-6%, $\sim$9-18%, $\sim$11-21%, and $\sim$20-23%, respectively, with variations arising from random differences in our three train/test sets.

We applied two continuum renormalization schemes to each of these four cases to check for improvement in our model results: a polynomial continuum fit as well as a continuum fit composed of summed sine and cosine functions. To select the polynomial order used to

fit each spectrum, we tested polynomial fits with n, the number of free parameters, ranging from $n = 1 - 10$. For each spectrum, we chose the n value corresponding to the lowest reduced χ^2. We applied the fitting function described in Casey et al. (2016) for our summed sin/cos renormalization tests.

As in previous tests, we ran each case in three iterations and used average values from these iterations to quantify the model performance. Thus, we ran a total of 24 train/test models in this section: two functional forms (polynomial and sin/cos) for each of the four pixel selection thresholds, and three iterations of each combination.

Because our continuum renormalizations do not use all pixels in each spectrum – only the roughly $3 - 23\%$ that are selected as "true" continuum pixels – the edges of our spectra are not generally included within the fits. As a result, we renormalized only the pixels between the minimum and maximum "continuum" pixels, setting our data-driven continuum fits equal to unity outside of these bounds. Figure 5.2 shows sample fits for the $N = 70$ case, with the polynomial fit shown in green and the sin/cos fit in purple. Generally, as in Figure 5.2, the two fitting methods closely trace each other and deviate most in wavelength ranges with few identified continuum pixels.

We found that, for both the polynomial and sin/cos renormalization, the $N = 70$ case produced the lowest reduced χ^2 value, with polynomial renormalization providing the best results. Both of these cases showed improvements over our original test case with no continuum renormalization, while $N = 50, 60,$ and 80 each produced slightly degraded results. Thus, we chose to adopt the $N = 70$ implementation with polynomial continuum renormalization in our continued single-order tests moving forward.

Based on the promising results of this test, we also applied data-driven continuum renormalization to our full model with $x_O = 10$, testing the $N = 70$ case with both the sin/cos and polynomial renormalizations. Using the same three 80%/20% splits as our pre-tuned models

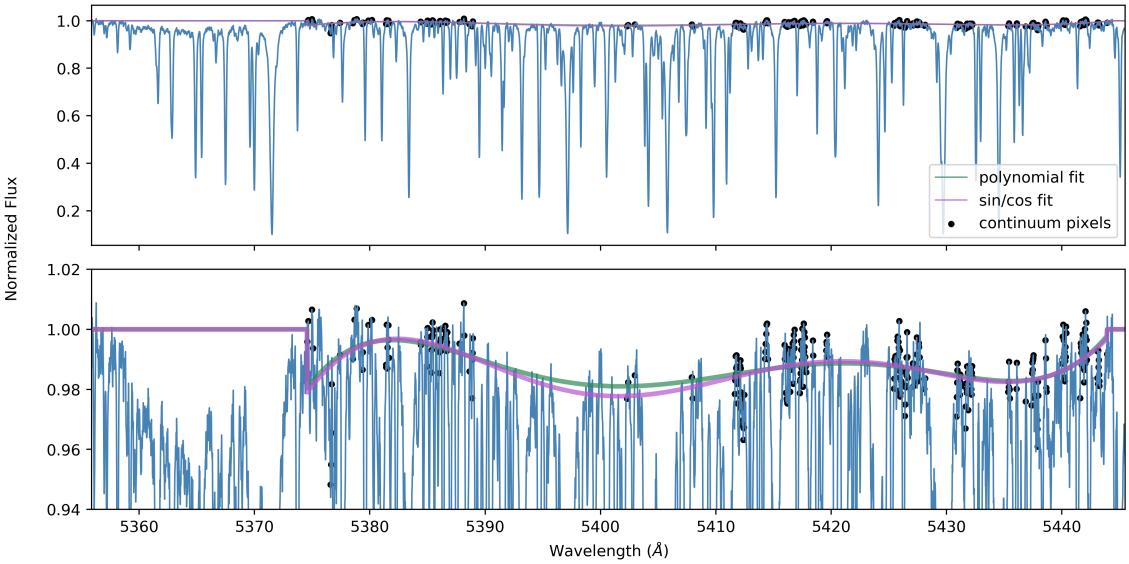

Figure 5.2: Sample $N = 70$ continuum renormalization fit over the spectrum of K0 star HD 22072 shown in blue. Here, the polynomial fit is shown in green while the sin/cos fit is in purple. Continuum pixels are denoted by black markers. The top panel shows the full spectrum over echelle order 10, while the bottom panel zooms in for a clearer comparison between continuum fits.

for a direct comparison, we again found that both renormalization schemes improved our results, and the polynomial renormalization provided the best results. As a result, we chose to include a polynomial continuum renormalization in our final version of the model and in continued tests.

5.5.3.2 Telluric Masking

Telluric lines are spectral imprints of the Earth's atmosphere superimposed onto all spectra taken by ground-based telescopes. The presence and variation of these lines over time can produce noise in a spectrum that is difficult to disentangle from the astrophysical signal of interest. Thus, our next step in improving our model performance is to mask out telluric lines to avoid introducing false correlations into our model.

In each spectrum, the locations of telluric lines remain stationary while the stellar lines are shifted to their rest-frame wavelengths using a barycentric correction and the radial velocity

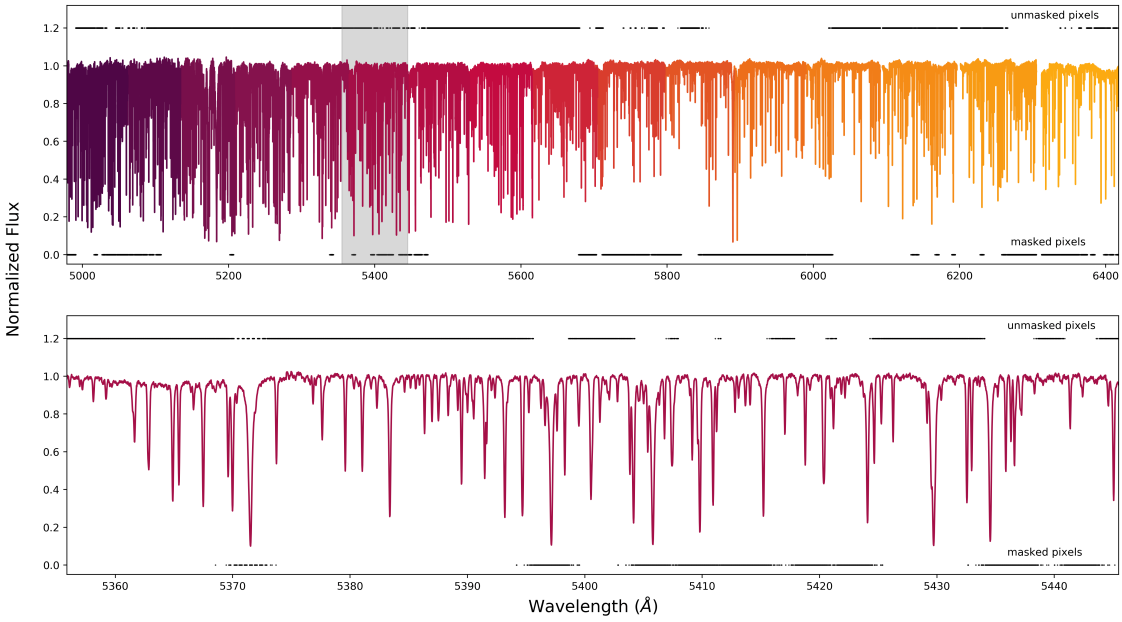

Figure 5.3: Top: Full continuum-renormalized spectrum of sample star HD 22072, with each of the 16 echelle orders shown in a different color. Both panels show the continuum renormalization with $N = 70$, corresponding to our all-orders best fit. The portion of the spectrum corresponding to the lower panel is highlighted in gray. Bottom: Zoom-in of only echelle order 10, ranging from $5355 - 5445$ Å. Black markers denote the telluric ("masked") pixels at the bottom of each panel, as well as the non-telluric ("unmasked") pixels at the top of each panel.

of the host star. As a result, the locations of the telluric lines do not perfectly align in every spectrum. To account for this effect, we determined the locations of all known tellurics in each spectrum and created a corresponding mask for each. Telluric masks were created by selecting pixels below 99% of the continuum in the NSO solar atlas telluric spectrum (Wallace et al., 2011) and rescaling the masks to the resolution of our spectra.

We then combined these masks to create a final, uniform mask applied to all spectra. We visualize the resulting mask in Figure 5.3, where masked pixels are denoted with black markers below the spectrum, while the unmasked pixels are displayed above the spectrum for comparison.

The top panel of Figure 5.3 displays the telluric mask applied to all 16 wavelength orders

placed side-by-side, while the bottom panel zooms in to our best-performing single order. Every wavelength order is shown in a different color, and the sample spectrum has been continuum renormalized using the methods described in Section 5.5.3.1. With this method, we found that 40,218 of the full 64,336 pixels remained unmasked after telluric masking.

In our single best-performing echelle order alone, 1419 pixels of the full 4021 were masked. Despite this substantial masking, which removed roughly 35% of pixels from the model, our model's performance improved due to the removal of confusion from telluric signals. As a result, we continued to implement telluric masking in our ongoing single-order tests.

Because of the nonuniform distribution of telluric lines across echelle orders, testing the performance of a single telluric-masked order is insufficient to determine the overall effect of telluric masking on the performance of *The Cannon*. Thus, we also trained *The Cannon* on our three train/test splits using the full telluric-masked spectrum, with all 16 orders input together. Building upon our previous best case with $x_O = 10$ and polynomial renormalization, we found that *The Cannon* returned further improved results with the implemented telluric masking in place for all 16 orders. Thus, we continued to use this telluric masking in ongoing testing and in our final model configuration.

5.5.3.3 Censoring

Censoring allows the user to select which individual labels contribute to the model's flux in each pixel, providing a method to incorporate prior knowledge of known features that correlate with each label. We use a data-driven approach to apply censoring within our models in a similar manner to our continuum pixel selection implemented in Section 5.5.3.1. This allows us to circumvent problems arising from the use of individual element line lists, since, as illustrated in Ting et al. (2019), abundances may have complex correlations due to the presence of molecules in stellar atmospheres. By applying our data-driven methods, we remain agnostic to the cause of the observed correlations, instead focusing only on the

ability of our model to reproduce these features.

To test different thresholds, we first trained our model with all pixels included. Then, for each label, we selected the top (1) 5%, (2) 15%, (3) 50%, (4) 85%, and (5) 95% of nonzero pixels – pixels that were not masked out as telluric lines – with coefficient values furthest from zero, indicating strong variations in flux with changes to that label's value. These maximally varying pixels are most directly impacted by the corresponding stellar label and should thus generally correspond to relevant stellar lines or features. We retrained the model, allowing each label to contribute flux to only its selected, most highly varying pixels, then applied the trained model to our test set to check its performance.

We completed this process for two different cases: one in which we applied censoring only to the primary 4 labels describing a star ($\log g$, $T_{\rm eff}$, $v \sin i$, and [Fe/H]) and another in which we censored all 18 labels. Sample pixel selections in the 4-label case are depicted in Figure 5.4, which shows the 85% and 15% most highly-varying pixels as the top and bottom "unmasked" row of each color, respectively. Conversely, all pixels in Figure 5.4 that are not "unmasked" are labeled as "masked" below the spectrum to more clearly visualize the distribution of pixels included and excluded. With censoring implemented, each label has its own independent mask such that all labels contribute flux to only the pixels with which they vary the most.

We again completed our testing process with three train/test splits for each case to reduce the effect of stochastic variations resulting from different randomly selected train/test samples. We ultimately ran 30 total single-order train/test iterations: three iterations for each of the two censoring thresholds applied to the five choices in the number of labels censored.

Our single-order models performed best with little to no censoring, and censoring only the 4 primary labels produced consistently more accurate results than censoring all 18. However, our best-performing single-order censoring run – our 95% case with 4 labels censored –

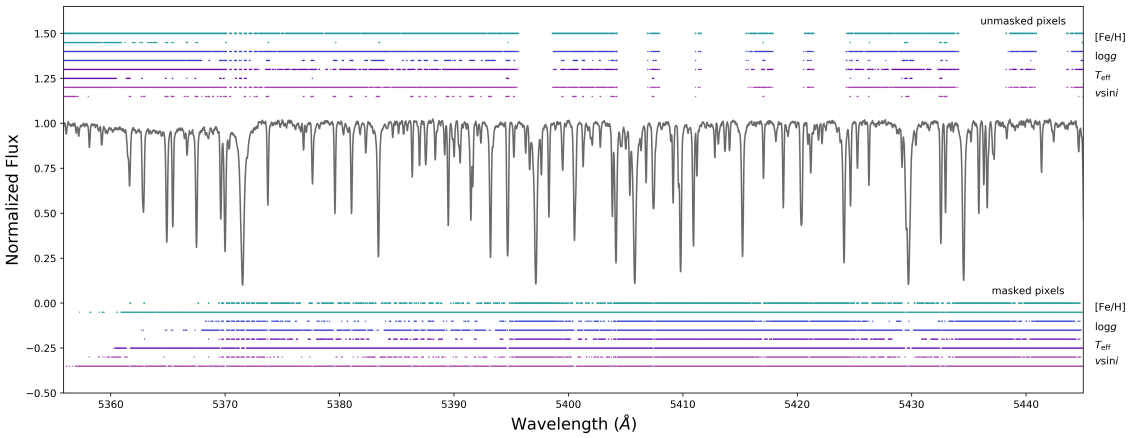

Figure 5.4: Censored wavelengths for sample star HD 22072, selected for the primary four stellar labels: [Fe/H] (green), log$g$ (blue), $T_{\rm eff}$ (violet), and $v \sin i$ (purple). The unmasked pixels corresponding to each label are shown above the spectrum, and the masked, unused pixels are below. The spectrum shown has been continuum renormalized with $N = 70$, the best-performing single-order continuum renormalization. Masks for each label are provided in pairs, where the upper line in each color corresponds to the 85% mask, while the lower line corresponds to the 15% mask. "Unmasked" pixels are included in the analysis for that label, while "masked" pixels are excluded.

performed slightly worse than our model with no censoring implemented. We tested this case with all orders included, as well, and found that, as in the single-order case, our results degraded slightly. Furthermore, censoring within *The Cannon* causes a substantial increase in the model training time. We concluded that the loss of information from removing even 5% of pixels from each label's training set was greater than the gain from censoring in our model, and we elected to use a less time-intensive version of our model with no censoring for our final model training.

5.5.3.4 L1 Regularization

Lastly, we explored the use of L1 regularization, or lasso regression, to enforce sparsity within our models. In practice, this means that we include a penalty term in our cost function which scales with the summed absolute value of coefficients for all labels. Models

are then "encouraged" to take on a simpler form in which coefficients tend towards zero values, and the severity of the penalty term determines the simplicity enforced for the model. This penalty term is denoted in Equation 5.6 as β, and the strength of the regularization is set by the parameter Λ.

$$\beta = \Lambda \sum_{q=1}^{Q-1} |\theta_q| \tag{5.6}$$

We sum over the Q components in the coefficient vector θ, excluding the zeroth term that provides the baseline spectrum of the model. Thus, the full model with regularization is given by

$$\theta_j, s_j^2 \leftarrow \underset{\theta, s}{\text{argmin}} \left[\sum_{n=0}^{N-1} \ln p(y_{jn}|\theta, l_n, s^2) + \beta \right] \tag{5.7}$$

Because our model is high-dimensional, with 18 different parameters, a sparser model may prevent over-fitting and thus lead to improved performance on the test set. We tested for this possibility by implementing regularization within our single-order model, with test values $\Lambda = 1$, 10, 100, 1,000, and 10,000. The effect of each Λ value on the distribution of coefficient strengths in one of our three trained models is shown in Figure 5.5.

We found that these enforcements of regularization substantially increased the model training time, and higher regularization values degraded the accuracy of our test set label recovery in all cases except $\Lambda = 1$, where we found slight improvement over the $\Lambda = 0$ case. However, this improvement was minor ($\chi^2 = 3.72$ by comparison with $\chi^2 = 3.74$), and applying $\Lambda = 1$ to the all-orders model increases the projected training time by a factor of over 300, making the model much less flexible and more tedious to retrain. Furthermore, Figure 5.5 shows that $\Lambda = 1$ only marginally increases the sparsity of the model, resulting in minimal changes from the $\Lambda = 0$ case. As a result, we chose not to

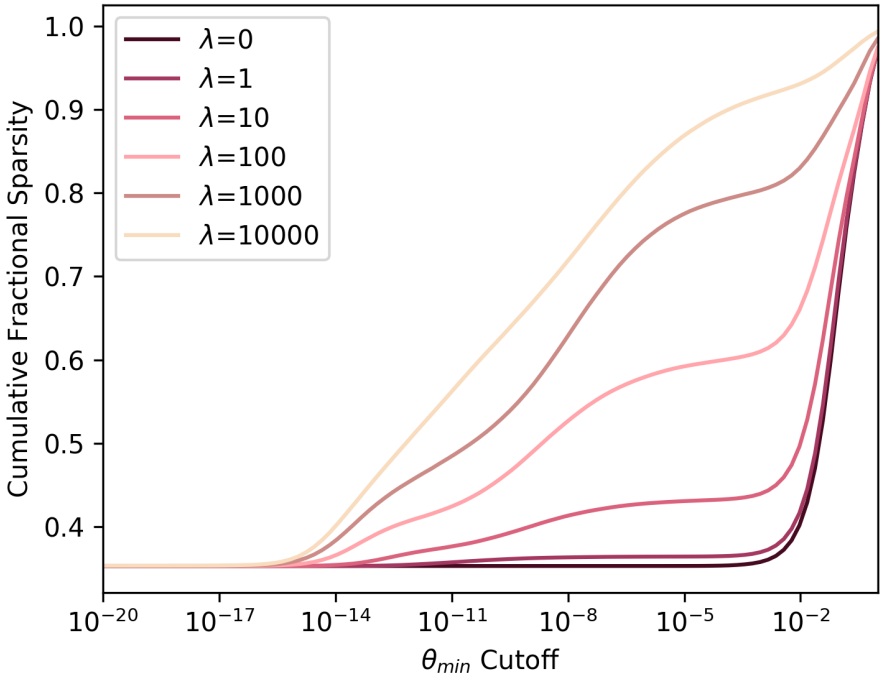

Figure 5.5: Cumulative fractional sparsity at each tested regularization value in one of our three test cases. At each $\theta_{\min}$ value along the x-axis, the total fraction of coefficients with values smaller than $\theta_{\min}$ is given for each of our test cases $\Lambda = 1, 10, 100,$ 1,000, and $10,000$, as well as the case with no regularization incorporated ($\Lambda = 0$). All cumulative distributions bottom out at fractional sparsity 0.353 because this is the fraction of pixels set to zero by our telluric mask.

	Single order	All orders
x_O	none	10
Renormalization	N=70, polynomial	N=70, polynomial
Telluric masking	included	included
Censoring	none	none
Λ	1	0

Table 5.2: Final, optimized training configuration for both our single-order model and our model incorporating all echelle orders, developed to classify post-2004 Keck HIRES spectra. The single order spans wavelength range 5355-5445 Å.

incorporate regularization in our final model configurations.

5.5.4 Final Model Configurations

Our top-performing model configurations in both the single-order and all-orders case are summarized in Table 5.2. Both models use a polynomial continuum renormalization and include telluric masking, and neither includes censoring. The primary differences between these model configurations are the lack of regularization and the removal of a single outlier data point in the all-orders configuration, using $x_O = 10$.

To understand the performance of our model, it is informative to consider how individual spectral features are reflected in the corresponding, relevant pixel coefficients. For example, Figure 5.6 shows an observed HIRES solar spectrum – obtained for calibration by observing the bright asteroid Vesta – and three coefficient vectors of our final, all-orders model ($\theta_{\rm Mg}$, $\theta_{\log g}$, $\theta_{v \sin i}$) in the vicinity of the Mg Ib triplet. Deviations from the zero baseline of each coefficient vector signify that our model finds a correlation or anticorrelation between the flux at that pixel and the parameter value weighted by that coefficient. Therefore, more weight is placed on pixels further from the baseline during the label transfer process.

The Cannon appropriately infers that the pixel at the core of each Mg Ib triplet line includes substantial information content about the Mg abundance of the star, as shown by dips

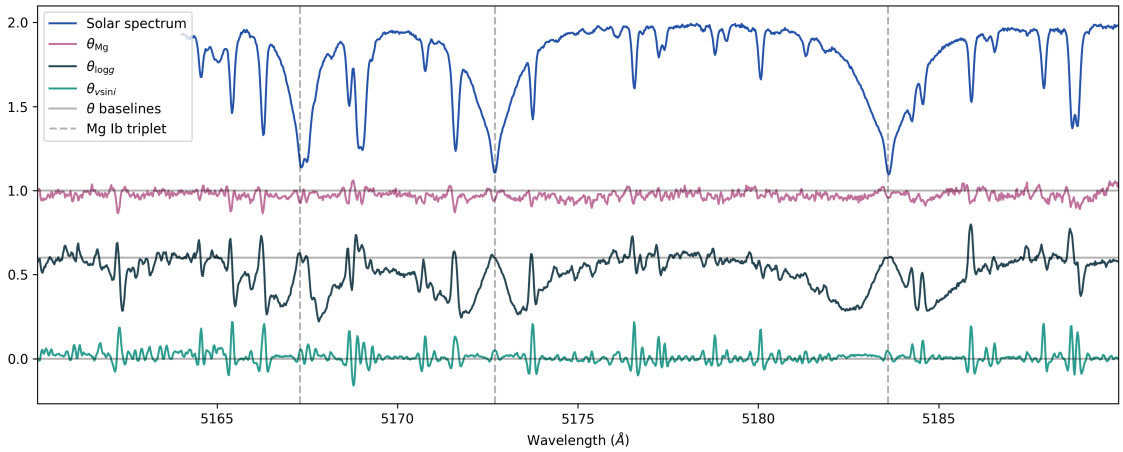

Figure 5.6: Comparison of the solar spectrum (top) with three coefficient vectors ($\theta_{\rm Mg}$, $\theta_{v\sin i}$, $\theta_{\log g}$) in the same wavelength segment for comparison. The $\theta_{\log g}$ and $\theta_{\rm Mg}$ coefficient vectors are vertically displaced to baselines of y=0.6 and y=1.0, respectively, for visual clarity; all coefficient baselines are demarcated with solid gray horizontal lines. We focus on the region directly surrounding the Mg Ib triplet to show how the primary pixel correlations deduced by *The Cannon* correspond to spectral features at the same wavelengths. While $\theta_{\rm Mg}$ directly correlates with the cores of the three Mg lines, which provide information about the Mg abundance, $\theta_{\log g}$ is more directly affected by the wings of the lines, which provide a metric for the star's surface gravity strength. Intermediate-depth lines are most heavily weighted by $\theta_{v\sin i}$, since these lines are typically neither saturated nor prone to blending together with the baseline. These features are independently identified by *The Cannon* through its training process, demonstrating that the correlations identified by the model correspond directly to known physical features.

in $\theta_{\rm Mg}$ at these pixels. Conversely, $\theta_{\log g}$ approaches zero at each of these line cores and deviates further from zero at the wings of each line, reflecting the physical phenomenon of line broadening with increased surface gravity. Information about the stellar spin can be gleaned from all spectral lines, as reflected in $\theta_{v \sin i}$; however, our model most heavily weights intermediate-depth lines. This is likely because deeper lines may be saturated, while shallower lines blend together and are washed out by noise, making intermediate-depth lines the most informative for determining $v \sin i$. As a whole, Figure 5.6 reflects that the most prominent correlations identified by *The Cannon* correspond directly to known physical features, improving confidence in our model's results.

We depict the test results from our final model configuration in Figure 5.7. The dotted, diagonal line in each panel corresponds to a perfect recovery of the expected label with *The Cannon*, while deviations from this line indicate scatter in the results. The scatter in results provides a per-label representative 1σ uncertainty estimate for our model. Our best-performing all-orders model returns average $\chi^2 = 5.89$ across our three test/train splits.

To best visualize the overall model performance, we removed three outliers from the [Al/H] panel, each of which lies far to the left of the panel shown in Figure 5.7. The outliers are all at the bottom end of the distribution of predicted [Al/H] values, and the poor estimates returned for these three stars are likely a result of the sparsity of comparable stars in the SPOCS catalog with low [Al/H]. This behavior is expected, since by design *The Cannon* performs best on spectra similar to the training set and is not expected to extrapolate beyond what the model has been taught. Estimated labels outside of the range of our training set are unreliable and must be treated with caution.

The parameter space spanned by SPOCS is provided for reference in Table 5.3, where we report only the reliable range of [Al/H] values ([Al/H] ≥ -0.66) without outliers. We also include in these "reliable ranges" only the range of stellar parameters reliably recovered by

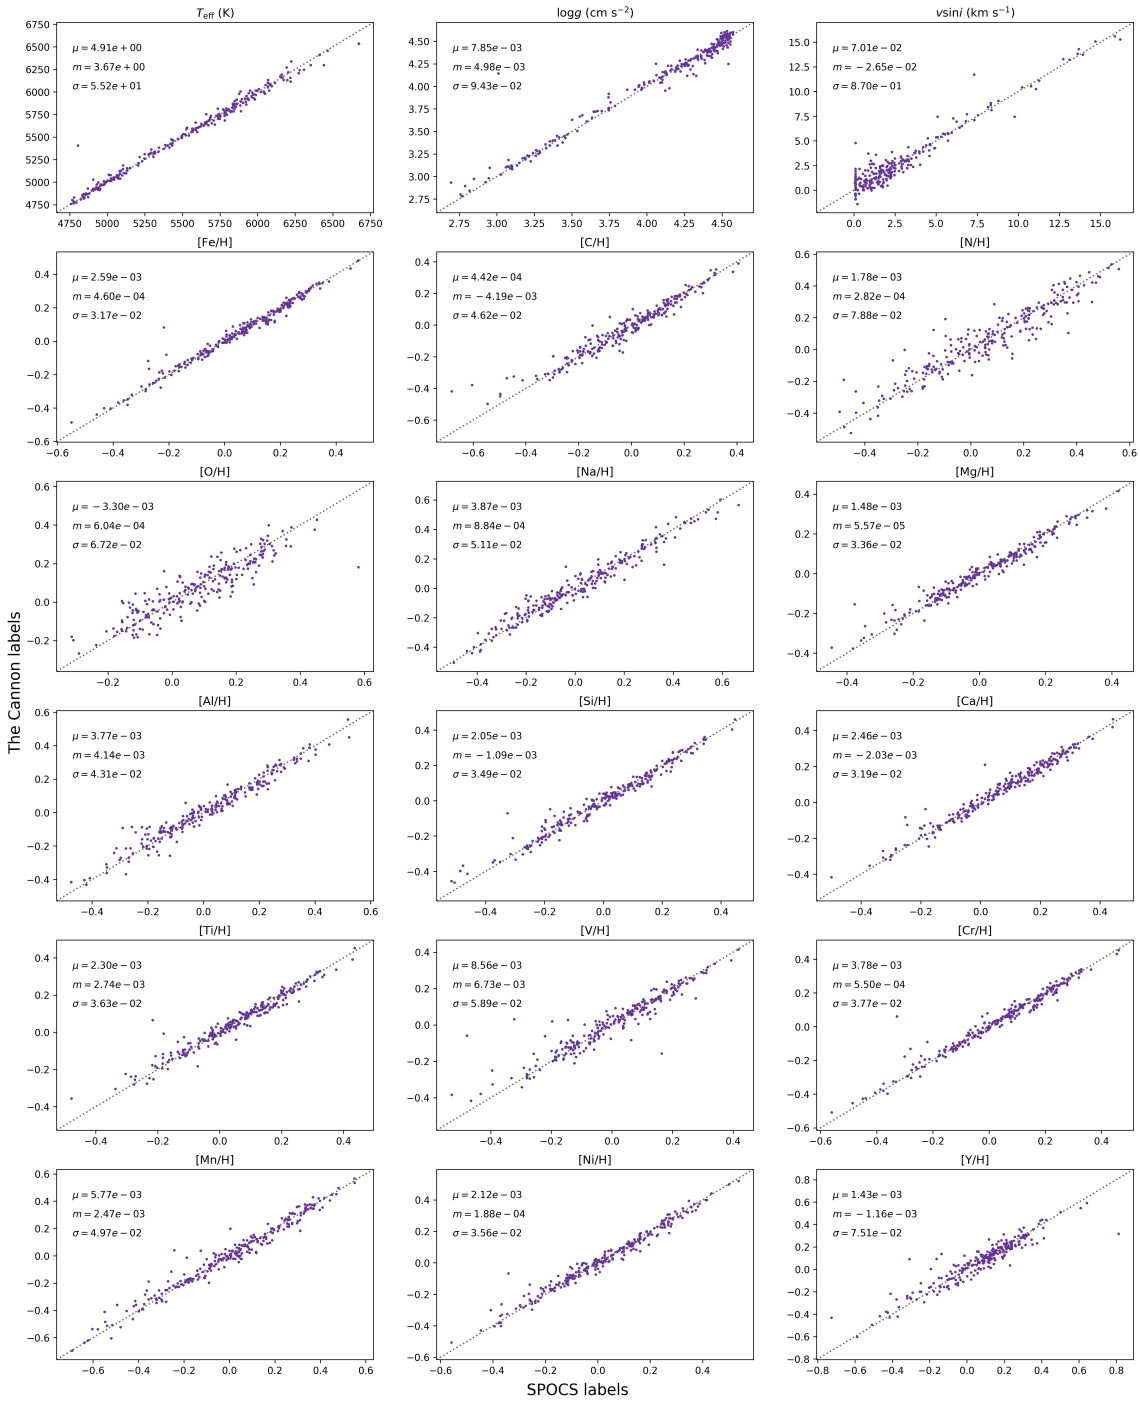

Figure 5.7: Post-2004 test results for all parameters with all echelle orders incorporated. In each panel, the mean μ, median m, and standard deviation σ from a perfect guess ($|x_i - E_i| = 0$) are provided in the top left. To most clearly visualize the bulk of our results, we exclude three outlier data points with SPOCS labels [Al/H] < -0.66 from the [Al/H] panel. We also do not include these [Al/H] outliers in our reported "reliable range" (see Table 5.2) or in the calculation of parameters in the top left.

our pipeline, meaning that these reported ranges are, in some cases, slightly smaller than the full range spanned by our training set. We note that $v \sin i$ has a cutoff at zero in the SPOCS database, resulting in the observed pileup at low $v \sin i$ in the top right panel of Figure 5.7. We do not impose an analogous condition with *The Cannon*.

We find that *The Cannon* reliably returns the expected stellar labels, with scatter, listed in Table 5.3, typically lower than but comparable to the scatter between different catalogs providing spectroscopic parameter estimates. For example, the Hypatia catalog finds roughly $0.1 - 0.2$ dex scatter between catalogs for each elemental abundance (Hinkel et al., 2014). The uncertainty in each stellar label returned by *The Cannon* is typically a factor of a few higher than that of the input SPOCS uncertainties (see Table 6 in Brewer et al., 2016). This makes sense, since our results cannot be more precise than the labels on which they are trained.

Our full model also returns lower scatter in $T_{\rm eff}$ and [Fe/H] than that obtained by `SpecMatch`, which reaches accuracies of 70 K in $T_{\rm eff}$ and 0.12 dex in [Fe/H] for stellar types K4 ($T_{\rm eff} \sim 4600$) and later (Yee et al., 2017). Furthermore, the scatter in temperature that we find is approximately equivalent to that of the combined infrared flux temperature measurements across different analyses (González Hernández & Bonifacio, 2009; Casagrande et al., 2010; Brewer et al., 2016).

This model is somewhat computationally expensive to train ($\sim$80 minutes for our best-performing all-orders configuration); however, once trained, it takes just a few ($\sim$3) seconds per spectrum to extract the 18 parameters of interest. This makes it a particularly powerful tool for large samples of stars, since the up-front model training procedure only needs to be performed once. We make our final model, which employs the top-performing all-orders configuration trained on all 1201 vetted stars in our $x_O = 10$ sample, publicly available at `https://github.com/malenarice/keckspec`. All other files required to run the code are also provided.

Label	$\sigma_{\rm 1order}$	$\sigma_{\rm full}$	Reliable Range
$T_{\rm eff}$ (K)	77	56	$4700 - 6674$
log$g$ (cm s^{-2})	0.13	0.09	$2.70 - 4.83$
$v \sin i$ (km s^{-1})	0.85	0.87	$0.0044 - 18.71$
[C/H]	0.09	0.05	$-0.60 - 0.64$
[N/H]	0.08	0.08	$-0.86 - 0.84$
[O/H]	0.07	0.07	$-0.36 - 0.77$
[Na/H]	0.09	0.05	$-1.09 - 0.78$
[Mg/H]	0.06	0.04	$-0.70 - 0.54$
[Al/H]	0.05	0.04	$-0.66 - 0.58$
[Si/H]	0.07	0.03	$-0.65 - 0.57$
[Ca/H]	0.04	0.03	$-0.73 - 0.54$
[Ti/H]	0.05	0.04	$-0.71 - 0.52$
[V/H]	0.07	0.06	$-0.85 - 0.46$
[Cr/H]	0.06	0.04	$-1.07 - 0.52$
[Mn/H]	0.10	0.05	$-1.40 - 0.66$
[Fe/H]	0.05	0.03	$-0.99 - 0.57$
[Ni/H]	0.07	0.04	$-0.97 - 0.63$
[Y/H]	0.07	0.08	$-0.87 - 1.35$

Table 5.3: Summary of our post-2004 Keck HIRES model performance for each parameter, including 1σ scatter in each label and the parameter space spanned by our training dataset over which results are considered reliable. Scatter in our best-performing single-order and all-orders results is reported as $\sigma_{\rm 1order}$ and $\sigma_{\rm full}$, respectively.

5.6 Application to Pre-2004 Spectra

With our optimized model for current Keck spectra at hand, we then developed a framework to classify archival, pre-2004 data. Our archival dataset includes 831 Keck spectra, each continuum-normalized in the same manner as the SPOCS dataset, obtained from 810 different stars prior to Keck's detector upgrade. We set aside the 8 stars in our sample with multiple spectra available for a separate analysis of the scatter in results obtained from *The Cannon*, described in Section 5.6.2.

There are 337 single-spectrum stars in our pre-2004 archival dataset that were also observed after 2004 and are accordingly included in the SPOCS dataset. We use this overlapping sample as our test set to check and optimize model performance; we train on post-2004 spectra of the 865 stars that were not observed prior to 2004, then test our results using the pre-2004 spectra of our 337 test set stars. By construction, therefore, we no longer randomly sample our train/test sets from the same larger pool of spectra. As a result, we completed only one iteration for each test case in this section. Ultimately, we applied our optimized model to report newly obtained stellar labels for the remaining 473 single-spectrum stars, as well as 4 multi-spectrum stars that were not characterized in Brewer et al. (2016).

All spectra needed to be recalibrated prior to training and testing with *The Cannon* due to structural differences between the pre- and post-2004 Keck HIRES detectors. Each echelle order of the pre-2004 spectra includes either 2047 or 2048 pixels, rather than the 4021 pixels per echelle order sampled in our post-2004 training set. In addition, the pre-2004 echelle orders do not span exactly the same wavelength ranges as the post-2004 echelle orders.

For each echelle order, we found the broadest wavelength range covered by all available spectra by directly comparing pre- and post-2004 HIRES echelle orders with the maximum

wavelength overlap. We then interpolated our training set, as well as all archival spectra, onto a 2048-pixel scale spanning this wavelength range for a uniform comparison across samples. From this process, we obtained 12 overlapping echelle orders covered by both the pre- and post-2004 HIRES spectra, for a total of 24,576 pixels modeled for each star.

5.6.1 Extrapolating the Model to Pre-2004 Spectra

Our goal in this section is to find a best-performing model that incorporates all 12 echelle orders in order to extract new labels for the 477 unlabeled stars. Due to the differing systematics across spectrographs, it is not necessarily true that the same optimization found in Section 5.5 will provide the best results in our interpolated model. We therefore repeated the model tuning steps described in Section 5.5.3 to find an optimal configuration for our interpolated model. We refer the interested reader to Appendix A for a detailed discussion of this process, which closely parallels that described in Section 5.5.

Our best-performing all-orders model built in this way – by progressively accepting or rejecting each potential alteration one by one – is characterized by outlier removal with $x_O = 3$ (10 stars removed; 7 from the training set and 3 from the test set), a sin/cos continuum renormalization with $N = 70$, and no telluric masking, censoring, or regularization. However, when we compared this model with our post-2004 optimization (see Table 5.2), we found that we obtained the best results when applying the post-2004 configuration to the interpolated spectra. Our final model is therefore trained with the same all-orders optimized hyperparameters described in Table 5.2. With $x_O = 10$, the performance of this model was verified using 337 test set stars and 864 training set stars. Our best-performing all-orders model returns $\chi^2 = 4.37$.

We visually display the performance of our model in Figure 5.8 and report our final 1σ uncertainties and reliable ranges in Table 5.4. Because it is trained using the same sample of stars, our pre-2004 model spans the same parameter space as our post-2004 model. The

Label	σ_{full}	Reliable Range
T_{eff} (K)	42	$4700 - 6674$
log$g$ (cm s^{-2})	0.05	$2.70 - 4.83$
$v \sin i$ (km s^{-1})	0.98	$0.0044 - 18.71$
[C/H]	0.05	$-0.60 - 0.64$
[N/H]	0.09	$-0.86 - 0.84$
[O/H]	0.07	$-0.36 - 0.77$
[Na/H]	0.06	$-1.09 - 0.78$
[Mg/H]	0.03	$-0.70 - 0.54$
[Al/H]	0.05	$-0.66 - 0.58$
[Si/H]	0.03	$-0.65 - 0.57$
[Ca/H]	0.03	$-0.73 - 0.54$
[Ti/H]	0.03	$-0.71 - 0.52$
[V/H]	0.04	$-0.85 - 0.46$
[Cr/H]	0.03	$-1.07 - 0.52$
[Mn/H]	0.05	$-1.40 - 0.66$
[Fe/H]	0.02	$-0.99 - 0.57$
[Ni/H]	0.03	$-0.97 - 0.63$
[Y/H]	0.07	$-0.87 - 1.35$

Table 5.4: Summary of our pre-2004 Keck HIRES model performance for each parameter, including 1σ scatter in each label (σ_{full}), as well as the parameter space spanned by this model. Because our sample of stars is the same as in the post-2004 model, the reliable range remains unchanged.

model performs remarkably well given that it has been trained using data from a different spectrograph from the test set. We demonstrate that, even with an interpolated set of spectra taken using a different instrument, our model returns the expected labels with high fidelity.

After optimizing the model hyperparameters, we re-trained our final model on the full $x_O = 10$ SPOCS dataset with a total of 1201 stars. We then applied this model to our set of 473 unlabeled single-spectrum stars to obtain all 18 labels for each star in the pre-2004 dataset, reported in Table 5.5. We searched these returned labels for outliers that do not fall within the parameter ranges of our training set and that therefore may be unreliable. In total, 128 of our 473 stars had at least one predicted parameter that fell outside of these

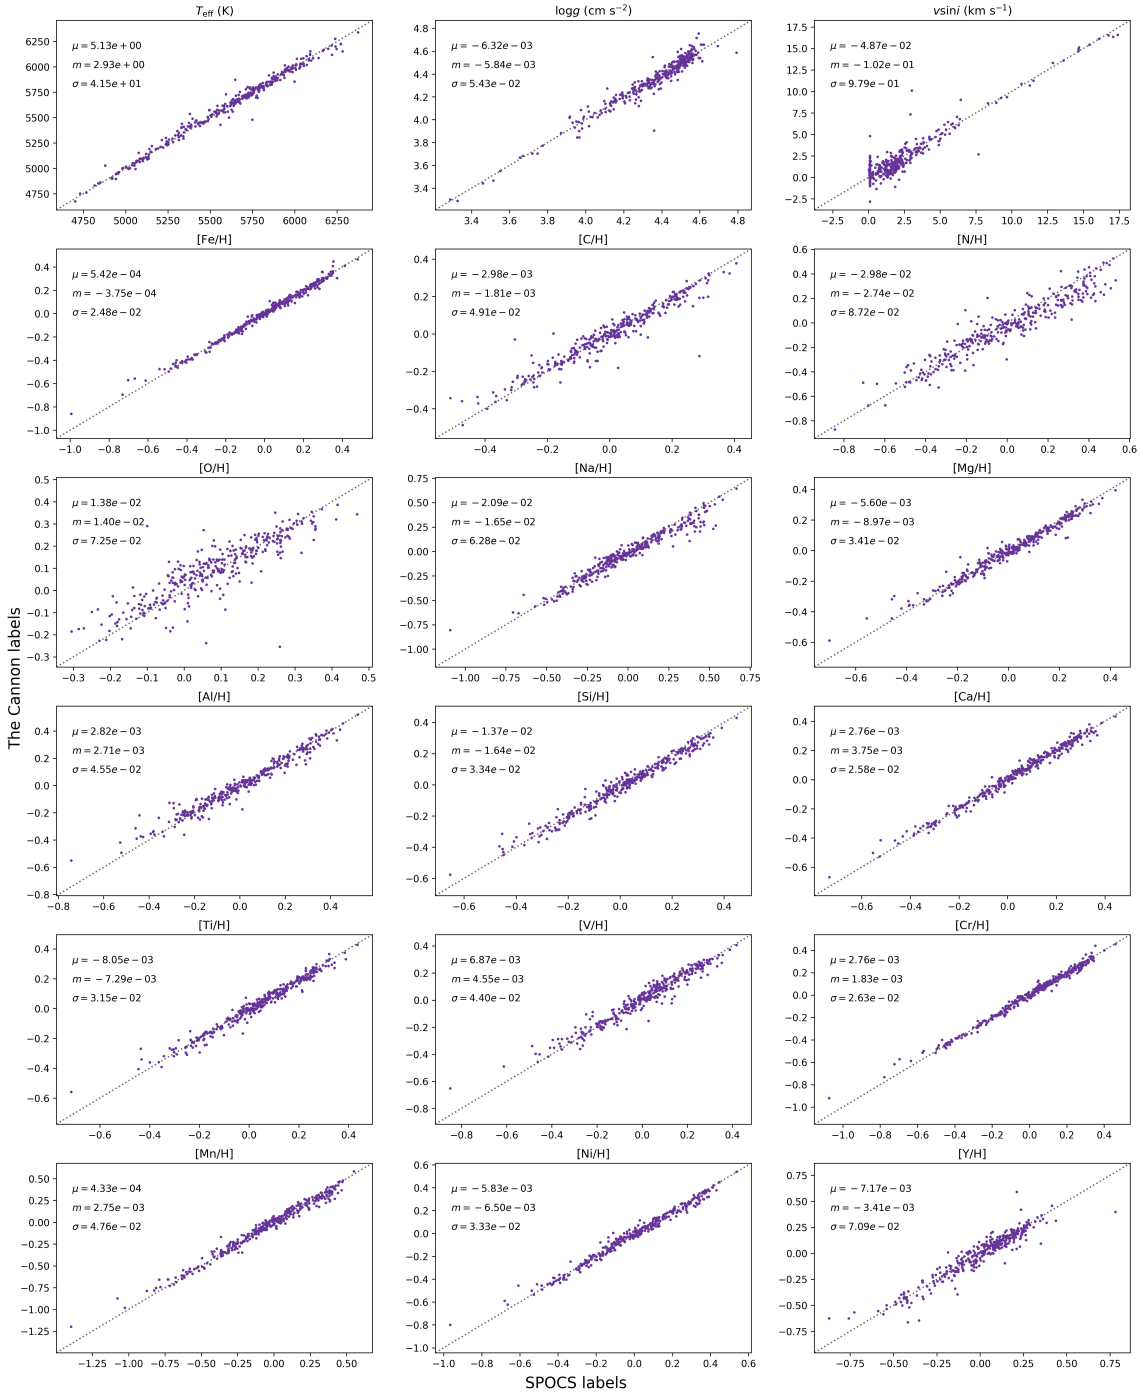

Figure 5.8: Pre-2004 test results for all parameters with all 12 echelle orders incorporated. In each panel, the mean μ, median m, and standard deviation σ from a perfect guess ($|x_i - E_i| = 0$) are provided in the top left. An outlier below our reported reliable range has been left out of the [N/H] panel for visual clarity.

training set ranges. We flag these stars in the rightmost column of Table 5.5, where "y" indicates that a star has at least one predicted parameter outside of the reliable range.

Star	$T_{\rm eff}$	log$g$	$v \sin i$	[C/H]	[N/H]	[O/H]	[Na/H]	[Mg/H]	[Al/H]	[Si/H]	[Ca/H]	[Ti/H]	[V/H]	[Cr/H]	[Mn/H]	[Fe/H]	[Ni/H]	[Y/H]	Flagged?
HD 98744	6157	3.98	2.15	-0.24	-0.09	0.01	-0.16	-0.19	-0.35	-0.19	-0.10	-0.10	-0.16	-0.15	-0.39	-0.16	-0.24	-0.12	-
HD 18144	5514	4.47	0.49	0.05	-0.02	0.11	0.04	0.05	0.07	0.04	0.06	0.05	0.07	0.07	0.07	0.07	0.04	0.05	-
HD 91204	5935	4.32	2.17	0.21	0.18	0.26	0.27	0.21	0.27	0.21	0.23	0.22	0.19	0.25	0.28	0.24	0.25	0.19	-
HD 230409	5388	4.67	-0.19	-0.49	-0.82	-0.14	-0.78	-0.58	-0.57	-0.55	-0.64	-0.54	-0.62	-0.86	-1.13	-0.81	-0.77	-0.66	y
HD 150437	5742	4.21	2.67	0.20	0.45	0.21	0.49	0.24	0.27	0.26	0.25	0.23	0.20	0.26	0.37	0.27	0.35	0.27	-
HD 25825	6020	4.43	6.93	0.11	0.15	0.13	0.04	0.11	0.06	0.13	0.18	0.16	0.17	0.18	0.12	0.19	0.12	0.25	-
HD 43587	6194	4.63	14.22	0.25	-0.17	0.35	0.05	0.22	0.04	0.17	0.37	0.31	0.32	0.34	0.27	0.31	0.19	0.30	-
HD 213575	5710	4.28	0.58	-0.05	-0.13	0.09	-0.12	-0.03	0.04	-0.07	-0.08	0.01	-0.04	-0.16	-0.27	-0.14	-0.13	-0.15	-
HD 139324	5889	4.23	1.32	0.10	0.07	0.19	0.13	0.10	0.10	0.10	0.13	0.10	0.13	0.13	0.17	0.14	0.14	0.11	-
HD 188510	5926	4.62	-2.47	-0.54	-0.54	-0.41	-0.82	-0.64	-0.73	-0.65	-0.71	-0.60	-0.67	-0.85	-1.11	-0.77	-0.80	-0.84	y
HD 157172	5481	4.49	2.46	0.10	0.23	0.10	0.23	0.13	0.10	0.13	0.12	0.12	0.13	0.14	0.20	0.14	0.16	0.08	-
HD 98618	5861	4.44	0.65	0.04	0.01	0.06	0.04	0.06	0.08	0.04	0.06	0.05	0.08	0.04	0.04	0.05	0.05	0.03	-
HD 7727	6080	4.34	2.84	0.06	0.11	0.16	0.06	0.07	0.04	0.08	0.13	0.09	0.09	0.12	0.08	0.12	0.07	0.10	-
HD 202108	5732	4.56	0.53	-0.18	-0.30	-0.10	-0.27	-0.19	-0.22	-0.19	-0.16	-0.17	-0.16	-0.19	-0.32	-0.20	-0.27	-0.17	-
HD 13825	5711	4.38	0.46	0.15	0.24	0.14	0.30	0.18	0.23	0.18	0.19	0.17	0.19	0.18	0.25	0.17	0.20	0.10	-
HD 152792	5738	4.05	0.56	-0.25	-0.30	-0.14	-0.30	-0.24	-0.27	-0.28	-0.20	-0.20	-0.25	-0.30	-0.47	-0.27	-0.31	-0.20	-
HD 204587	4556	4.46	1.31	0.31	-0.29	0.22	0.15	-0.06	0.01	0.10	0.10	-0.03	-0.07	-0.14	-0.08	-0.02	0.00	-0.22	y
HD 139457	6084	4.10	0.64	-0.28	-0.25	-0.11	-0.37	-0.26	-0.38	-0.29	-0.23	-0.18	-0.25	-0.33	-0.64	-0.32	-0.39	-0.36	-
HD 104067	4884	4.49	1.39	0.14	-0.09	0.12	0.08	0.06	0.09	0.10	0.12	0.09	0.10	0.06	0.08	0.10	0.07	-0.04	-
HD 106116	5657	4.35	-0.36	0.08	0.09	0.12	0.12	0.11	0.15	0.09	0.12	0.10	0.12	0.12	0.16	0.13	0.12	0.11	y
HD 208801	4834	3.48	1.35	0.10	0.20	0.20	0.09	0.11	0.25	0.01	0.11	0.18	0.10	0.09	0.19	0.12	0.14	-0.02	-
HD 120066	5830	4.14	1.20	0.01	0.04	0.13	0.02	0.10	0.13	0.07	0.13	0.12	0.10	0.08	0.04	0.10	0.09	0.15	-
HD 88218	5849	4.10	0.33	-0.10	-0.14	0.03	-0.14	-0.10	-0.08	-0.13	-0.07	-0.07	-0.10	-0.12	-0.23	-0.11	-0.13	-0.07	-
HD 141103	6213	4.11	3.83	-0.23	-0.04	-0.02	-0.18	-0.16	-0.34	-0.17	-0.12	-0.08	-0.11	-0.17	-0.42	-0.16	-0.25	-0.22	-
HD 213628	5580	4.50	0.46	0.00	-0.06	0.01	-0.05	0.02	0.05	-0.00	0.01	0.02	0.05	0.00	-0.01	0.01	0.00	-0.05	-
HD 122303	4972	4.17	-2.87	0.10	-0.22	-0.27	-0.23	-0.05	0.56	-0.22	-0.09	0.05	-0.13	-0.08	0.03	-0.09	0.00	0.14	y
HD 181655	5674	4.44	3.60	0.01	-0.11	0.06	-0.04	0.01	0.04	0.02	0.08	0.03	0.08	0.06	0.01	0.06	-0.00	0.11	-
HD 47127	5611	4.34	0.38	0.07	0.05	0.14	0.07	0.09	0.14	0.07	0.10	0.10	0.10	0.09	0.09	0.09	0.08	0.06	-
HD 144253	4790	4.25	8.48	-0.13	-0.24	-0.27	0.03	-0.16	0.03	0.07	0.07	-0.06	-0.03	-0.08	0.01	-0.01	-0.03	-0.30	-
HD 90125	4816	2.99	0.60	-0.25	-0.14	-0.03	-0.29	-0.16	-0.08	-0.29	-0.11	-0.07	-0.17	-0.17	-0.12	-0.10	-0.12	0.02	-
HD 150554	6080	4.42	2.58	0.07	0.04	0.06	0.05	0.04	-0.04	0.05	0.05	0.03	0.03	0.08	0.04	0.08	0.04	0.14	-
HD 9331	5580	4.27	1.00	0.13	0.08	0.22	0.14	0.20	0.22	0.18	0.20	0.17	0.21	0.18	0.22	0.19	0.20	0.19	-
HD 183650	5654	4.13	1.10	0.23	0.36	0.21	0.45	0.23	0.33	0.25	0.28	0.26	0.23	0.25	0.34	0.26	0.31	0.19	-
HIP 84099	5161	4.30	-1.79	0.10	-0.56	-0.30	-0.39	-0.11	0.22	-0.22	-0.18	-0.01	-0.33	-0.16	-0.17	-0.18	-0.15	-0.17	y

Table 5.5: Stellar labels returned by our trained model for archival, pre-2004 Keck HIRES spectra.

NOTE—See Rice & Brewer (2020) for a downloadable version of the full table.

Star Name	# of spectra
HD 178911B	4
HD 141937	2
HD 11964A	3
HD 212291	2
HIP 76901	3
HD 207740	3
HD 92222A	2
HD 92222B	2

Table 5.6: Stars in our pre-2004 dataset with multiple archival spectra available. The first four stars in the table are also included in the labeled SPOCS dataset, while the last four are not.

Figure 5.9 illustrates the distribution of stellar parameters returned for our full set of 473 single-spectrum stars that were not included in the SPOCS dataset. Regions outside of the reliable parameter space (see Table 5.4) are shaded in light red. The distribution of all predicted stellar labels is in light green, and the corresponding distribution of only stars for which all labels fall in the "reliable" parameter space is overlaid in dark green.

5.6.2 Scatter in Results: Stars with Multiple Spectra

Eight stars in our archival dataset have more than one spectrum available, and we use a subset of these for a separate check on the precision of parameters reported for our model.

All stars with multiple archival spectra available are listed in Table 5.6, along with the total number of spectra available for each star. The first four stars in Table 5.6 are included in our training set, while the last four are not. Therefore, only the last four stars (HIP 76901, HD 207740, HD 92222A, and HD 92222B) are included in Table 5.5, where we report results for all spectra of each star for thoroughness. We use the first four stars, for which we have SPOCS labels, to study the performance of our model.

In Figure 5.10, we show the spread in results for HD 178911B, HD 141937, HD 11964A,

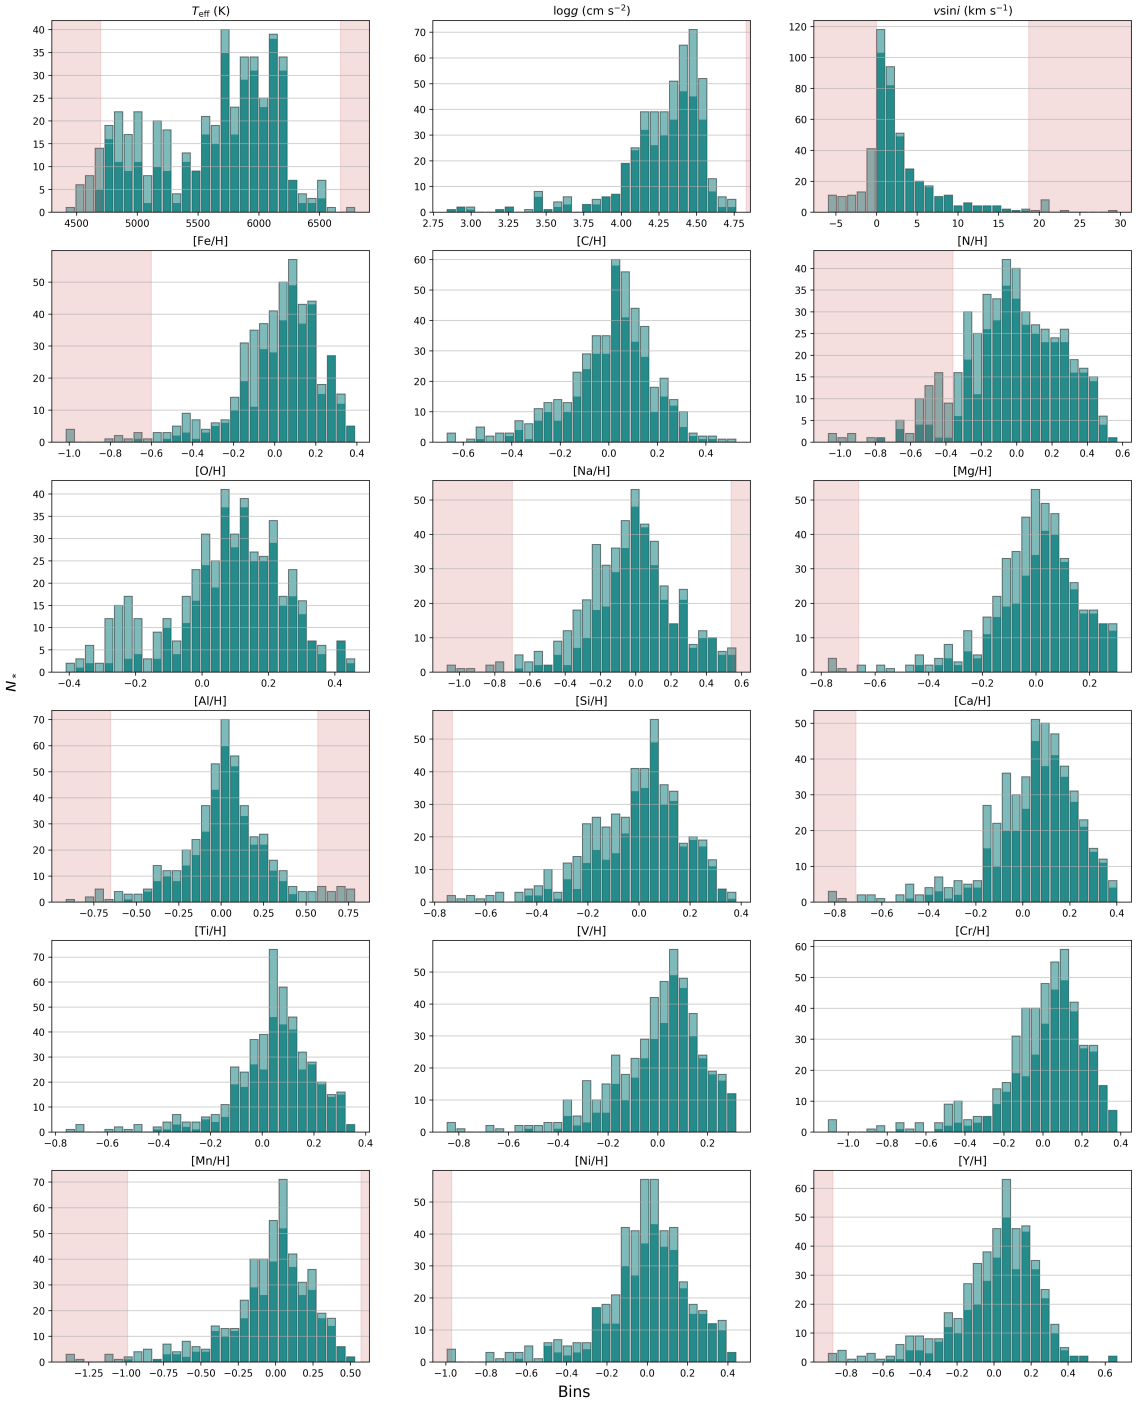

Figure 5.9: Distribution of labels attributed by *The Cannon* to the 473 single-spectrum, previously unlabeled archival Keck stars. Regions of parameter space outside the reliable range detailed in Table 5.4 are shaded in light red. The full distribution of all stars, including those with at least one unreliable parameter estimate, is provided for each label in light green. The dark green overlaid distribution includes only stars with all parameters falling within the reliable range.

and HD 212291, our sample stars with multiple archival spectra available and with known SPOCS labels, in order to display the precision of our model (how consistent its predictions are with each other) as well as its accuracy (how these predictions compare with the corresponding SPOCS values). Each star is represented by a color given in the legend, and estimates measured from different spectra of the same star are connected with a line. Values are reported relative to the "correct" SPOCS labels. Thus, points that lie to the right of the zero line (shown as a vertical dashed black line) are overestimated relative to the SPOCS labels, while points to the left of the zero line are underestimated. All points within the shaded gray regions fall within 1σ of the corresponding SPOCS value, and stars are vertically separated for visual clarity.

We find that most, but not all, of our labels from *The Cannon* fall within 1σ of the "correct" SPOCS label. For a more conservative uncertainty estimate, therefore, these 1σ uncertainties may be multiplied by a factor of 1.5 to 2. The typical scatter in spectral properties of an individual star is fairly small, and labels returned by different observations of the same star are generally consistent with each other within our error bars. Stars may also have some intrinsic variability such that labels will not necessarily stay exactly the same over time.

5.7 Potential Biases and Systematics

After quantifying the overall performance of *The Cannon* with our pre-labeled SPOCS dataset, we also searched for trends in the model results that could be indicative of systematic biases in the labels returned by SME. We report these trends and their potential origins, and, where possible, we propose methods to eliminate these trends in future analyses. We chose to complete this process using an early version of our model, with no tuning implemented, to ensure that any observed trends result from our input labels, rather than

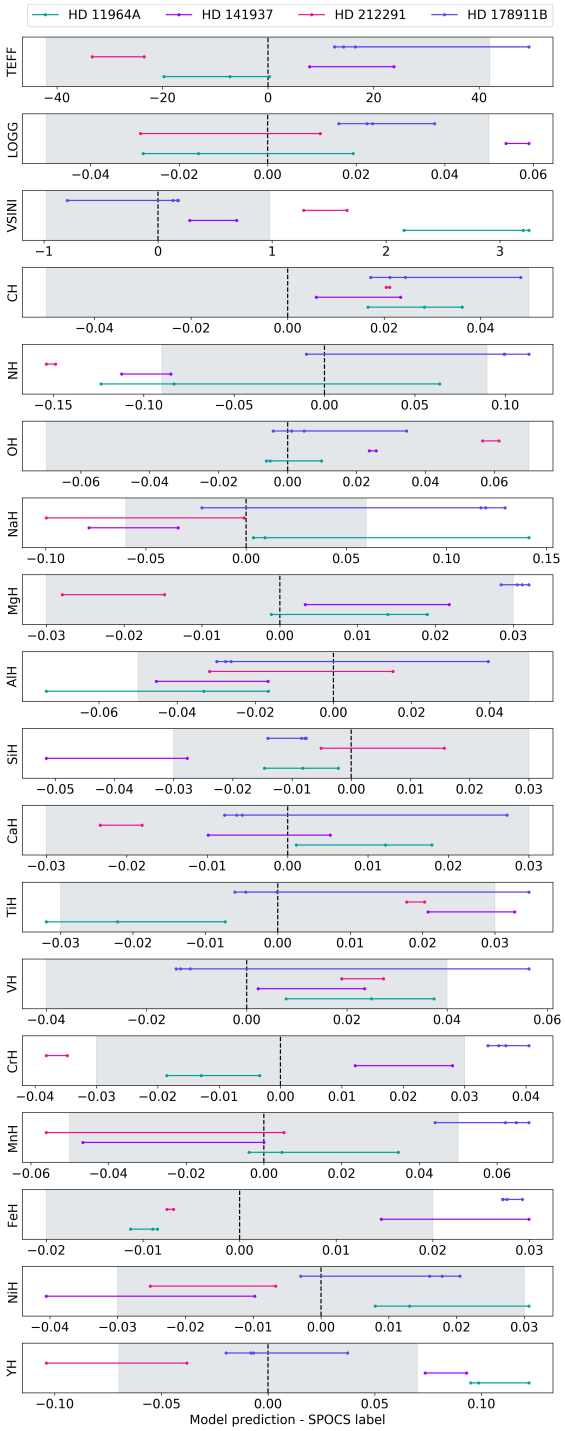

Figure 5.10: Multi-spectrum stars in our archival sample with available SPOCS labels. An exact match to the SPOCS label is represented by the vertical black dashed line, and the gray region represents the 1σ parameter space for each label. Each predicted label from *The Cannon* is represented by a point, and we connect points that are associated with the same star for visual clarity.

adjustments in our model setup. Accordingly, all trends described in this section are based on results from our $x_O = 10$ all-orders post-2004 model with no additional tuning.

Throughout this section, we show that *The Cannon* can be used to draw out systematic patterns across an input catalog. This demonstrates that *The Cannon* may, more broadly, serve as a useful tool to search and correct for systematic trends across a given stellar catalog.

5.7.1 Metallicity Correlations

While examining the offset between our test results and corresponding SPOCS labels, we observed a clear gradient with metallicity in most abundance estimates, as shown in Figure 5.11. Figure 5.11 displays [Mg/H] as a representative example, demonstrating that elemental abundances deduced by *The Cannon* tend to be overestimated at low metallicity and underestimated at high metallicity relative to the corresponding SPOCS abundances. As shown in Figure 5.11, the model tends to perform well at solar metallicity; however, our results deviate further from those in the SPOCS catalog for stars with metallicity further from solar. This match at solar metallicity is likely a result of pre-processing in Brewer et al. (2016) that calibrated the VALD-3 line list relative to solar using the National Solar Observatory solar flux atlas (Wallace et al., 2011). We include a linear regression of all points in Figure 5.11, as well as the baseline at zero corresponding to a "perfect" parameter recovery, for reference.

This systematic offset, consistent across abundances, most likely results from a systematic bias in the input dataset, since a generative model returns parameters analogous to those that it accepts as an input – including learning any biases inherent in that input model. In particular, the anticorrelation between [Fe/H] and all other elemental abundances suggests a bias in the value of [M/H].

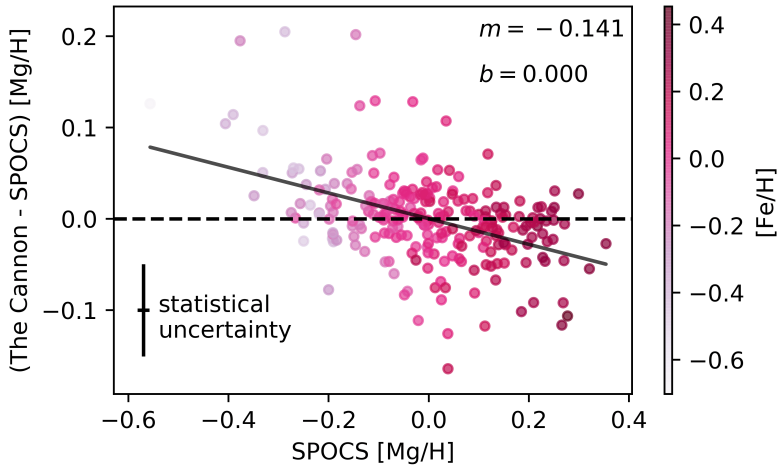

Figure 5.11: [Mg/H] from the SPOCS catalog as a function of the difference between the [Mg/H] labels predicted by *The Cannon* and those of the input SPOCS sample. A linear fit to the data is shown in black, with slope m and y-intercept b, to quantify this downward trend. The statistical uncertainty of each point, obtained for the full population based on scatter in test results with *The Cannon* and the SPOCS label uncertainties reported in Brewer et al. (2016), is provided in the bottom left.

In theory, [M/H] represents the summed abundance of all metals as compared to the solar value and should, therefore, trace [Fe/H] with a slight offset to account for all other metals. In SME, [M/H] is estimated from [Fe/H] and used to choose a model atmosphere grid to build a spectral model. Abundances are then determined by modeling lines using radiative transfer through that atmosphere with the current model parameters. If the parameters do not change substantially, the same atmosphere may be used to sample a range of possible abundances. As a result, the value of [M/H] is not forced to be in agreement with the value of [Fe/H] in our input dataset.

To further explore this bias, we studied the correlation of [M/H] with [Fe/H] in our input dataset. In Figure 5.12, we show [Fe/H] as a function of the offset between [M/H] and [Fe/H]. There is a clear downward slope in this offset value, best fit by a line with slope $m = -0.086$ and y-intercept $b = -0.013$. The coolest stars in the sample follow a steeper downwards slope, while the hotter stars, which dominate the sample, follow a shallower slope.

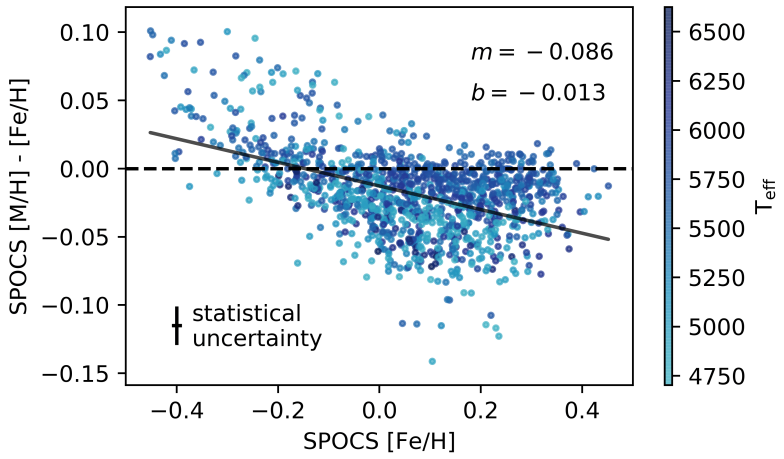

Figure 5.12: SPOCS dataset correlation with [M/H] with [Fe/H]. A linear fit to the data is shown in black, with slope m and y-intercept b, to quantify the downward trend. The statistical uncertainty of each point is provided in the bottom left.

Interestingly, the best-fitting linear interpolation that we found does not produce a perfect fit at solar metallicity. We conclude that the observed offset is likely due to degeneracies between $T_{\rm eff}$ and log$g$ within the model atmosphere selection grid in SME. Although [M/H] is not a parameter in our model with *The Cannon*, it tends towards solar values in our results based on the anticorrelation observed in Figure 5.11. The offset in [Mg/H] calculated from *The Cannon* and SPOCS roughly tracks the trend of [M/H] with [Fe/H] from the SPOCS dataset, showing that *The Cannon* reproduces this inherited pattern.

The systematic trend in [M/H] with [Fe/H] is a problem inherent to our input dataset which, in turn, produces a bias in our results from *The Cannon* as shown in Figure 5.11. In the process of determining labels with SME, a stellar atmosphere model is selected from a coarse grid with steps of 250 K in temperature, 0.5 dex in logg, and 0.1 dex in [M/H] for values -0.3 to +0.3, or 0.5 dex outside. The presence of this systematic problem indicates that a reanalysis of these spectra in SME with a finer atmospheric grid – or a different atmospheric grid altogether – may be warranted.

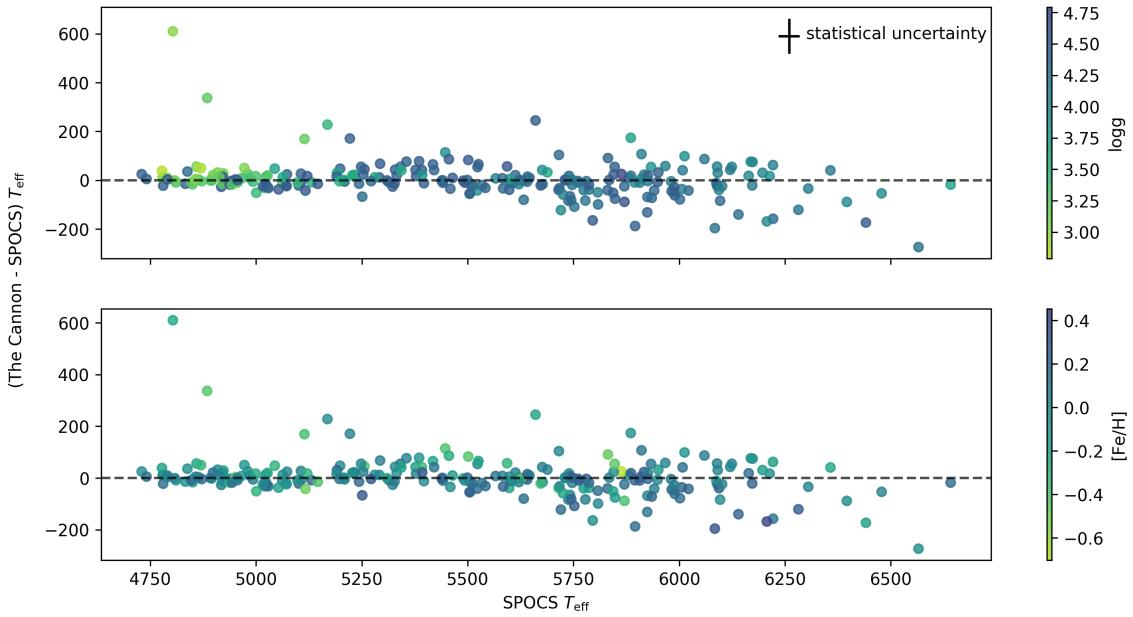

Figure 5.13: Effective temperature as a function of the difference between the labels predicted by *The Cannon* and those of the input SPOCS sample. The top panel is colored by log$g$, while the bottom panel is colored by [Fe/H]. The statistical uncertainty of each point is provided in the top right.

5.7.2 Systematics in $T_{\rm eff}$

We also explored potential systematics that may be present in the distribution of $T_{\rm eff}$ values across our stellar sample. Figure 5.13 displays the SPOCS $T_{\rm eff}$ as a function of the offset between the model and input (SPOCS) $T_{\rm eff}$ values. The top panel of Figure 5.13 shows an increase in the dispersion of offsets at high log$g$ and at high $T_{\rm eff}$. The lowest log$g$ values are clustered at the coolest stars, reflecting the inclusion of sub-giants and a few giant stars in the sample. We find no clear trend in the scatter of $T_{\rm eff}$ with metallicity in the lower panel of Figure 5.13.

5.8 Additional Applications

The methods explored in this work have a wide range of potential applications beyond the scope of this paper. The model that we have developed may be applied to any current and

future Keck HIRES spectra of individual stars to obtain not only the four primary stellar labels, but also all 15 elemental abundances determined in Brewer et al. (2016). With the abundance of new planets around bright stars being discovered by, for example, the Transiting Exoplanet Survey Satellite (TESS; Ricker et al., 2015; Wang et al., 2019; Huang et al., 2018), the number of promising targets for follow-up radial velocity observations is regularly increasing. A uniform, large-scale statistical analysis of these host star properties would allow for detailed population studies of the growing sample of known planets.

Once trained, the saved model can be quickly loaded and takes seconds to classify a spectrum, meaning that it is possible to determine all 18 labels with minimal delay after spectra have been obtained, reduced, and normalized. Thus, our model provides a powerful method to rapidly and precisely determine the properties of stars in the northern hemisphere. This may open new doors to study, for example, correlations between host star properties and planet size, composition, multiplicity, and other properties. Furthermore, it may allow for efficient stellar characterization soon after observing in order to quickly obtain stellar parameters and inform ongoing observations.

In theory, this model could be further applied to spectra obtained from other telescopes and instruments with an overlapping wavelength coverage. We demonstrated in this work that, by interpolating our spectra to a new wavelength grid, we reliably recovered the properties of stars observed with Keck's older, pre-2004 HIRES detector – a separate instrument with different systematics as compared to the newer, current detector. We have also completed preliminary tests extending this concept, showing in Worku et al. (submitted) that our model can successfully recover the primary stellar labels from spectra obtained with the Automated Planet Finder (APF; Vogt et al., 2014). Our findings suggest that, with further refinement, it may be possible to obtain all 18 labels from non-HIRES spectra using interpolated versions of our model, though potentially with higher uncertainties due to the differing systematics across instruments.

5.9 Conclusions

Throughout this work, we demonstrated applications of *The Cannon* to obtain 18 stellar labels from Keck HIRES stellar spectra using the SPOCS catalog as a training set. We explored several methods to optimize the model's performance, including outlier removal, data-driven continuum renormalization, telluric masking, label censoring, and L1 regularization. The primary outcomes of this work are as follows:

- We developed and tested a novel, efficient open-source tool that takes current (post-2004) Keck HIRES spectra as its inputs and outputs 18 stellar labels, including $T_{\rm eff}$, log$g$, $v \sin i$, and 15 stellar abundances: C, N, O, Na, Mg, Al, Si, Ca, Ti, V, Cr, Mn, Fe, Ni, and Y. The corresponding uncertainties for each parameter, which are comparable to the scatter in values across catalogs, are described in Table 5.3.

- We demonstrated that an interpolated model trained on the SPOCS catalog can return accurate stellar parameters for spectra spanning a similar parameter space and wavelength range, but obtained from a separate spectrograph.

- We applied our interpolated, re-optimized model to create a catalog of 18 stellar labels for 477 stars observed with Keck HIRES prior to its 2004 detector upgrade. These archival spectra could not be processed uniformly with the rest of the SPOCS sample with the SME program due to the older detector's more limited wavelength range. Our results are provided in Table 5.5 and can be found in full in the online of this paper.

In addition to quickly delivering stellar properties for individual stars, the high precision and robustness of parameters obtained with *The Cannon* make it a particularly powerful tool for population studies of stars. Studies comparing these stellar properties to trends in system architecture hold great potential to reveal a more comprehensive understanding of planetary

systems and their underlying correlations. Our code's capability to rapidly determine stellar parameters from individual spectra makes it possible to efficiently and uniformly analyze large samples of stars, rendering such studies much more computationally tractable than in the past. Applications to a broader range of stellar spectra may further extend this work and provide a more holistic view of the relationship between stars' properties and their surrounding environments.

5.10 Acknowledgements

We thank Andy Casey, Melissa Ness, and Debra Fischer for helpful conversations over the course of this work. M.R. is supported by the National Science Foundation Graduate Research Fellowship Program under Grant Number DGE-1752134. This work has made use of the VALD database, operated at Uppsala University, the Institute of Astronomy RAS in Moscow, and the University of Vienna. The authors wish to recognize and acknowledge the very significant cultural role and reverence that the summit of Mauna Kea has always had within the indigenous Hawaiian community. We are most fortunate to have the opportunity to conduct observations from this mountain.

`numpy` (Oliphant, 2006; Walt et al., 2011), `matplotlib` (Hunter, 2007), *The Cannon* (Ness et al., 2015; Casey et al., 2016)

5.11 Appendix: Pre-2004 Model Optimization

Here we detail the process of model optimization used to obtain our best-fitting model in Section 5.6.1. We note that our final model adopts the hyperparameters described in Table 5.2, rather than the final model obtained in Table 5.7 of this Appendix. However, the same wavelength ranges and telluric mask described in this section are also used in the final model.

5.11.1 Outlier Removal

We repeated the outlier removal procedure discussed in Section 5.5.2, again testing our model performance with no outliers removed and with $x_O = 1.5$, 3, and 10. As before, we ran extensive tests with a single wavelength order and used our results to inform a narrower set of tests for our full model that includes all 12 wavelength orders. With this approach, we were able to examine a wider range of models before the associated computation time became prohibitive.

We found that echelle order 6, spanning $5366 - 5432$ Å, performed best in these tests with $x_O = 1.5$. We accordingly adopted this configuration as our representative single-order base test case moving forward. Unsurprisingly, this order fully overlaps with the best-performing order from our post-2004 tests, indicating the high information content of this wavelength range.

The $x_O = 3$ case returned the lowest χ^2 value in our interpolated, pre-2004 model with all 12 echelle orders included. For both our single-order and all-orders interpolated models, we obtained the lowest χ^2 value when implementing a more liberal outlier removal criterion than in the previous model optimized for current Keck HIRES data. This may reflect the sparser training set used in this section. The training set used for pre-2004 model testing contains significantly fewer stars (865 before outlier removal) than the post-2004 case, where 80% of the pre-labeled vetted SPOCS stars were used for training (961 before outlier removal). The removal of outliers reduces the parameter space over which a model can be considered reliable; however, it also decreases the chance that the edges of the parameter space, where stars are poorly sampled, are falsely included within the reliable parameter range.

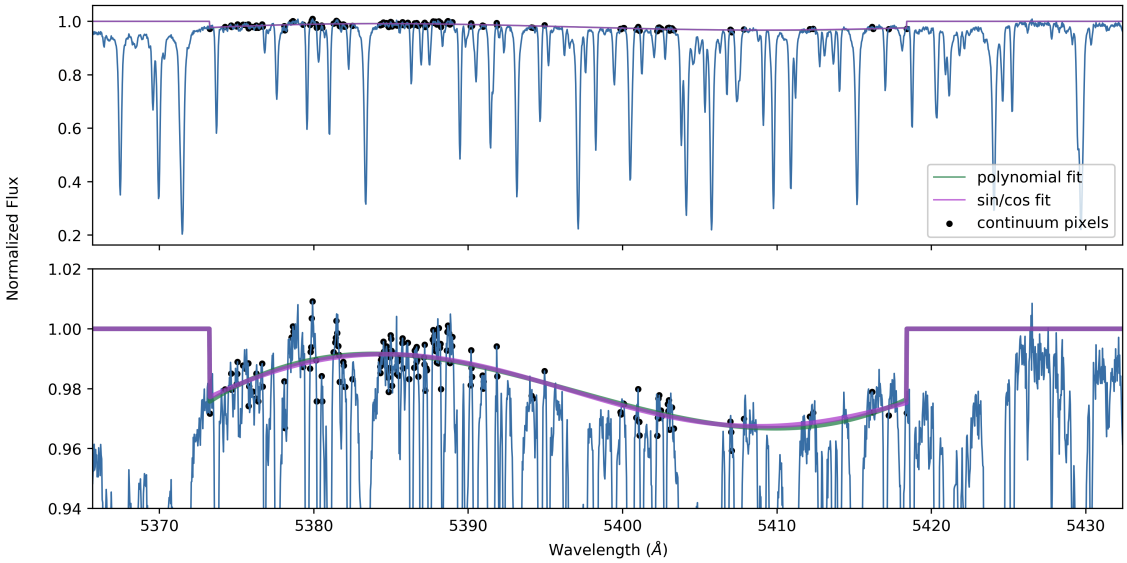

Figure 5.14: Sample continuum renormalization fit over the spectrum of HD 36130, shown in blue, for the best-fitting single order wavelength range of our pre-2004 Keck HIRES data. As in Figure 5.2, the polynomial fit is shown in green and the sin/cos fit is in purple. Continuum pixels are denoted by black markers, where here we show the $N = 50$ case.

5.11.2 Data-driven Continuum Renormalization

To explore several possible model configurations, we again used four different thresholds for our data-driven continuum pixel selection: $N = 50$, 60, 70, and 80. As in Section 5.5.3.1, we first trained our model to find the $N\%$ of pixels with coefficients closest to zero for each of our four primary labels. We then selected the pixels that overlapped between these sets and fell within 1.5% of the continuum baseline.

For per-label pixel cuts at the 50th, 60th, 70th, and 80th percentile, the final percentage of pixels identified as "true" continuum pixels in our single-order fit was roughly 7%, 12%, 16%, and 20%, respectively, with some variation occurring from spectrum to spectrum. A sample spectrum of G0 star HD 36130 is shown in Figure 5.14 with $N = 50$, which, for this spectrum, results in 7.6% of all pixels being selected as continuum pixels. Figure 5.14 displays both the selected continuum pixels and the two corresponding fits for comparison. As in Section 5.5.3.1, and as illustrated by Figure 5.14, the two functional forms – sin/cos

and polynomial fits – typically provide similar results.

We found that both the polynomial and sin/cos renormalization functions consistently improved our single-order fit for all N values. The best-performing single-order model, which we adopted for our ongoing testing, used the $N = 50$ threshold with a sin/cos renormalization.

To apply these results to our full model with all 12 orders, we again tested the $N = 50$ case with both a sin/cos and polynomial renormalization. We found that both renormalization schemes degraded our results with all orders included. However, upon closer examination of the individual per-order renormalization fits, we found that the $N = 50$ threshold resulted in a very sparse and potentially unreliable continuum fit in several echelle orders. To determine whether the inclusion of more "continuum" pixels would improve our fit, we also tested the $N = 70$ case, which provided only marginally less reliable results in the single-order case and which resulted in our best fits in Section 5.5.3.1.

As suspected, the $N = 70$ case did improve our results for both a polynomial and sin/cos fit with all echelle orders incorporated. This suggests that, while order 6 ($5366 - 5432$ Å) performs best with $N = 50$, sufficiently few pixels are selected in other echelle orders with $N = 50$ such that the continuum fit is not consistently reliable across orders. With $N = 70$, a substantially larger number of pixels is incorporated into the continuum fit, leading to a better approximation of the underlying baseline. While both fits improved our results, we obtained the lowest χ^2 value with a sin/cos renormalization applied to all echelle orders. As a result, we implemented sin/cos renormalizations with $N = 50$ in our single-order models and with $N = 70$ in our all-orders models moving forward.

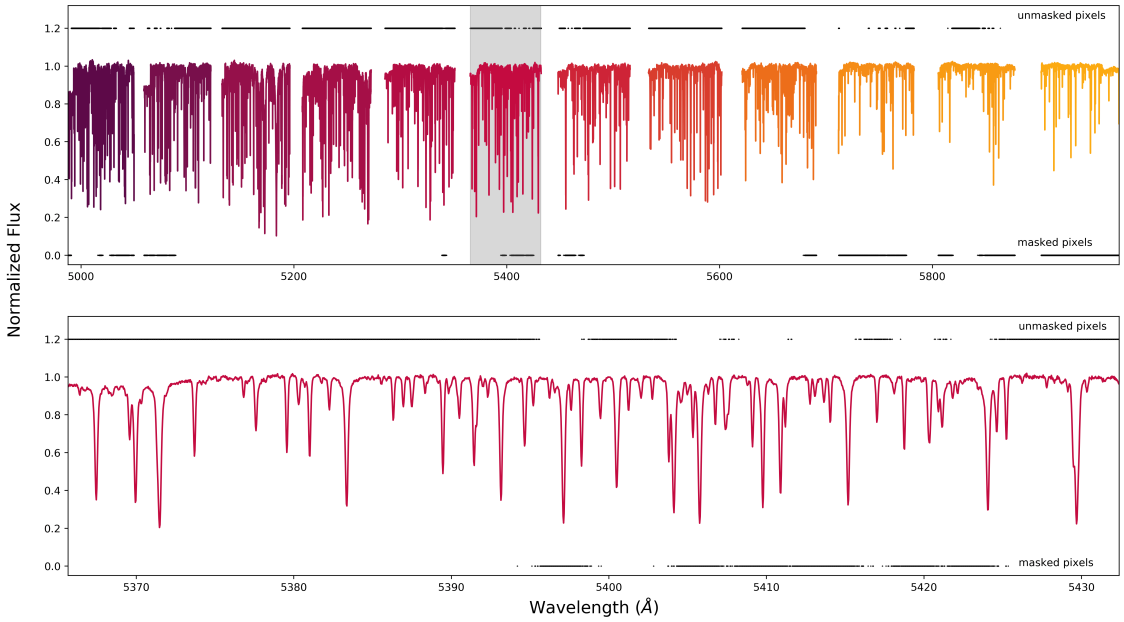

Figure 5.15: Top: Full, continuum renormalized spectrum of sample star HD 36130, showing the 12 overlapping wavelength regions shared across our pre- and post-2004 spectra. The portion of the spectrum corresponding to the lower panel is highlighted in gray. Bottom: Zoom-in of only echelle order 6, ranging from $5366 - 5432$ Å. Black markers denote the telluric pixels ("masked pixels") at the bottom of each panel, as well as the non-telluric pixels ("unmasked pixels") at the top of each panel.

5.11.3 Telluric Masking

We also applied telluric masking in our model configuration tests. First, we created a new telluric mask by finding the locations of telluric lines in each spectrum from our pre-2004 training set and creating a mask for each spectrum. The telluric lines do not match up exactly in every spectrum due to the differing barycentric corrections and radial velocities of each observed star. Thus, we combined all of our individual masks to create one master mask that we applied to all spectra. This mask is visualized in Figure 5.15 with a sample spectrum from HD 36130 shown for reference.

With our telluric mask in place, 648 of the 2048 pixels in our single, best-performing echelle order were excluded from our fit, with the masked pixels shown in Figure 5.15. Despite the loss of information, we found that this masking slightly improved our results.

This reflects the tradeoff between removing noise and removing signal with substantial masking implemented. We continued to apply telluric masking in our ongoing single-order tests to account for the net improvement observed in our label recovery.

In our all-orders tests, we instead found that telluric masking degraded our results, and we chose not to include it within our final model as a result. This suggests that, in our interpolated model with all orders included, telluric masking reduces the signal more than it reduces the noise in our model, leading to poorer performance overall.

5.11.4 Censoring

As in Section 5.5.3.3, we also applied censoring at the 5%, 15%, 50%, 85%, and 95% levels – meaning that, for example, only the most highly varying 5% of all nonzero pixels for each label were used when fitting at the 5% censoring level. We ran two sets of tests in which we censored (1) all labels or (2) only the primary 4 labels ($T_{\rm eff}$, log$g$, $v \sin i$, and [Fe/H]), resulting in a total of 10 test cases. Two of these cases – 85% and 15%, each with only the primary four stellar labels censored – are illustrated in Figure 5.16, which depicts the masked/unmasked pixel locations for the each case.

Ultimately, we found that both the 85% and 95% test cases with 4 labels censored improved our single-order model results. The 85% test case performed slightly better and we therefore used this case moving forward. In general, heavier censoring – using smaller samples of pixels to fit each label – led to less reliable results than lighter censoring.

We then tested these two best-performing cases in our all-orders model to determine whether they would produce improvements in our label recovery. We found that, with all orders included in the model, censoring 4 labels at either the 85% or 95% level provided no substantial improvements to our model performance. This reflects the tradeoff between removing noisy pixels and eventually removing pixels that provide useful information.

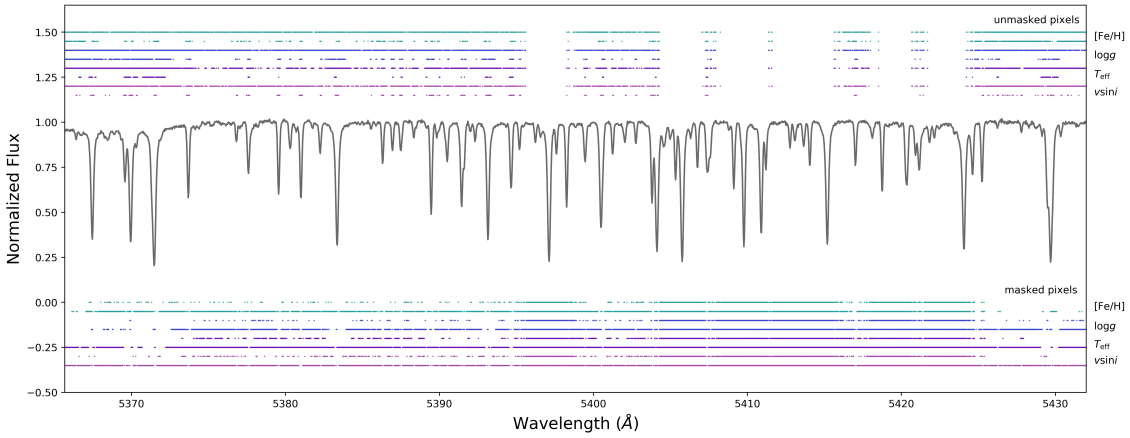

Figure 5.16: Sample censored wavelengths for sample star HD 36130, selected for the primary four stellar labels: [Fe/H] (green), log$g$ (blue), $T_{\rm eff}$ (violet), and $v \sin i$ (purple). The unmasked pixels corresponding to each label are shown above the spectrum, and the masked, unused pixels are below. Masks for each label are provided in pairs, where the upper line in each color corresponds to the 85% mask, while the lower line corresponds to the 15% mask. "Unmasked" pixels are included in the analysis for that label, while "masked" pixels are excluded.

Thus, we did not incorporate censoring into our final, all-orders model and instead elected to use it only in our single-order model.

5.11.5 L1 Regularization

Lastly, we applied regularization to our best-fitting individual order with $\Lambda = 1$, 10, 100, 1,000, and 10,000, resulting in label density distributions almost identical to those in Figure 5.5. We found that lower regularization values generally provided better results than higher ones, but that any of our tested regularization values degraded the model results relative to the case with no regularization. Thus, we chose not to include regularization in either our single-order or all-orders model configuration.

While we did not find that the tested regularization values led to an improved χ^2 value, this does not imply that *no* values of regularization would improve our results. Of our tested Λ values, we obtained the best results with $\Lambda = 100$. This suggests that, if any Λ value

exists that would improve our model results, it is likely between $\Lambda = 10$ and $\Lambda = 1000$. Behmard et al. (2019) also tested regularization values on a grid spanning $\Lambda = 10^{-6}$ to $\Lambda = 10^{2}$ and found that no tested Λ values improved their model results. Given that our test set results already had low scatter and that additional benefits from fine-tuning would likely be only marginal, we found that it was not practical for our purposes to sample a finer grid of possible values.

5.11.6 Final Model Configuration

Our best-performing pre-2004 model configurations for both the single-order and all-orders cases are provided in Table 5.7. Both the best-fitting single-order and all-orders models are characterized by strict outlier thresholds, removing several stars from the training/test sets. This improves our model performance with the tradeoff that our model spans a smaller parameter space and cannot be applied to as wide a range of stars. Our final all-orders model obtained in this section, with $x_O = 3$, ultimately includes 334 test set stars and 858 training set stars, returning $\chi^2 = 4.79$. We emphasize that the final configuration used to obtain the catalog in Table 5.5 is not this model, but, rather, one that applies the same hyperparameters as the optimized post-2004 model. Our final model does, however, use the same telluric mask and the same wavelength ranges described throughout this Appendix.

	Single order	All orders
x_O	1.5	3
Renormalization	N=50, sin/cos	N=70, sin/cos
Telluric masking	included	not included
Censoring	85%, 4 labels	none
Λ	0	0

Table 5.7: Optimized training configuration for both our single-order model and our model incorporating all orders, developed to classify pre-2004 Keck HIRES spectra. The single-order run spans wavelength range $5366 - 5432$ Å. We note that our final model does not use this configuration, since the hyperparameters found in our analysis of current Keck spectra provided further improved results.

Chapter 6

A Stellar Obliquities Survey as a Probe of the Dynamical History of Hot Jupiters

Originally published as:

Rice, M., Wang, S., Howard, A.W., Isaacson, H., Dai, F., Wang, X.-Y., Beard, C., Behmard, A., Brinkman, C., Rubenzahl, R.A., & Laughlin, G. 2021, The Astronomical Journal, 162, 182

6.1 Abstract

Obliquity measurements for stars hosting relatively long-period giant planets with weak star-planet tidal interactions may play a key role in distinguishing between formation theories for shorter-period hot Jupiters. Few such obliquity measurements have been made to date due to the relatively small sample of known wide-orbiting, transiting Jovian-mass planets and the challenging nature of these targets, which tend to have long transit durations and orbit faint stars. We report a measurement of the Rossiter-McLaughlin effect across the transit of K2-140 b, a Jupiter-mass planet with period $P = 6.57$ days orbiting a $V = 12.6$ star. We find that K2-140 is an aligned system with projected spin-orbit angle $\lambda = 0.5\,^{\circ} \pm 9.7\,^{\circ}$, suggesting a dynamically cool formation history. This observation builds towards a population of tidally detached giant planet spin-orbit angles that will enable a direct comparison with the distribution of close-orbiting hot Jupiter orbital configurations, elucidating the prevalent formation mechanisms of each group.

6.2 Introduction

While the obliquities of many hot-Jupiter-hosting stars have been determined through measurements of the Rossiter-McLaughlin effect (Rossiter, 1924; McLaughlin, 1924), only a handful of these measurements have been made for systems with planets orbiting at larger distances from their host star, where star-planet tidal interactions are too weak to significantly influence the planets' orbital alignment. We refer to these wide-orbiting planets as "tidally detached." Because the probability p of transit for a given planet with semimajor axis a falls as $p \propto 1/a$, fewer tidally detached transiting giant planets have been discovered to date as compared with closer-in hot Jupiters. Due to the volume-limited sample of bright stars, those tidally detached Jupiters that have been discovered are often challenging Rossiter-McLaughlin targets, orbiting stars too faint to obtain a sufficient

number of high-resolution spectroscopic observations across a single transit with any but the largest existing ground-based telescopes.

The vast majority of spin-orbit angle measurements have, consequently, been made for planets on tight orbits with $a/R_* \lesssim 12$. However, the obliquities of stars hosting wider-separation giant planets may provide the evidence necessary to distinguish between the various proposed formation mechanisms for hot Jupiters (Winn & Fabrycky, 2015; Dawson & Johnson, 2018). Furthermore, in the classical quiescent formation framework for hot and warm Jupiters, the long tidal realignment timescales of these systems enable a direct measurement of the primordial dispersion in protoplanetary disk misalignments.

We present a Rossiter-McLaughlin measurement across the transit of K2-140 b with the High Resolution Echelle Spectrometer (HIRES; Vogt et al., 1994) on the 10-meter Keck I telescope. K2-140 (EPIC 228735255) is a $V = 12.6$ G5 star hosting a 1.019±0.070$\mathrm{M_J}$ planet, K2-140 b, with a $P = 6.57-$day orbital period (Giles et al., 2018). With $a/R_* = 12.88$, K2-140 b is one of the widest-separation tidally detached planets to date with a measured spin-orbit angle. This is the first measurement in our Stellar Obliquities in Long-period Exoplanet Systems (SOLES) survey designed to expand the sample of Rossiter-McLaughlin measurements for wide-separation exoplanets.

K2-140 b was first characterized using photometry from the *K2* mission (Howell et al., 2014), a re-purposed extension of the *Kepler* mission (Borucki et al., 2010), after two of its reaction wheels failed. Giles et al. (2018) used radial velocity (RV) observations from CORALIE (Queloz et al., 2000) and the High Accuracy Radial velocity Planet Searcher (HARPS; Pepe et al., 2000) to confirm the planetary nature of K2-140 b and to constrain its physical and orbital properties.

The planet was independently characterized in Korth et al. (2019), which incorporated RV data from the FIbre-fed Échelle Spectrograph (FIES) (Telting et al., 2014) within their

analysis. While Giles et al. (2018) found a low but nonzero eccentricity for the planet ($e = 0.120^{+0.056}_{-0.046}$), Korth et al. (2019) found that the planet was consistent with $e = 0$. Data from both of these past studies is incorporated within our joint analysis.

6.3 Observations

We obtained 18 radial velocity measurements of K2-140 with the Keck/HIRES instrument from 9:25-15:40 UT on Feb 24, spanning a full transit of K2-140 b. Conditions were favorable throughout most of the observing period, with typical seeing ranging from $1.0'' - 1.3''$. A spike in humidity during the post-transit baseline observations led to a telescope dome closure at 14:40 UT, resulting in a ∼40-minute gap in data before the last radial velocity measurement.

All RV observations were obtained using the C2 decker ($14'' \times 0.861''$, $R = 60,000$) and an iodine absorption cell, which imprints a dense forest of molecular iodine features onto each spectrum to enable high Doppler precision (Butler et al., 1996). The $14''$ length of the C2 decker allows for direct sky subtraction, improving RV precision for faint stars such as K2-140. The median exposure time was $1,119$ seconds, with ∼34k exposure meter counts per spectrum.

Our dataset was reduced using the California Planet Search pipeline outlined in Howard et al. (2010), and we obtained a typical signal-to-noise ratio of 71 per pixel from the reduced spectra. The HIRES radial velocity results and uncertainties can be found in Table 6.1 and are shown in the rightmost panel of Figure 6.1. We include the S-index and associated uncertainty at each observation in Table 6.1 for reference.

We also obtained a 45-minute iodine-free HIRES exposure of K2-140 using the B3 decker ($14.0'' \times 0.574''$, $R = 72,000$) two nights after the measurement of the Rossiter-McLaughlin effect, during UT Feb 26. This template observation was used to calibrate our RVs and to

precisely determine stellar parameters (see Section 6.5). Conditions were favorable during this measurement, and seeing was $1.2''$. Our reduced template had a signal-to-noise ratio of 123 per pixel (106k exposure meter counts).

Table 6.1: HIRES radial velocities for the K2-140 system.

Time (BJD)	RV (m/s)	$\sigma_{\rm RV}$ (m/s)	S-index	σ_S
2459269.903807	17.62	2.48	0.165	0.001
2459269.917697	1.43	2.59	0.166	0.001
2459269.931633	17.18	2.55	0.160	0.001
2459269.944029	23.36	2.58	0.168	0.001
2459269.956264	23.63	2.38	0.161	0.001
2459269.968059	22.22	2.47	0.166	0.001
2459269.978707	14.54	2.35	0.159	0.001
2459269.989634	11.66	2.34	0.160	0.001
2459270.000977	-2.85	2.31	0.165	0.001
2459270.013142	-15.76	2.19	0.167	0.001
2459270.026233	-14.19	2.41	0.163	0.001
2459270.040065	-18.89	2.45	0.160	0.001
2459270.054024	-23.71	2.41	0.159	0.001
2459270.069384	-10.01	2.63	0.161	0.001
2459270.083054	2.645	2.96	0.155	0.001
2459270.097753	-20.13	2.85	0.152	0.001
2459270.110694	-14.47	3.08	0.147	0.001
2459270.153058	-14.61	2.71	0.155	0.001

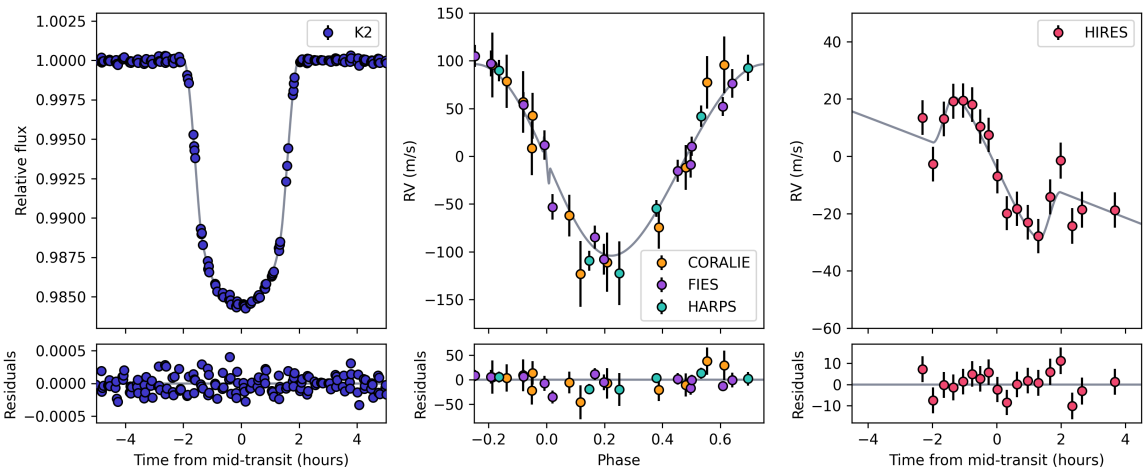

Figure 6.1: Joint fit to photometry, out-of-transit RV data, and the in-transit Rossiter-McLaughlin RV data obtained for K2-140 b. The model is shown in gray, while data is provided in color with modeled constant offsets and jitter terms included. The associated residuals are provided below each panel.

6.4 Obliquity Modeling

To determine the sky-projected spin-orbit angle λ for K2-140 b, we used the `allesfitter` Python package (Günther & Daylan, 2021) to jointly model the in-transit HIRES radial velocity data together with photometry from $K2$ and archival radial velocity datasets available from the FIES, CORALIE, and HARPS spectrographs.

All fitted parameters listed in Table 6.2 were allowed to vary, and each parameter was initialized with uniform priors. Initial guesses for P, T_0, $\cos i$, $R_p/R_\star$, $(R_\star + R_p)/a$, K, $\sqrt{e}\cos\omega$, and $\sqrt{e}\sin\omega$ were obtained using values from Giles et al. (2018). The two limb darkening coefficients q_1 and q_2 were each initialized with values of 0.5. We accounted for potential radial velocity offsets between each separate spectrograph, with priors bounded by ± 1000 m/s. Jitter terms were modeled separately for each instrument and added in quadrature to the instrumental uncertainties. λ was allowed to vary between $-180\,^\circ$ and $+180\,^\circ$.

We ran an affine-invariant Markov Chain Monte Carlo (MCMC) analysis with 100 walkers

to sample the posterior distributions of all model parameters. The best-fit model parameters and their associated 1σ uncertainties were extracted after obtaining 500,000 accepted steps per walker. Our results are listed in Table 6.2 and are in good agreement with the associated values obtained by Giles et al. (2018) and Korth et al. (2019).

The best-fit joint model is shown in Figure 6.1 together with each dataset included in the analysis, as well as the residuals of each fit. The fitted and derived parameters corresponding to this model are provided in Table 6.2. We obtain a low but nonzero eccentricity $e = 0.069^{+0.042}_{-0.028}$ for K2-140 b, in agreement with the value derived by Giles et al. (2018). K2-140 is consistent with alignment, with $\lambda = 0.5\,^{\circ} \pm 9.7\,^{\circ}$ and $v \sin i_* = 2.51 \pm 0.38$ km/s.

6.5 Stellar Parameters

An understanding of host star properties can help to contextualize the evolutionary pathways through which a system may have reached its current state. We extracted stellar parameters from our Keck/HIRES template spectrum of K2-140 using the data-driven spectroscopic modeling program *The Cannon* (Ness et al., 2015; Casey et al., 2016), following the methods of Rice & Brewer (2020).

Given a set of uniformly processed input training spectra and associated stellar "labels" – that is, stellar parameters and elemental abundances – *The Cannon* constructs a generative model describing the probability density function of flux at each wavelength as a function of the labels. The model can then be applied to a new set of spectra, uniformly processed in the same manner as the training set, to obtain the associated stellar labels.

We trained *The Cannon* using the uniformly analyzed Spectral Properties of Cool Stars (SPOCS) catalogue (Brewer et al., 2016) of 18 stellar labels, including 3 global stellar parameters ($T_{\rm eff}$, $\log g$, $v \sin i_*$), and 15 elemental abundances: C, N, O, Na, Mg, Al, Si, Ca, Ti, V, Cr, Mn, Fe, Ni, and Y. The sample of 1202 Keck/HIRES spectra vetted in Rice &

Parameter	Description	Priors	Value	$+1\sigma$	-1σ
Fitted Parameters:					
$R_p/R_\star$	Planet-to-star radius ratio	$\mathcal{U}(0;1)^*$	0.1163	0.0012	0.0011
$(R_\star + R_p)/a$	Sum of radii divided by the orbital semimajor axis	$\mathcal{U}(0;1)$	0.0960	0.0043	0.0039
$\cos i$	Cosine of the orbital inclination	$\mathcal{U}(0;1)$	0.0518	0.0053	0.0048
T_0	Mid-transit epoch (BJD)	$2458435.7138^\dagger$	2458429.1426	0.0015	0.0015
P	Orbital period (days)	$\mathcal{U}(5.569188; 7.569188)$	6.569199	1.2e-05	1.2e-05
K	Radial velocity semi-amplitude (m/s)	$\mathcal{U}(0; 1000)$	106.5	4.7	4.7
$\sqrt{e}\cos\omega$	Eccentricity parameter 1	$\mathcal{U}(-1.0; 1.0)$	-0.156	0.067	0.053
$\sqrt{e}\sin\omega$	Eccentricity parameter 2	$\mathcal{U}(-1.0; 1.0)$	0.189	0.098	0.15
q_1	Quadratic limb darkening coefficient 1	$\mathcal{U}(0.0; 1.0)$	0.74	0.15	0.14
q_2	Quadratic limb darkening coefficient 2	$\mathcal{U}(0.0; 1.0)$	0.175	0.072	0.055
$\Delta_{\mathrm{RV,FIES}}$	RV offset, FIES (m/s)	$\mathcal{U}(-1000.0; 1000.0)$	1.1287	0.0044	0.0042
$\Delta_{\mathrm{RV,CORALIE}}$	RV offset, CORALIE (m/s)	$\mathcal{U}(-1000.0; 1000.0)$	1.2141	0.0083	0.0083
$\Delta_{\mathrm{RV,HARPS}}$	RV offset, HARPS (m/s)	$\mathcal{U}(-1000.0; 1000.0)$	1.2459	0.0053	0.0053
$\Delta_{\mathrm{RV,HIRES}}$	RV offset, HIRES (m/s)	$\mathcal{U}(-1000.0; 1000.0)$	0.0051	0.0034	0.0034
λ	Sky-projected spin-orbit angle (°)	$\mathcal{U}(-180.0; 180.0)$	0.5	9.7	9.7
$v \sin i_*$	Sky-projected stellar rotational velocity (km/s)	$\mathcal{U}(0.0; 20.0)$	2.51	0.38	0.38
Derived Parameters:					
R_p	Planetary radius ($\mathrm{R_J}$)	...	1.203	0.076	0.076
M_p	Planetary mass ($\mathrm{M_J}$)	...	1.13	0.12	0.11
b	Impact parameter	...	0.574	0.022	0.022
T_{14}	Transit duration (h)	...	3.946	0.016	0.016
δ	Transit depth	...	0.015490	$4.2e-05$	$4.2e-05$
a	Semimajor axis (au)	...	0.0575	0.0046	0.0043
i	Inclination (°)	...	87.03	0.27	0.30
e	Eccentricity	...	0.069	0.042	0.028
ω	Argument of periastron (°)	...	131	38	20
u_1	Limb darkening parameter 1	...	0.300	0.088	0.080
u_2	Limb darkening parameter 2	...	0.56	0.15	0.16

Table 6.2: System properties derived for K2-140.

* $\mathcal{U}(a; b)$ is a uniform prior with lower and upper limits a and b, respectively.

† We provided a reference value for T_0. During the fit, `allesfitter` can shift epochs to the data center to derive an optimal T_0.

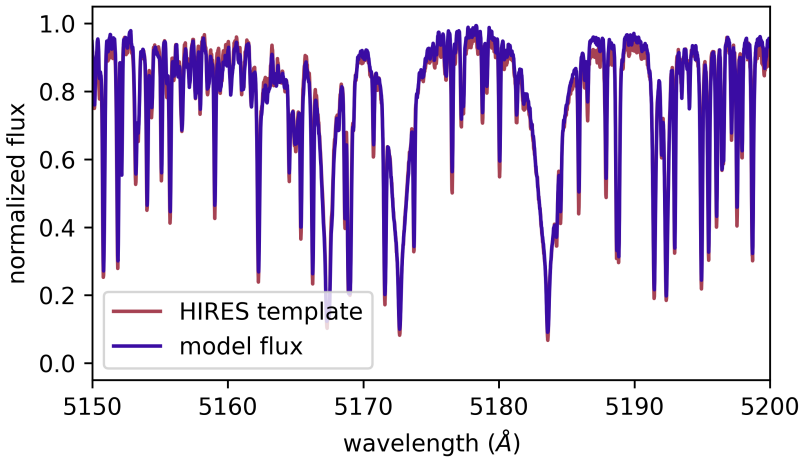

Figure 6.2: Sample segment of the model spectrum returned by *The Cannon*, shown alongside the K2-140 HIRES template spectrum in the vicinity of the Mg Ib triplet (lines at 5167, 5172, and 5183 Å).

Brewer (2020) was applied as our training/test set.

The continuum baseline of each spectrum was uniformly fit and divided out using the iterative polynomial fitting procedure outlined in Valenti & Fischer (2005). Then, we split the SPOCS sample into an 80%/20% training/test split, applied the telluric mask from Rice & Brewer (2020) to all spectra, and trained the model, using the scatter of the test set results to determine the uncertainties of each extracted parameter. Finally, the trained model was applied to the newly acquired K2-140 Keck/HIRES template spectrum. A segment of the obtained model spectrum is shown in comparison with the HIRES template data in Figure 6.2.

Our results are provided in Table 6.3 together with estimates from previous works for reference. Our values for $T_{\rm eff}$, $\log g$, and [Fe/H] are in agreement with those acquired in previous studies. The $v \sin i_*$ value obtained using *The Cannon* is lower than previous estimates but is in agreement with the RM joint fit results and Korth et al. (2019) within 1σ. We find that K2-140 is metal-enriched relative to solar abundances, consistent with past evidence showing that short-period giant planets are more common around metal-rich stars (Fischer & Valenti, 2005).

Parameter	Unit	The Cannon	RM joint fit	Giles+ 2019	Korth+ 2019
$T_{\rm eff}$	K	5610 ± 59	-	5654 ± 55	5585 ± 120
log$g$	cm/s^2	4.4 ± 0.1	-	$4.452^{+0.010}_{-0.009}$	4.4 ± 0.2
$v \sin i_*$	km/s	2.31 ± 1.07	2.51 ± 0.38	3.8 ± 0.2	3.6 ± 1.0
[Fe/H]	dex	0.18 ± 0.04	-	0.12 ± 0.045	0.10 ± 0.10
[C/H]	dex	0.14 ± 0.08	-	-	-
[N/H]	dex	0.05 ± 0.09	-	-	-
[O/H]	dex	0.06 ± 0.09	-	-	-
[Na/H]	dex	0.13 ± 0.07	-	-	0.12 ± 0.10
[Mg/H]	dex	0.12 ± 0.04	-	-	0.27 ± 0.10
[Al/H]	dex	0.22 ± 0.13	-	-	-
[Si/H]	dex	0.18 ± 0.05	-	-	-
[Ca/H]	dex	0.24 ± 0.04	-	-	0.12 ± 0.10
[Ti/H]	dex	0.24 ± 0.05	-	-	-
[V/H]	dex	0.17 ± 0.06	-	-	-
[Cr/H]	dex	0.19 ± 0.05	-	-	-
[Mn/H]	dex	0.18 ± 0.06	-	-	-
[Ni/H]	dex	0.18 ± 0.05	-	-	0.20 ± 0.10
[Y/H]	dex	0.23 ± 0.12	-	-	-

Table 6.3: Stellar parameters for K2-140.

6.6 Discussion

6.6.1 Implications of the Low λ of K2-140 b

At $T_{\rm eff} = 5585$ K, K2-140 should have a convective envelope. In the framework of equilibrium tides, the planet's timescale for realignment from turbulent friction would thus follow

$$\tau_{CE} = \frac{10^{10}\text{yr}}{(M_p/M_*)^2}\left(\frac{a/R_*}{40}\right)^6, \tag{6.1}$$

where τ_{CE} is the realignment timescale for host stars with convective envelopes, and M_p/M_* is the planet-to-star mass ratio (Zahn, 1977; Albrecht et al., 2012). For K2-140 b, $\tau_{CE} = 1.2 \times 10^{13}$ yr, longer than the age of the Universe. As a result, we conclude that K2-140 b was likely aligned at the time of protoplanetary disk dispersal.

The theory of equilibrium tides presented in Equation 6.1 is employed as a simplified heuristic for a broader theoretical framework that has substantially advanced over recent

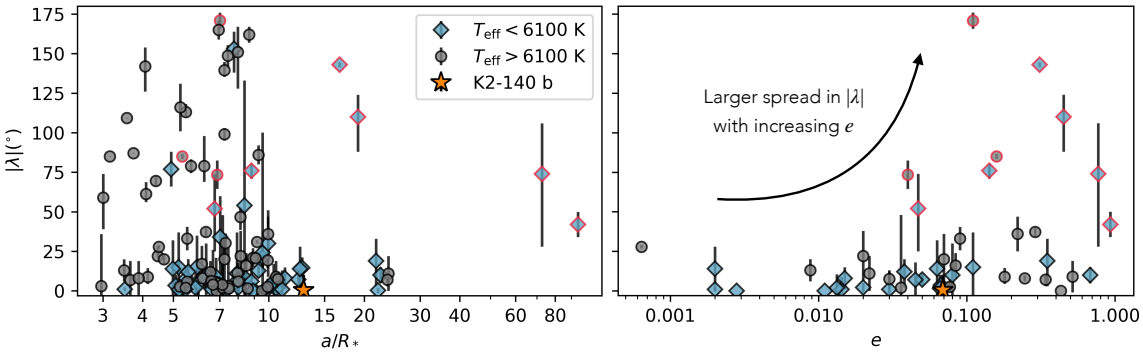

Figure 6.3: *Left:* K2-140 b relative to other Rossiter-McLaughlin measurements of giant planets ($M > 0.3 M_J$) in the TEPCat catalogue (Southworth, 2011), with a and R_* obtained through cross-matching with the NASA Exoplanet Archive and missing values filled in from the Extrasolar Planets Encyclopaedia. In a/R_* space, K2-140 b lies just exterior to most planets with measured spin-orbit angles. This makes it one of only a handful of planets with a tidal realignment timescale $\tau >>$ stellar age, such that it would not have had time to realign if it had been misaligned at the time of protoplanetary disk dispersal. *Right:* Distribution of measured obliquities as a function of eccentricity for planets with $e > 0$. The range of measured λ values is larger at high e for planets orbiting stars both above and below the Kraft break. Planets with $|\lambda| > 40\,^\circ$ and $e > 0.01$ are outlined in red in both panels. Among the planets with $e \neq 0$, the spread in misalignments is larger at higher eccentricities.

years (Ogilvie, 2014). A key problem in the equilibrium tides framework is that, under standard assumptions, hot Jupiters should experience rapid orbital decay. Lai (2012) demonstrated that one component of the tidal potential (the "obliquity tide") can excite inertial waves in the convective envelopes of cool stars, which, when damped, can enhance dissipation of the stellar obliquity without shrinking the companion's orbit. Obliquity tides have been further explored in additional work (Ogilvie, 2013; Lin & Ogilvie, 2017; Anderson et al., 2021). While obliquity tides are not included in our analysis, previous work has demonstrated that Equation 6.1 is a useful heuristic revealing that low-obliquity systems tend to have shorter tidal timescales than high-obliquity systems (Albrecht et al., 2012).

K2-140 b is one of only 13 giant planets with a measured spin-orbit angle at $a/R_* > 12$, as shown in the left panel of Figure 6.3. Of these planets, it is one of only a few with

an aligned orbit despite its long tidal realignment timescale. The alignment of K2-140 b suggests that at least some hot Jupiters form through quiescent pathways, such as in-situ formation or disk migration in an initially aligned disk.

The highly misaligned planets at large a/R_* each have high eccentricities (right panel of Figure 6.3), indicative of strong dynamical interactions that may have produced both elevated obliquities and eccentricities. The coexistence of this population with the dynamically quiescent K2-140 system at large a/R_* suggests that there are multiple hot Jupiter formation channels.

No neighboring planets have yet been found in the K2-140 system, despite previous observations showing that inner, coplanar companions are common for slightly longer-period warm Jupiters ($10 < P < 200$ days; Huang et al., 2016) and predictions that outer, mutually-inclined companions with $P \lesssim 100$ days should be regularly produced by in-situ hot Jupiter formation via core accretion (Batygin et al., 2016). However, additional low-mass or distant planets in the system cannot be ruled out by the existing observations, with radial velocity residuals of tens of m/s and an observing baseline of only $\sim$4 years.

6.6.2 Motivation for Additional Obliquity Measurements in Tidally Detached Systems through SOLES

Existing Rossiter-McLaughlin measurements of wide-orbiting planets suggest that this population may have an intrinsically different spin-orbit angle distribution from the shorter-period hot Jupiter population. In particular, the trend of systematically lower obliquities observed for stars at temperatures below the Kraft break ($T_{\rm eff} \approx 6100$ K; Kraft, 1967; Winn et al., 2010; Schlaufman, 2010), which may result from tidal damping (e.g. Wang et al., 2021), is not immediately evident at longer orbital periods, as shown in Figure 6.4.

Conversely, Figure 6.4 reveals tentative (2.79σ) evidence that relatively long-period planets

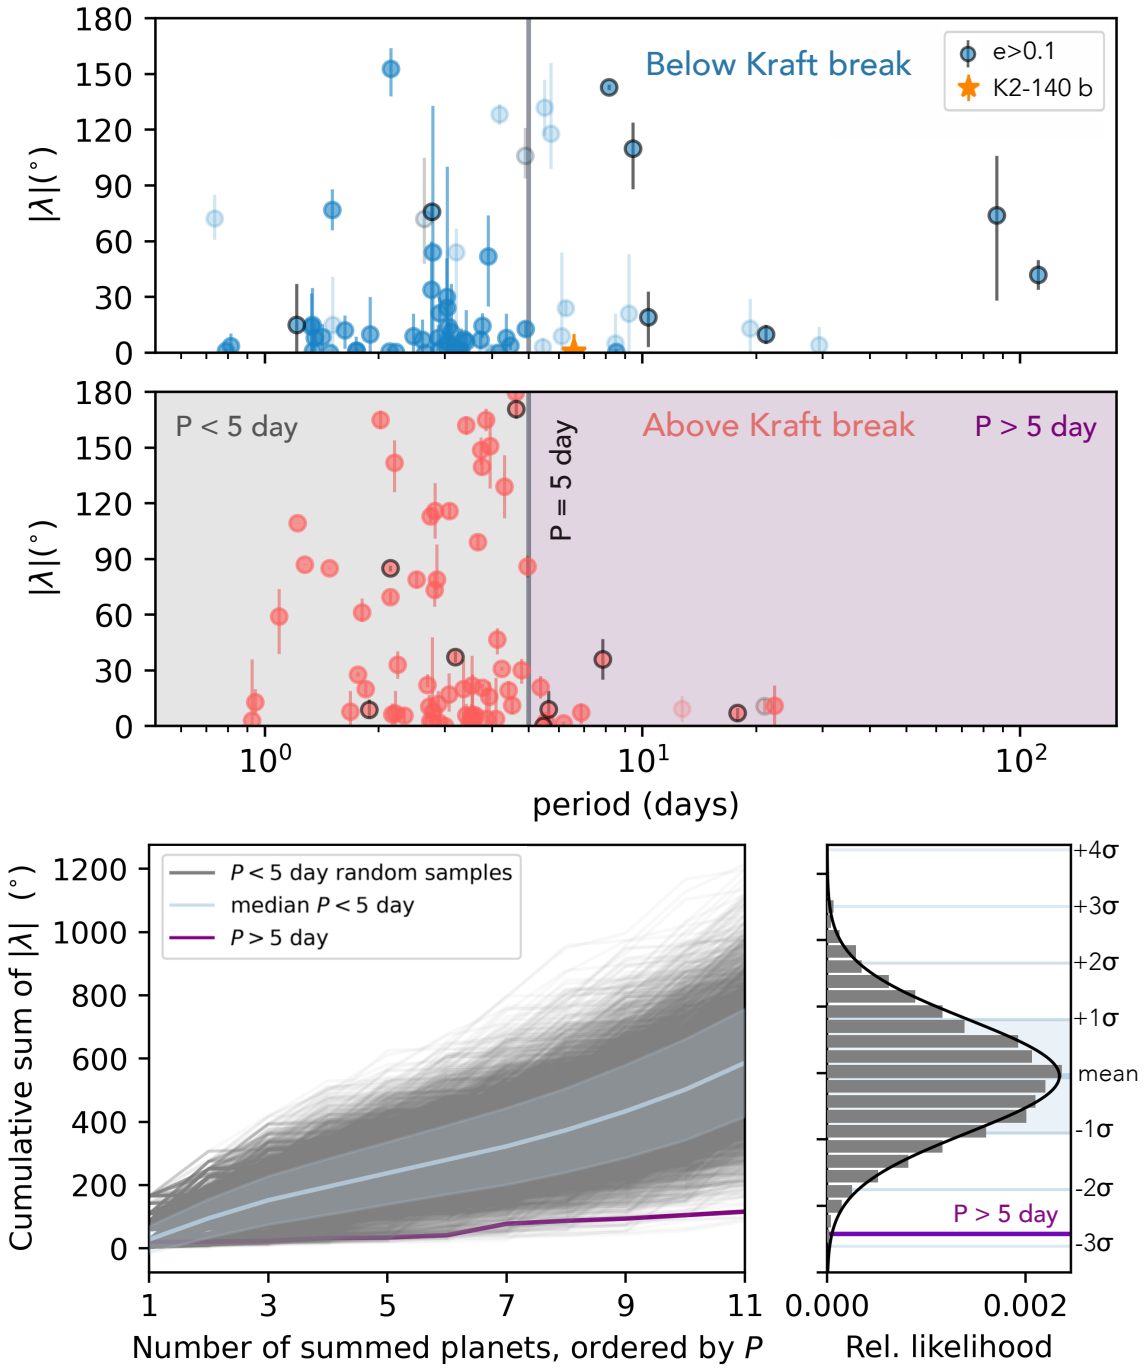

Figure 6.4: Obliquity distribution of all stars below (top panel) and above (middle panel) the Kraft break with measured λ values in the TEPCat catalogue, provided as a function of the companion planet's orbital period. Measurements for companion planets with $e > 0.1$ are bordered in black. Measurements for companion planets with $M < 0.3 M_J$ are shown at lower opacity. The cumulative sums in the bottom panel include only planets above the Kraft break, comparing the $P > 5$ day population (purple) with 5000 $P < 5$ day samples (gray), randomly sampled without replacement. The histogram on the bottom right shows a vertical cut through the final cumulative sum distribution. The $P > 5$ day cumulative sum is a 2.79σ outlier from the random draw distribution.

($P > 5$ days) around hot stars are preferentially *less* misaligned than their shorter-period counterparts. Furthermore, the most misaligned $P > 5$ day planet in the middle panel, KELT-6 b ($P = 7.8$ days, $|\lambda| = 36 \pm 11\,^{\circ}$; Damasso et al., 2015), orbits a star directly bordering the nominal Kraft break ($T_{\rm eff} = 6102\,{\rm K}$; Collins et al., 2014), making it an ambiguous member of its group. We note that hot stars typically have larger radii and thus stronger tidal dissipation at a given orbital period; however, a similar trend has also been previously suggested for a smaller population of planets at large a/R_* (Yu et al., 2018).

Longer-period planets have comparatively long tidal alignment timescales, especially around hot stars. The small spin-orbit angles observed for exoplanets around hot stars, therefore, may suggest that protoplanetary disks tend to be aligned at the time of gas dispersal (in contrast with a primordial spin-orbit misalignment from a tilted disk; Batygin, 2012; Spalding & Batygin, 2015). In this case, hot Jupiters would need to obtain their misalignments after the protoplanetary disk has dispersed, favoring high-eccentricity migration as an important formation mechanism (e.g. Wu & Murray, 2003; Fabrycky & Tremaine, 2007). Most of the misaligned planets at $P > 5$ days are on eccentric orbits, suggesting that interactions with stellar or planetary companions in the systems may have induced both elevated obliquities and eccentricities.

If a growing sample of spin-orbit angles at $P > 5$ days reveals that cool stars below the Kraft break are preferentially more misaligned than hotter stars above the Kraft break, this may also be suggestive of interactions with additional companions. In the presence of an external, inclined Jovian-mass companion, spin-orbit misalignments can be excited preferentially in cool star systems due to a secular resonance between the host star's spin axis precession frequency and nodal precession induced by interactions with the companion (Anderson & Lai, 2018). Cool stars spin down over time due to magnetic braking, enabling this resonant excitation. Hot stars, by contrast, lack a convective envelope and continue to rapidly rotate over their lifetimes. As a result, irrespective of the external giant planet

companion rate around hot stars, this mechanism should occur only in cool star systems.

The substantially inclined companions necessary to induce this mechanism would also be capable of exciting the inner planet's orbital eccentricity. This may be reflected by the apparent increase in both misalignments and eccentricities for long-period planets around cool stars, shown in the top panel of Figure 6.4. Conversely, the alignment of K2-140 b and similar planets on low-eccentricity orbits indicates that they should not have nearby giant planet companions with large ($> 10\,^{\circ}$) mutual inclinations within ~ 2 au.

If the trend of large misalignments for cool stars hosting tidally detached planets persists, while misalignments remain small for their hot star counterparts, further monitoring would be warranted to constrain the long-period giant planet companion rate for misaligned cool star systems. Additional observations are needed to parse the emerging relationship between stellar temperature, eccentricity, and obliquity for wide-separation planets.

6.7 Acknowledgements

We thank Konstantin Batygin and Kassandra Anderson for helpful discussions that have refined this work. We also thank the anonymous referee for their helpful suggestions that have improved the quality of this manuscript. M.R. is supported by the National Science Foundation Graduate Research Fellowship Program under Grant Number DGE-1752134. The data presented herein were obtained at the W. M. Keck Observatory, which is operated as a scientific partnership among the California Institute of Technology, the University of California and the National Aeronautics and Space Administration. The Observatory was made possible by the generous financial support of the W. M. Keck Foundation. This work is supported by Astronomical Big Data Joint Research Center, co-founded by National Astronomical Observatories, Chinese Academy of Sciences and Alibaba Cloud. This research has made use of the Keck Observatory Archive (KOA), which is operated by the W.

M. Keck Observatory and the NASA Exoplanet Science Institute (NExScI), under contract with the National Aeronautics and Space Administration. This research has made use of the NASA Exoplanet Archive, which is operated by the California Institute of Technology, under contract with the National Aeronautics and Space Administration under the Exoplanet Exploration Program.

`numpy` (Oliphant, 2006; Walt et al., 2011; Harris et al., 2020), `matplotlib` (Hunter, 2007), `pandas` (McKinney et al., 2010), `scipy` (Virtanen et al., 2020), `allesfitter` (Günther & Daylan, 2021), `emcee` (Foreman-Mackey et al., 2013)

Keck: I (HIRES), Exoplanet Archive, Extrasolar Planets Encyclopaedia

Chapter 7

Origins of Hot Jupiters from the Stellar Obliquity Distribution

7.1 abstract

The obliquity of a star, or the angle between its spin axis and the average orbit normal of its companion planets, provides a unique constraint on that system's evolutionary history. Unlike the Solar System, where the Sun's equator is nearly aligned with its companion planets, many hot Jupiter systems have been discovered with large spin-orbit misalignments, hosting planets on polar or retrograde orbits. We demonstrate that, in contrast to stars harboring hot Jupiters on circular orbits, those with eccentric companions follow no population-wide obliquity trend with stellar temperature. This finding can be naturally explained through a combination of high-eccentricity migration and tidal damping. Furthermore, we show that the joint obliquity and eccentricity distributions observed today are consistent with the outcomes of high-eccentricity migration, with no strict requirement to invoke the other hot Jupiter formation mechanisms of disk migration or in-situ formation. At a population-wide level, high-eccentricity migration can consistently shape the dynamical evolution of hot Jupiter systems.

7.2 Introduction

The sky-projected obliquities, λ, of over 140 exoplanet-hosting stars have been determined to date by observing the Rossiter-McLaughlin effect (Rossiter, 1924; McLaughlin, 1924) in which a transiting planet blocks out different components of a rotating star's light as it passes across the stellar profile. These measurements have revealed a diversity of projected angles between the stellar spin axis and the orbit normal vectors of neighboring planets, with systems spanning the full range of possible configurations from prograde to polar and retrograde orbits.

Because Rossiter-McLaughlin observations require a transiting geometry and at least 10-12

high-resolution in-transit spectra, they are limited to only a subset of the known population of exoplanets and have typically been made for hot Jupiters – giant planets on tight orbits. Several channels have been proposed for hot Jupiter formation (see Dawson & Johnson (2018) for a comprehensive overview), including (1) high-eccentricity migration, in which planets born on wide orbits reach extremely high eccentricities before tidally circularizing to their current orbits (e.g. Wu & Murray, 2003; Fabrycky & Tremaine, 2007; Wu et al., 2007; Nagasawa et al., 2008; Beaugé & Nesvorný, 2012); (2) disk migration, in which planets born on wide orbits migrate inwards within the disk plane (e.g. Goldreich & Tremaine, 1980; Lin & Papaloizou, 1986; Lin et al., 1996); and (3) in-situ formation, in which planets form on similar orbits to those on which they currently lie (e.g. Batygin et al., 2016; Boley et al., 2016). The stellar obliquity distribution may provide compelling evidence to distinguish between hot Jupiter formation mechanisms: hot Jupiters formed through high-eccentricity migration, should commonly attain both high eccentricities and large misalignments early in their evolution. On the other hand, hot Jupiters formed in-situ or through disk migration should typically be aligned in the absence of disk- or star-tilting perturbers, with no requirement to reach high eccentricities or inclinations at any point in their evolution.

The primary observational result from previous studies of exoplanet host star obliquities is that hot stars hosting hot Jupiters span a wider range of obliquities than their cool star counterparts (Winn et al., 2010; Schlaufman, 2010). The transition point occurs at the Kraft break ($T_{\rm eff} \sim 6100$ K), a rotational discontinuity above which stars rotate much more quickly and lack thick convective envelopes (Kraft, 1967). The observed discontinuity in obliquities at the Kraft break is commonly attributed to differences in the tidal realignment timescale for stars above and below the Kraft break (Winn et al., 2010; Albrecht et al., 2012), and alternative explanations invoking magnetic braking (e.g. Dawson, 2014; Spalding & Batygin, 2015), internal gravity waves (Rogers et al., 2012), and differences in the external

companion rate (Wang et al., 2022) have also been proposed. To date, however, this trend has only been demonstrated for the full population of giant-exoplanet-hosting stars with measured obliquities, without delineating the role of the companions' orbital eccentricity, e.

In this work, we show that the population of exoplanets on eccentric orbits reveals no evidence for this well-established transition at the Kraft break. While we recover this discontinuity for the population of exoplanets on circular orbits, it is not present for exoplanets on eccentric orbits. This discrepancy supports high-eccentricity migration as a key hot Jupiter formation mechanism, where the final obliquities of hot Jupiter systems are shaped by tidal dissipation.

7.3 Population-Wide Obliquity Analysis

We compared two populations: "eccentric" planets with $e \geq 0.1$ – ranging from $e = 0.1$ to $e = 0.93$ in our sample – as well as "circular" ($e = 0$) planets with a reported eccentricity of exactly zero. As in Wang et al. (2021), the cutoff for "eccentric" planets was set to $e = 0.1$ to remove systems with the most modest eccentricities, which may be more analogous to the "circular" population due to the Lucy-Sweeney bias (Lucy & Sweeney, 1971; Zakamska et al., 2011).

The samples were drawn from the set of planets included within both the NASA Exoplanet Archive "Confirmed Planets" list and the TEPCat catalogue (Southworth, 2011) of sky-projected spin-orbit angles, both downloaded on 10/20/2021. For planets with multiple spin-orbit angle measurements available through TEPCat, only the most recent measurement was included within this analysis, with the exception of systems for which previous observations were much more precise. We included only planets with pericenter distances $q < 0.1$ au, which can be tidally circularized on relatively short timescales necessary for high-eccentricity migration. The distribution of obliquities for both populations as a function of

stellar temperature is provided in the top panel of Figure 7.1.

A set of cumulative sums comparing the eccentric and circular obliquity distributions, where misalignments are accumulated as a function of stellar temperature, reveals that the population-wide change in obliquities at the Kraft break is present for stars hosting planets on circular orbits but does not extend to systems with eccentric orbits (Figure 7.1). The $e \geq 0.1$ cumulative sum is shown in purple in Figure 7.1, where the absolute value of the projected spin-orbit angle, $|\lambda|$, is used to uniformly display the deviation of each system from perfect alignment. To compare this result with the $e = 0$ population, we divided the circular sample into populations with host star $T_{\rm eff}$ below and above the Kraft break. Random planets were selected from each set to match the number of planets on either side of the Kraft break in the $e \geq 0.1$ population, with 5000 iterations to sample the full parameter space. Then, the $e = 0$ planet samples were ordered by $T_{\rm eff}$ and cumulatively summed.

We recovered the previously reported trend in obliquity as a function of stellar temperature, confirming that the trend is stronger when excluding lower-mass ($M < 0.3 M_J$) planets from the sample. However, this relation holds *only for the population of planets on circular orbits.* At the Kraft break, the eccentric cumulative sum is a 6.5σ outlier from the circular distribution with all planets included (middle panel of Figure 7.1) and an 8.7σ outlier from the circular giant planets distribution (bottom panel of Figure 7.1).

To determine the likelihood that the $e = 0$ trend with stellar temperature is absent for $e \geq 0.1$ planets, we compared two models: a linear model and a running median of the $e = 0$ population, which represents the null result. The fit of each model to the cumulatively summed data was evaluated using the Bayesian Information Criterion (BIC; Schwarz et al., 1978) and Akaike Information Criterion (AIC; Akaike, 1973) metrics.

The BIC is given by

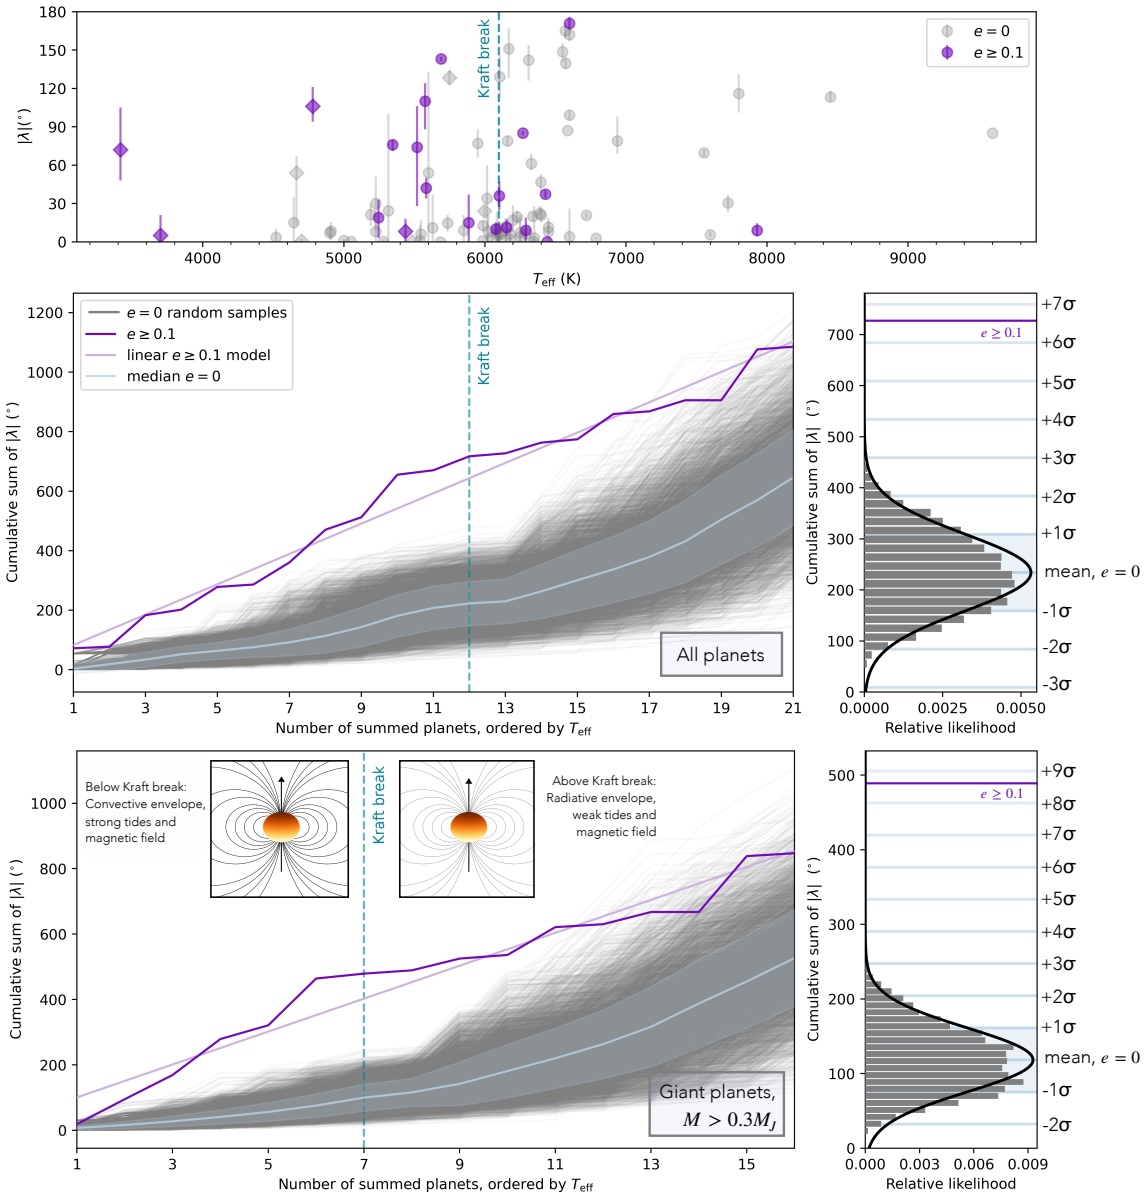

Figure 7.1: Comparison of the obliquity distributions for stars hosting exoplanets on circular vs. eccentric orbits. Top: Full sample of spin-orbit angles included in this study. Planets with $M < 0.3 M_J$ are shown with diamond markers. The data behind this panel is available together with all other planet parameters used in this work, drawn from archival studies. Middle and bottom: Cumulative sums of $|\lambda|$ for eccentric exoplanets, compared with 5000 randomly sampled sets of circular exoplanets (sampled without replacement). Histograms on the right provide vertical cuts through the sums at the Kraft break. In each panel, a linear model fitting the $e \geq 0.1$ cumulative sum is shown in light purple, while the running median of the $e = 0$ population is provided in light blue together with the shaded region within 1σ of the median.

$$\mathrm{BIC} = k \ln N - 2 \ln \mathcal{L}, \tag{7.1}$$

where k is the total number of parameters estimated by the model, N is the number of planets in the cumulative sum, and $\mathcal{L}$ is the likelihood function. For the likelihood function, we used a reduced χ^2 metric comparing the cumulative sum obtained from the data (y_{data}) with the corresponding values for each model (y_{model}):

$$\mathcal{L} = \frac{1}{N} \sum (y_{\mathrm{data}} - y_{\mathrm{model}})^2. \tag{7.2}$$

We adopted the corrected AIC metric (AICc), which includes an adjustment for small sample sizes with $N < 40$ (Hurvich & Tsai, 1989).

$$\mathrm{AIC} = 2k - 2N \ln \mathcal{L} \tag{7.3}$$

$$\mathrm{AICc} = \mathrm{AIC} + \frac{2k(k+1)}{N-k-1} \tag{7.4}$$

Both the BIC and the AICc strongly favor a linear model over the median $e = 0$ cumulative sum (see Figure 1), with ΔBIC$= 82$ (52) and ΔAICc$= 83$ (53) for the all- (giant-) planets fit. Based on the AICc for each model, the null result, which increases in gradient at the Kraft break, is 8×10^{-19} (3×10^{-12}) times as likely as the linear model to minimize the information loss in the all- (giant-) planet fit. The eccentric and circular populations appear to be distinct.

A direct comparison of each examined subpopulation is provided in Figure 7.2. The full sample is segmented by temperature (above/below the Kraft break) and by eccentricity, where we consider the circular and eccentric populations separately. Systems with hot host

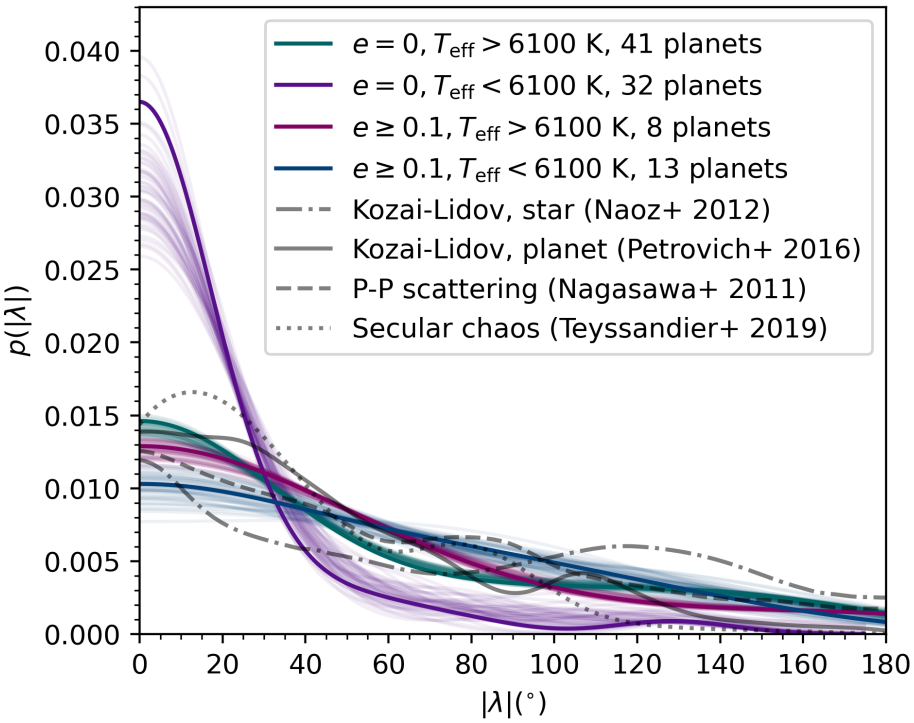

Figure 7.2: Distribution of $|\lambda|$ values for each population, segmented by eccentricity and host star temperature. The central $|\lambda|$ distribution for each set of planets is shown in bold. To quantify uncertainties, we iteratively drew from a Gaussian distribution around each measured $|\lambda|$ value and recalculated the probability distribution based on these new draws, plotting the resulting distributions with lower opacity. Fifty iterations are displayed per set of planets. The theoretical projected obliquity distributions are overlaid for the Kozai-Lidov effect with a stellar (Naoz et al., 2012) or planetary (Naoz et al., 2011; Petrovich & Tremaine, 2016) companion, secular chaos (Teyssandier et al., 2019), and planet-planet scattering (Nagasawa & Ida, 2011).

stars or eccentric planets span a wide range of stellar obliquities. By contrast, the circular distribution around cool stars, in purple, is heavily weighted towards aligned systems ($|\lambda| \sim 0^\circ$).

7.4 Tidal Damping in Hot Jupiter Systems

We considered the effects of tidal damping to investigate the potential origins of the different stellar obliquity distributions in Figure 7.2. Once a system becomes misaligned, interactions with the host star continually act to damp that misalignment. All bound planets are affected

by interactions with their host stars. However, the extent to which those interactions alter the planet's orbit varies strongly with the system parameters. Stars below the Kraft break have convective envelopes and efficient magnetic dynamos, whereas stars above the Kraft break have radiative envelopes and weaker magnetic braking (Dawson, 2014). As a result, cool stars can much more efficiently damp out tidal oscillations and realign their companions (Winn et al., 2010).

In the classical equilibrium tide theory, the tidal realignment timescale for cool stars is given by

$$\tau_\lambda = k(M_{\rm cz}/M_p)(a/R_*)^6(1-e^2)^{9/2}(1+3e^2+(3/8)e^4)^{-1}, \tag{7.5}$$

where M_{cz} is the mass of the convective envelope, M_p is the planet mass, R_* is the stellar radius, a is the planet semimajor axis, and k is a constant (Winn et al., 2010). While this model is a simplified heuristic of a more nuanced tidal theory (Ogilvie, 2014; Lin & Ogilvie, 2017), rather than an exact relation, it captures the global properties of the system's behavior.

Figure 7.3 shows the cumulative sum of spin-orbit angles as a function of τ_λ for planets orbiting cool stars ($T_{\rm eff} < 6100$ K), with systems segmented by damping timescale (rather than eccentricity as in Figure 7.1). We use $k = 10^3$ as a calibrator, such that systems with ages $\tau_\lambda < 10^{10}$ years are typically aligned. The measured $|\lambda|$ values were cumulatively summed for the 20 systems with the longest obliquity damping timescales ($\tau_\lambda > 10^{11}$ years) and compared to random draws without replacement from the population of planets with shorter timescales $\tau_\lambda < 10^{11}$ years.

As in Figure 7.1, the median of the randomly sampled distribution is shown in light blue together with the region within 1σ of the median. Systems with $N_{\rm pl} > 2$, which, within our sample, are "peas-in-a-pod" systems (Millholland et al., 2017; Weiss et al., 2018), were excluded from Figure 7.3. To determine the mass of the convective envelope $M_{\rm cz}$,

we applied a previously calculated model (Pinsonneault et al., 2001) relating the stellar $T_{\rm eff}$ to $M_{\rm cz}$. All other parameters were drawn from the NASA Exoplanet Archive and supplemented with values from the Extrasolar Planets Encyclopaedia.

The cumulative sum in Figure 7.3 reveals that, at a 7.6σ confidence level, planets with longer tidal realignment timescales tend to be observed with larger orbital misalignments. This supports the high-eccentricity migration framework in which hot Jupiter systems around cool stars often begin with large misalignments that are damped over time. Recent work has similarly found that high obliquities of giant exoplanet host stars are almost exclusively associated with wide-separation planets or hot stars (Wang et al., 2021), which have long tidal realignment timescales.

High-eccentricity migration is initialized by N-body interactions in systems with three or more constituent masses. Dynamical interactions push one planet onto an extremely eccentric orbit, which is gradually recircularized through tidal interactions with the host star. These interactions can simultaneously account for the elevated eccentricities and spin-orbit angles of $e \geq 0.1$ exoplanets orbiting stars both above and below the Kraft break. They can also produce orbits with large initial eccentricities and misalignments that have subsequently tidally circularized, and that have, in some cases, realigned. N-body mechanisms that are capable of exciting high eccentricities and large spin-orbit misalignments include secular chaos (Wu & Lithwick, 2011; Hamers et al., 2017; Teyssandier et al., 2019), Kozai-Lidov interactions (Wu & Murray, 2003; Petrovich, 2015; Anderson et al., 2016; Vick et al., 2019), and planet-planet scattering (Rasio & Ford, 1996; Beaugé & Nesvorný, 2012).

A combination of these processes, together with differential tidal dissipation in hot and cool star systems, can account for the currently observed $|\lambda|$ distributions in Figure 7.2. In Figure 7.2, theoretical $|\lambda|$ distributions produced by each of these mechanisms are provided alongside the observed distributions. We propose that all four observed distributions in Figure 7.2 are consistent with an origin from the same set of high-eccentricity formation

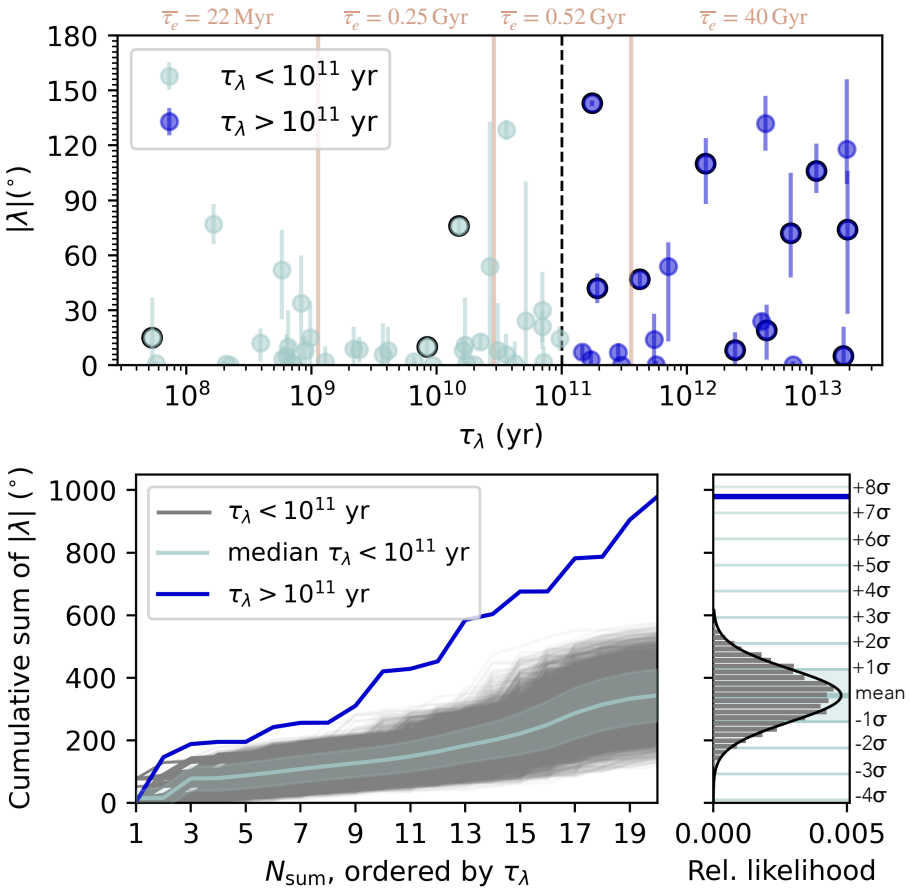

Figure 7.3: Demonstration that systems with long obliquity damping timescales around cool stars are significantly more misaligned than those with shorter damping timescales. Top: Tidal realignment timescales for cool ($T_{\rm eff} < 6100$ K) star systems in our sample. Planets with $e \geq 0.1$ are outlined with a black border. The mean eccentricity damping timescale $\overline{\tau_e}$ for each of the four evenly-sized, 15-planet bins is provided along the top of the panel. Bottom: Stars with the 20 longest tidal realignment timescales ($\tau_\lambda > 10^{11}$ years) as compared with random draws without replacement from the population of systems with shorter tidal timescales.

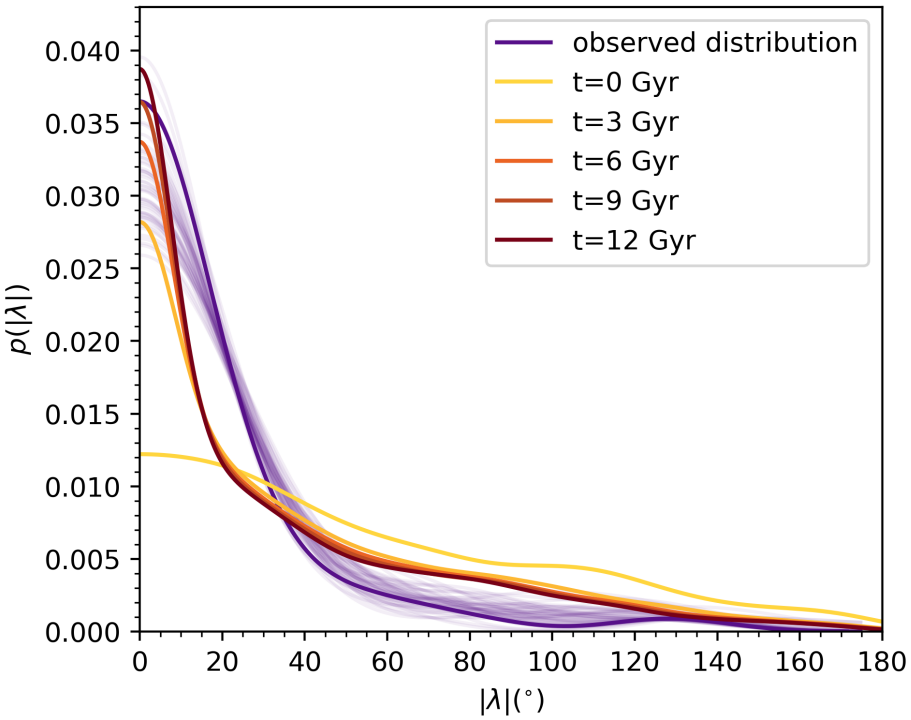

Figure 7.4: Damping evolution for a sample primordial $|\lambda|$ distribution. The initial obliquity model is comprised of 20% stellar Kozai-Lidov systems, 40% planet Kozai-Lidov systems, 10% secular chaos systems, and 30% planet-planet scattering systems. Planets are initialized with the same distribution of semimajor axis, planet mass, age, and stellar $T_{\rm eff}$ as the current set of cool stars with $e = 0$ planets and measured λ.

channels, and that differences in these distributions that are observed today are the natural consequence of obliquity damping.

To demonstrate the effects of tidal damping, we first focused on the obliquity evolution of $e = 0$ planets orbiting cool stars. We applied Equation 7.5 to evolve a mixture model in which 20%, 40%, 10%, and 30% of planets obtain their obliquities through stellar Kozai-Lidov, planet Kozai-Lidov, secular chaos, and planet-planet scattering, respectively, using the starting distributions in Figure 7.2. This distribution is consistent with the high frequency of distant giant perturbers (Ngo et al., 2015; Bryan et al., 2016) that have been proposed to excite the inclinations of their shorter-period companions (Wang et al., 2022).

We fit a kernel density estimation (KDE) to the distributions of host star $T_{\rm eff}$, age, M_p, and a/R_* for the 33 $e = 0$ planets in our sample that orbit stars below the Kraft break, then drew random values from each of these smoothed distributions to produce a set of 10,000 simulated systems. All systems were initialized with $e = 0$. We assumed a linear damping rate, and we set $k = 10^3$ in in accordance with Figure 7.3.

We ultimately found that the distribution evolves along the pathway shown in Figure 7.4 as a result of tidal damping. The theoretical KDE at $t = 12$ Gyr shows a peak at low obliquities analogous to that of the observed distribution. Minor discrepancies at moderate $|\lambda|$ values may result from the small number of misaligned planets (five planets with $|\lambda| > 35°$) that shape the tail of the smoothed, observed sample, without necessarily indicating a true disagreement between the two distributions.

The sky-projected obliquities were directly evolved under the implicit assumption that $\lambda = \psi$. At low $|\lambda|$ values, $|\lambda|$ acts as a lower limit on $|\psi|$, whereas $|\psi|$ is more likely to be close to $|\lambda|$ for larger $|\lambda|$ values (Fabrycky & Winn, 2009). This bias indicates that some fraction of systems observed with low sky-projected obliquities should actually have larger 3D obliquities. Because the probability density peaks at $\lambda = \psi$ even for low

measured λ values (see Figure 3 of Fabrycky & Winn (2009)), the distribution should not change dramatically if the true distribution $p(\psi)$ was evolved rather than the sky-projected distribution $p(\lambda)$.

Our proof-of-concept shows that tidal damping can reproduce the current distribution of observed $e = 0$ cool star obliquities based on an initial model comprised of secular mechanisms, without the requirement of invoking disk migration or in-situ formation. We emphasize that we do not rule out contributions from these mechanisms, but, rather, we show that they are not stricly required to account for the stellar obliquity distribution. Minor adjustments to the weighting of secular processes can reproduce the circular hot star distribution and the two eccentric distributions in Figure 7.2, each of which has been relatively unaffected by damping.

7.5 The Role of Orbital Eccentricity

The eccentricities of misaligned systems provide an independent test of high-eccentricity migration. While our analysis up to this point has focused on the stellar obliquity damping timescale τ_λ, we can also consider the timescale, $\tau_e \sim e/(de/dt)$, for eccentricity evolution driven by tidal dissipation within the planet. Under the effects of tidal dissipation, the evolution of a planet's eccentricity is given by

$$\frac{de}{dt} = \frac{dE}{dt}\frac{a(1-e^2)}{GMme}\,, \tag{7.6}$$

where M is the host star mass, m is the mass of the planet, and E is the orbital energy of the planet.

The rate of energy dissipation for a synchronously rotating planet is

$$\frac{dE}{dt} = \frac{21 k_2 G M^2 \Omega r^5}{2Qa^6} \zeta(e). \tag{7.7}$$

Here, Q is the planet's effective tidal dissipation parameter, k_2 is the planet's Love number, and Ω is the pseudosynchronous rotation rate, given by

$$\Omega = \frac{1 + \frac{15}{2}e^2 + \frac{45}{8}e^4 + \frac{5}{16}e^6}{(1 + 3e^2 + \frac{3}{8}e^4)(1 - e^2)^{3/2}} n \tag{7.8}$$

where n is the mean motion of the planet's orbit. We set $Q = 10^5$ and $k_2 = 0.3$. The corrective factor $\zeta(e)$, derived in Wisdom (2008), is defined as

$$\zeta(e) = \frac{2}{7}\Big[\frac{f_0(e)}{\beta^{15}} - \frac{2 f_1(e)}{\beta^{12}} + \frac{f_2(e)}{\beta^9}\Big], \tag{7.9}$$

where

$$f_0(e) = 1 + \frac{31}{2}e^2 + \frac{255}{8}e^4 + \frac{185}{16}e^6 + \frac{25}{64}e^8 \tag{7.10}$$

$$f_1(e) = 1 + \frac{15}{2}e^2 + \frac{45}{8}e^4 + \frac{5}{16}e^6 \tag{7.11}$$

$$f_2(e) = 1 + 3e^2 + \frac{3}{8}e^4 \tag{7.12}$$

$$\beta = \sqrt{(1 - e^2)}. \tag{7.13}$$

Separating the 60 planets in Figure 7.3 into four evenly sized bins of 15 planets each, ordered by τ_λ, we demonstrate that systems with shorter τ_λ also preferentially have shorter

τ_e such that their eccentricities and obliquities should be jointly damped over the lifetime of the system. τ_e is typically shorter than τ_λ such that hot Jupiters should often circularize before realigning. In contrast, systems with the longest τ_λ, which include most of the $e \geq 0.1$ population, have τ_e values exceeding the age of the system. Both timescales, which are contributed by two independent processes – damping within the planet and damping within the star – are therefore consistent with the observed distributions. If all systems formed through high-eccentricity migration, the observed λ and e distributions would look as they do today.

7.6 Implications for Hot Jupiter Formation Theory

Our results provide two key constraints on the obliquity distribution of Rossiter-McLaughlin targets, which are primarily hot Jupiter host stars. The first constraint, which is a variation on previous findings (Winn et al., 2010; Schlaufman, 2010), is the observation that stars hosting circular hot Jupiters span a wider range of obliquities above the Kraft break than at lower temperatures. The second is the absence of this pattern in the eccentric sample, where obliquities are consistent with no change at the Kraft break. Together, these observations demonstrate that the population of hot Jupiters is consistent with formation through high-eccentricity migration and suggest that dissipative mechanisms are vital for shaping the obliquity distribution of hot Jupiter host stars.

The absence of a significant change in obliquities at the Kraft break for eccentric systems indicates that dissipative mechanisms have not had time to sculpt the eccentric population in the same way that they have shaped the circular population. In the framework of Kozai capture (Naoz et al., 2012), eccentric planets are either experiencing ongoing eccentric Kozai-Lidov oscillations or they have had these oscillations suppressed by apsidal precession in the past. Our sample includes well-characterized systems such as that of HD

80606, in which the transiting planet's high obliquity and eccentricity are cleanly recovered through Kozai migration (Wu & Murray, 2003).

A subset of eccentric planets that have two or more planetary companions may be undergoing secular chaos. Secular chaos transfers angular momentum outwards due to the overlap of resonances in a multiplanet system, elevating the orbital inclination and eccentricity of the innermost planet. In either excitation framework, tidal dissipation, which acts differentially in hot and cool star systems, has not had time to globally alter the $|\lambda|$ distribution of the eccentric population.

For hot Jupiter systems with circular orbits, tidal dissipation has played a more important role in shaping the currently observed stellar obliquity distribution. These systems are consistent with N-body interactions that were suppressed early on such that the companion orbits were able to fully tidally circularize and, in some cases, realign. Planets orbiting cool stars realign quickly, while those orbiting hot stars have much longer tidal realignment timescales and remain closer to their primordial spin-orbit angles.

Systems with both large misalignments and high eccentricities may instead be produced by a combination of primordial disk misalignments and planet-planet or planet-disk interactions (Duffell & Chiang, 2015; Anderson & Lai, 2017; Frelikh et al., 2019; Anderson et al., 2020; Debras et al., 2021). Our results do not rule out these alternative scenarios, but, rather, they provide a relatively simple framework that is fully consistent with the observed λ and e distributions. Previous work has revealed that long-period ($P > 5$ day) planets orbiting hot stars tentatively demonstrate a trend towards alignment (Rice et al., 2021). Because these planets are exceptionally difficult to realign through tidal interactions, this trend, if confirmed, may suggest that protoplanetary disks are typically aligned and that the misalignments of hot Jupiters are attained through dynamical interactions after the disk has dispersed.

7.7 Conclusions

Our analysis establishes that the observed distribution of hot Jupiter host star obliquities can arise naturally from a combination of high-eccentricity migration and obliquity damping mechanisms. Cool stars hosting circular planets have had the most strongly damped obliquities, while hot stars and hosts of eccentric planets have experienced weaker damping. We predict that, under our proposed framework, the observed difference between the eccentric and $e = 0$ cumulative sums will grow with additional observations. We conclude that the stellar obliquity distribution for hot Jupiter systems is consistent with having been crafted primarily by high-eccentricity migration and tidal damping, with no requirement to appeal to disk migration or in-situ formation at the population level.

7.8 Acknowledgements

M.R. is supported by the National Science Foundation Graduate Research Fellowship Program under Grant Number DGE-1752134. This research has made use of the NASA Exoplanet Archive, which is operated by the California Institute of Technology, under contract with the National Aeronautics and Space Administration under the Exoplanet Exploration Program.

`numpy` (Oliphant, 2006; Walt et al., 2011; Harris et al., 2020), `matplotlib` (Hunter, 2007), `pandas` (McKinney et al., 2010), `scipy` (Virtanen et al., 2020), `emcee` (Foreman-Mackey et al., 2013)

Exoplanet Archive, Extrasolar Planets Encyclopaedia, Open Exoplanet Catalogue

7.9 Appendix: Adopted Parameters

Our full samples of parameters, drawn from archival studies, are provided as supplementary data for Figure 7.1. All stellar and planetary parameters other than stellar multiplicity and λ were drawn directly from the NASA Exoplanet Archive, with ages supplemented by the Extrasolar Planets Encylopaedia. The stellar multiplicity of each system, provided for reference, was determined through cross-matching with the Catalogue of Exoplanets in Binary Star Systems (Schwarz et al., 2016) and the Open Exoplanet Catalogue.

7.10 Appendix: Orbital Period vs. Eccentricity

One alternative possibility is that eccentricity acts as a proxy for a different trend in the dataset. The tidal circularization timescale of a short-period planet scales with semimajor axis as $\tau_{\rm cir} \propto a^{13/2}$ (Murray & Dermott, 1999), meaning that small differences in semimajor axis correspond to dramatically different tidal circularization timescales. Planets on eccentric orbits, by definition, have not completed the tidal circularization process. This means that they may also tend to have larger semimajor axes, or, equivalently, longer orbital periods (P) as compared with $e = 0$ planets.

To address this possibility, we carried out the same analysis as a function of orbital period (comparing the $P > 5$ day and $P < 5$ day populations) and as a function of orbital separation (comparing the $a/R_* > 12$ and $a/R_* < 12$ populations), with results shown in Figure 7.5. If the observed effect is predominantly due to a correlation between obliquity and orbital period (orbital separation), rather than eccentricity, the population should show a stronger increase in misalignments with increasing P (a/R_*) than e.

In both cases, the significance of our result was substantially weaker than when dividing the sample by eccentricity (6.5σ), with only a 3.3σ signal when segmenting by P and a 4.9σ

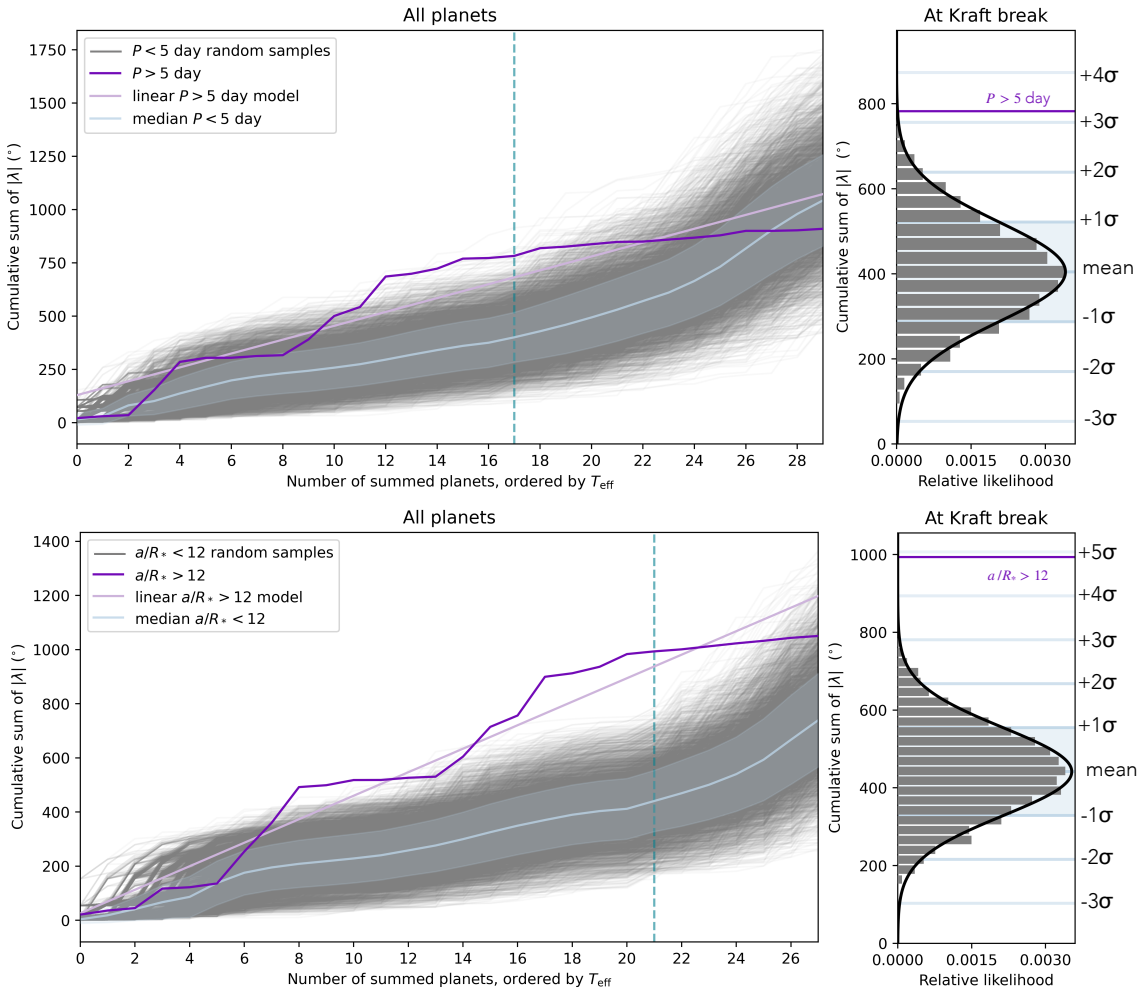

Figure 7.5: Cumulative sums as a function of orbital period P (top) and orbital separation a/R_* (bottom), for comparison with the eccentricity cumulative sum in Figure 7.1.

signal when segmenting by a/R_*. Figure 7.5 also shows substantial structure below the Kraft break in the $P > 5$ day and $a/R_* > 12$ sums. This suggests that divisions by P or a/R_* may produce a more heterogeneous population than that produced by our eccentricity cut, where planets smoothly follow a relatively consistent upward trend in $\sum |\lambda|$.

Finally, both panels of Figure 7.5 reveal a nearly flat trend in $\sum |\lambda|$ above the Kraft break. This demonstrates that relatively long-period systems around hot stars have typically been observed to be aligned, as pointed out in Rice et al. (2021).

Chapter 8

Conclusions

8.1 Recent Advancements and Future Directions

The topics spanned by this dissertation jointly considered the various components of planetary systems, focusing on their relationship and how their interactions can build towards a holistic understanding of planetary systems. Here, we summarize some of the relevant advancements that have been made since the projects described here were conducted, as well as expected advancements in the immediate future.

8.1.1 Progress in the Search for Planet Nine

Since the initial proposition of a ninth distant solar system object, studies have been released to both contest the statistical significance of the apsidal alignment that led to Planet Nine's prediction (Napier et al., 2021) and to further argue for the robustness of this detection (Brown & Batygin, 2019). Part of the difficulty in ascertaining this significance is the nonuniformity of systematics and survey strategies that were used to discover each of the known long-term stable eTNOs.

All-sky surveys reaching the depths required to recover distant eTNOs will be necessary to more definitively demonstrate whether the observed apsidal alignment is a physical effect, rather than an observational bias. A large-scale application of the methods described in Chapter 3 will provide useful insights to distinguish between these two possibilities. Observations with the upcoming Vera Rubin Observatory will also improve minor planet detection limits across the southern sky, which will either strengthen or weaken the disputed apsidal alignment.

The predicted orbital parameters of Planet Nine have recently been more tightly constrained to pinpoint the expected region of the sky where the planet most likely lies, as well as its anticipated magnitude (Brown & Batygin, 2021). The new planet aligns well with the region where the galactic plane and the ecliptic plane meet, which may help to explain why,

if it exists, the object has not yet been found: the increased density of bright sources in this region of the sky leads to signal confusion, making dim objects more difficult to find amidst the many neighboring sources.

Over the past several months, the TESS spacecraft observed this segment of the sky as part of its extended mission. As a result, the algorithms described in Chapter 3 can be directly applied to search this high-probability area of parameter space for the presence of Planet Nine. This is an exciting future direction that is actively in progress at the time of writing.

A candidate planet with $a = 225 \pm 15$ au and $M = 3 - 5M_{\oplus}$ was recently identified based on an archival search for Planet Nine applying the Infrared Astronomical Satellite (IRAS) dataset (Rowan-Robinson, 2021). While the candidate planet does not share the properties required for the predicted Planet Nine and would not produce the observed apsidal alignment of eTNOs, it provides another compelling potential candidate to also search for within the TESS dataset and within other archival and upcoming surveys.

Planet Nine has not yet been found. The parameter space for the planet has, however, been dramatically narrowed by recent work, which ruled out 56% of the predicted parameter space for the planet (Brown & Batygin, 2022). Searches for Planet Nine using data from upcoming surveys, including that of both TESS and the Vera Rubin Observatory, will further reduce the remaining parameter space for this planet such that it should be discovered within the coming years if it exists in its proposed form.

8.1.2 The Discovery of the Second Interstellar Object

Two years after the discovery of 1I/'Oumuamua, the second interstellar object, 2I/Borisov, was discovered by Crimean amateur astronomer Gennady Borisov. Unlike 'Oumuamua, 2I/Borisov did demonstrate clear outgassing, with an observed coma and a reddish color consistent with the typical range of colors spanned by solar system comets (Guzik et al.,

2020). Other than its hyperbolic orbit, 2I/Borisov was not observed with properties that made it exceptionally distinct relative to other solar system comets.

While the unexpected properties of 'Oumuamua remain a mystery, observations of 2I/Borisov demonstrate that at least some interstellar objects likely are ejected cometary material from the outer regions of planetary systems, as in the scenario described within Chapter 4. Furthermore, the discovery of a second interstellar object demonstrates that these bodies comprise a population. This sets an exciting precedent for potential future discoveries by amateur astronomers and by upcoming large-scale surveys, including the Large Survey of Space and Time that will be conducted with the Vera Rubin Observatory.

A nonlinear shift-stacking survey with the TESS dataset will provide tighter constraints on the occurrence rates of interstellar objects, as well as an exciting avenue for future potential detections. The Vera Rubin Observatory will also likely discover one or more interstellar objects over the coming years. Future studies of interstellar objects will provide a more detailed understanding of the distribution of interstellar object properties, offering direct insights into the building blocks of extrasolar systems.

8.1.3 Composition as a Probe of Planet Formation Processes

The host star crafts the environment around which its neighboring planets are formed, setting the initial conditions that influence what planets are born in that system. As a result, correlations between stellar parameters and planetary system architectures can provide useful clues to unveil the key formation pathways for planets. This can be accomplished through direct comparisons between exoplanet demographics and stellar properties, including both global parameters ($T_{\rm eff}$, $v \sin i$, $\log g$, [Fe/H]) and abundances.

The detailed characterization of atmospheric properties constitutes a related major frontier for exoplanet studies in the immediate future. As new constraints are obtained for the

compositions of exoplanet atmospheres, it will become possible to demonstrate whether there are clear patterns in atmospheric composition as a function of stellar type, irradiation, host star elemental abundances, or additional environmental factors.

Atmospheric compositions can provide important clues regarding how exoplanets formed. For example, atmospheres should typically be rich in H and He, with a similar composition to that of the central star, if their gas directly accretes from the host protoplanetary disk. Alternatively, planets could instead host almost exclusively hydrogen envelopes produced by outgassing (Rogers & Seager, 2010). Previous work has, surprisingly, found no observed correlation between host star elemental abundances and the bulk metallicity of their neighboring planets (Teske et al., 2019). Additional atmospheric constraints, combined with constraints on the bulk compositions of both host stars and exoplanets, will enable future studies examining these trends in greater detail.

8.1.4 Population Studies of Stellar Obliquities

The catalogue of obliquity measurements remains too small to clearly delineate correlations between spin-orbit misalignments and eccentricity, orbital separation, planetary multiplicity, and stellar multiplicity for different exoplanet system types. The SOLES survey and analogous efforts will help to extend obliquity measurements into sparsely probed regions of parameter space.

While most obliquity constraints so far have been obtained in only 1D or 2D, measurements from multiple techniques can be combined to extract 3D obliquities in some systems. Recently, the population of stars with 3D obliquity constraints was analyzed to demonstrate that planets on polar orbits, passing sideways over the stellar equator, are unexpectedly common (Albrecht et al., 2021). If this trend is confirmed, it will provide a crucial constraint for planetary evolution models.

Polar protoplanetary disks can leave behind polar planets similar to those suggested by the stellar obliquity distribution (e.g. Childs & Martin, 2021). Protoplanetary disks around single stars can be tilted to a polar configuration due to resonance sweeping and instability in the later stages of planet formation (Petrovich et al., 2020). Furthermore, hydrodynamic simulation results have demonstrated that polar protoplanetary disks may be a common outcome of dynamical evolution around eccentric binary star systems (Martin & Lubow, 2017). A polar protoplanetary disk was recently discovered around a binary star system (Kennedy et al., 2019), demonstrating that at least some polar planets likely originate from polar protoplanetary disks.

Alternatively, planets may instead reach perpendicular orbits due to dynamical excitation once the protoplanetary disk has dispersed. The Kozai-Lidov mechanism may produce a pileup of roughly polar planets consistent with the observed distribution (Anderson et al., 2016). Tidal dissipation can produce an excess of perpendicular planets around cool stars with convective exteriors, as well (Lai, 2012; Rogers & Lin, 2013). Additional measurements demonstrating what types of systems typically contain perpendicular orbits – for example, whether this effect pervades consistently across both Jovian and Neptune-sized planet populations – will help to unveil the origins of polar planets.

8.2 Final Thoughts

Research is an ongoing process, and parallel paths of progress simultaneously pave the way towards a cohesive resolution. This dissertation has weaved together multiple lines of evidence to explore how the interactions between planets, stars, and minor planets can further our understanding of planetary systems. We hope that it will provide a rich set of new ideas and approaches as the field converges towards a unified set of planetary system evolution models.

Bibliography

Abell, P. A., Burke, D. L., Hamuy, M., et al. 2009, LSST Science Book, Version 2.0, Tech. rep.

Abolfathi, B., Aguado, D., Aguilar, G., et al. 2018, The Astrophysical Journal Supplement Series, 235, 42

Adams, E., López-Morales, M., Elliot, J., et al. 2011, The Astrophysical Journal, 741, 102

Akaike, H. 1973, Biometrika, 60, 255

Albrecht, S., Winn, J. N., Johnson, J. A., et al. 2012, The Astrophysical Journal, 757, 18

Albrecht, S. H., Marcussen, M. L., Winn, J. N., Dawson, R. I., & Knudstrup, E. 2021, The Astrophysical Journal Letters, 916, L1

Almenara, J., Damiani, C., Bouchy, F., et al. 2015, Astronomy & Astrophysics, 575, A71

Alvarez-Candal, A., Ortiz, J., Morales, N., et al. 2014, Astronomy & Astrophysics, 571, A48

Anderson, K. R., & Lai, D. 2017, Monthly Notices of the Royal Astronomical Society, 472, 3692

—. 2018, Monthly Notices of the Royal Astronomical Society, 480, 1402

Anderson, K. R., Lai, D., & Pu, B. 2020, Monthly Notices of the Royal Astronomical Society, 491, 1369

Anderson, K. R., Storch, N. I., & Lai, D. 2016, Monthly Notices of the Royal Astronomical Society, 456, 3671

Anderson, K. R., Winn, J. N., & Penev, K. 2021, The Astrophysical Journal, 914, 56

Andrews, S. M., Huang, J., Pérez, L. M., et al. 2018, The Astrophysical Journal Letters, 869, L41

Ansdell, M., Williams, J. P., van der Marel, N., et al. 2016, The Astrophysical Journal, 828, 46

Arimatsu, K., Tsumura, K., Usui, F., et al. 2019, Nature Astronomy, 1

Armitage, P. J., Bae, J., Benisty, M., et al. 2020, arXiv preprint arXiv:2009.04345

Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, Astronomy & Astrophysics, 558, A33

Bahcall, J. N., & Bahcall, S. 1985, Nature, 316, 706

Baillié, K., Colwell, J. E., Lissauer, J. J., Esposito, L. W., & Sremčević, M. 2011, Icarus, 216, 292

Bakos, G., Noyes, R., Kovács, G., et al. 2004, Publications of the Astronomical Society of the Pacific, 116, 266

Banda-Huarca, M., Camargo, J., Desmars, J., et al. 2019, The Astronomical Journal, 157, 120

Bannister, M. T., Bhandare, A., A, D. P., et al. 2019, Nature Astronomy, 3, 594

Barclay, T., Quintana, E. V., Raymond, S. N., & Penny, M. T. 2017, The Astrophysical Journal, 841, 86

Barry, M. A. T., Gault, D., Pavlov, H., et al. 2015, Publications of the Astronomical Society of Australia, 32, e031

Barucci, M., Cruikshank, D., Mottola, S., & Lazzarin, M. 2002, Asteroids III, 273

Batygin, K. 2012, Nature, 491, 418

Batygin, K., Adams, F. C., Brown, M. E., & Becker, J. C. 2019, Physics Reports, 805, 1

Batygin, K., Bodenheimer, P. H., & Laughlin, G. P. 2016, The Astrophysical Journal, 829, 114

Batygin, K., & Brown, M. E. 2016, The Astronomical Journal, 151, 22

—. 2021, The Astrophysical Journal Letters, 910, L20

Batygin, K., & Morbidelli, A. 2017, The Astronomical Journal, 154, 229

Baxter, E., Jain, B., Blake, C., et al. 2018, arXiv preprint arXiv:1812.08701

Beaugé, C., & Nesvornỳ, D. 2012, The Astrophysical Journal, 751, 119

Becker, A. C., Arraki, K., Kaib, N., et al. 2008, The Astrophysical Journal Letters, 682, L53

Becker, J., Khain, T., Hamilton, S. J., et al. 2018, The Astronomical Journal, 156, 81

Becker, J. C., Adams, F. C., Khain, T., Hamilton, S. J., & Gerdes, D. 2017, The Astronomical Journal, 154, 61

Behmard, A., Petigura, E. A., & Howard, A. W. 2019, The Astrophysical Journal, 876, 68

Bellini, A., Anderson, J., & Bedin, L. 2011, Publications of the Astronomical Society of the Pacific, 123, 622

Benedetti-Rossi, G., Sicardy, B., Buie, M. W., et al. 2016, The Astronomical Journal, 152, 156

Benedict, G. F., McArthur, B., Fredrick, L., et al. 2002, The Astronomical Journal, 124, 1695

Benisty, M., Bae, J., Facchini, S., et al. 2021, The Astrophysical Journal Letters, 916, L2

Bennett, D. P., Batista, V., Bond, I., et al. 2014, The Astrophysical Journal, 785, 155

Bernardinelli, P. H., Bernstein, G. M., Sako, M., et al. 2020, The Astrophysical Journal Supplement Series, 247, 32

Bernstein, G., & Khushalani, B. 2000, The Astronomical Journal, 120, 3323

Bernstein, G. M., Trilling, D. E., Allen, R., et al. 2004, The Astronomical Journal, 128, 1364

Bhattacherjee, D., Jyothi, S. A., Bozkurt, I. N., et al. 2018, arXiv preprint arXiv:1809.10897

Boley, A. C., Contreras, A. G., & Gladman, B. 2016, The Astrophysical Journal Letters, 817, L17

Borucki, W., Koch, D., Jenkins, J., et al. 2009, Science, 325, 709

Borucki, W. J., Koch, D., Basri, G., et al. 2010, Science, 327, 977

Bosh, A. S., Person, M., Levine, S., et al. 2015, Icarus, 246, 237

Boss, A. P. 1997, Science, 276, 1836

Brasser, R., Morbidelli, A., Gomes, R., Tsiganis, K., & Levison, H. F. 2009, Astronomy & Astrophysics, 507, 1053

Brewer, J. M., & Fischer, D. A. 2018, The Astrophysical Journal Supplement Series, 237, 38

Brewer, J. M., Fischer, D. A., Basu, S., Valenti, J. A., & Piskunov, N. 2015, The Astrophysical Journal, 805, 126

Brewer, J. M., Fischer, D. A., Valenti, J. A., & Piskunov, N. 2016, The Astrophysical Journal Supplement Series, 225, 32

Brewer, J. M., Wang, S., Fischer, D. A., & Foreman-Mackey, D. 2018, The Astrophysical Journal Letters, 867, L3

Brown, M. E., & Batygin, K. 2019, The Astronomical Journal, 157, 62

—. 2021, The Astronomical Journal, 162, 219

—. 2022, The Astronomical Journal, 163, 102

Brown, M. E., Trujillo, C., & Rabinowitz, D. 2004, The Astrophysical Journal, 617, 645

Bryan, M. L., Knutson, H. A., Howard, A. W., et al. 2016, The Astrophysical Journal, 821, 89

Buder, S., Asplund, M., Duong, L., et al. 2018, Monthly Notices of the Royal Astronomical Society, 478, 4513

Buie, M. W., & Keller, J. M. 2016, The Astronomical Journal, 151, 73

Buie, M. W., Olkin, C. B., Merline, W. J., et al. 2015, The Astronomical Journal, 149, 113

Burkhart, L. D., Ragozzine, D., & Brown, M. E. 2016, The Astronomical Journal, 151, 162

Burns, J. A. 1976, American Journal of Physics, 44, 944

Butler, R. P., Marcy, G. W., Vogt, S. S., et al. 2003, The Astrophysical Journal, 582, 455

Butler, R. P., Marcy, G. W., Williams, E., et al. 1996, Publications of the Astronomical Society of the Pacific, 108, 500

Butler, R. P., Vogt, S. S., Laughlin, G., et al. 2017, The Astronomical Journal, 153, 208

Cameron, A. C., Guenther, E., Smalley, B., et al. 2010, Monthly Notices of the Royal Astronomical Society, 407, 507

Casagrande, L., Ramírez, I., Melendez, J., Bessell, M., & Asplund, M. 2010, Astronomy & Astrophysics, 512, A54

Casey, A. R., Hogg, D. W., Ness, M., et al. 2016, arXiv preprint arXiv:1603.03040

Casey, A. R., Hawkins, K., Hogg, D. W., et al. 2017, The Astrophysical Journal, 840, 59

Chambers, K. C., Magnier, E., Metcalfe, N., et al. 2016, arXiv preprint arXiv:1612.05560

Chiang, E. I., Jordan, A., Millis, R., et al. 2003, The Astronomical Journal, 126, 430

Childs, A. C., & Martin, R. G. 2021, The Astrophysical Journal Letters, 920, L8

Chontos, A., Huber, D., Latham, D. W., et al. 2019, The Astronomical Journal, 157, 192

Christiaens, V., Cantalloube, F., Casassus, S., et al. 2019, The Astrophysical Journal Letters, 877, L33

Clarke, C. J., Tazzari, M., Juhász, A., et al. 2018, The Astrophysical Journal Letters, 866, L6

Collins, K. A., Eastman, J. D., Beatty, T. G., et al. 2014, The Astronomical Journal, 147, 39

Cook, N. V., Ragozzine, D., Granvik, M., & Stephens, D. C. 2016, The Astrophysical Journal, 825, 51

Cowan, N. B., Fuentes, P. A., & Haggard, H. M. 2013, Monthly Notices of the Royal Astronomical Society, 434, 2465

Cowan, N. B., Holder, G., & Kaib, N. A. 2016, The Astrophysical Journal Letters, 822, L2

Cumming, A. 2004, Monthly Notices of the Royal Astronomical Society, 354, 1165

Cumming, A., Marcy, G. W., & Butler, R. P. 1999, The Astrophysical Journal, 526, 890

Damasso, M., Esposito, M., Nascimbeni, V., et al. 2015, Astronomy & Astrophysics, 581, L6

Dark Energy Survey Collaboration. 2005, International Journal of Modern Physics A, 20, 3121

Dawson, R. I. 2014, The Astrophysical Journal Letters, 790, L31

Dawson, R. I., & Johnson, J. A. 2018, Annual Review of Astronomy and Astrophysics, 56, 175

de Bruijne, J., Perryman, M., Lindegren, L., et al. 2005, Gaia–JdB–022

Debras, F., Baruteau, C., & Donati, J.-F. 2021, Monthly Notices of the Royal Astronomical Society, 500, 1621

Devillepoix, H. A., Bland, P. A., Sansom, E. K., et al. 2018, Monthly Notices of the Royal Astronomical Society, 483, 5166

Do, A., Tucker, M. A., & Tonry, J. 2018, The Astrophysical Journal Letters, 855, L10

Ducati, J. 2002, VizieR Online Data Catalog, 2237

Duffell, P. C., & Chiang, E. 2015, The Astrophysical Journal, 812, 94

Eker, Z., Soydugan, F., Soydugan, E., et al. 2015, The Astronomical Journal, 149, 131

El-Nozahi, M., Sanchez-Sinencio, E., & Entesari, K. 2009, IEEE Transactions on Microwave Theory and Techniques, 57, 1054

Elliot, J. L., Dunham, E., Bosh, A., et al. 1989, Icarus, 77, 148

Elliot, J. L., Dunham, E., & Mink, D. 1977, Nature, 267, 328

Elliot, J. L., Person, M., Zuluaga, C., et al. 2010, in Bulletin of the American Astronomical Society, Vol. 42, 991

Endl, M., Cochran, W. D., Kürster, M., et al. 2006, The Astrophysical Journal, 649, 436

Esteves, L. J., De Mooij, E. J., & Jayawardhana, R. 2013, The Astrophysical Journal, 772, 51

Fabrycky, D., & Tremaine, S. 2007, The Astrophysical Journal, 669, 1298

Fabrycky, D. C., & Winn, J. N. 2009, The Astrophysical Journal, 696, 1230

Fernández, Y. R., Jewitt, D., & Ziffer, J. E. 2009, The Astronomical Journal, 138, 240

Fernández, Y. R., Sheppard, S. S., & Jewitt, D. C. 2003, The Astronomical Journal, 126, 1563

Fienga, A., Laskar, J., Manche, H., & Gastineau, M. 2016, Astronomy & Astrophysics, 587, L8

Fienga, A., Manche, H., Laskar, J., Gastineau, M., & Verma, A. 2014, arXiv preprint arXiv:1405.0484

Fischer, D. A., Marcy, G. W., Butler, R. P., Vogt, S. S., & Apps, K. 1999, Publications of the Astronomical Society of the Pacific, 111, 50

Fischer, D. A., & Valenti, J. 2005, The Astrophysical Journal, 622, 1102

Folkner, W., Jacobson, R. A., Park, R., & Williams, J. G. 2016, in AAS/Division for Planetary Sciences Meeting Abstracts, Vol. 48, AAS/Division for Planetary Sciences Meeting Abstracts #48, 120.07

Folkner, W. M., Williams, J. G., Boggs, D. H., Park, R. S., & Kuchynka, P. 2014, Interplanet. Netw. Prog. Rep, 196, 1

Foreman-Mackey, D. 2016, The Journal of Open Source Software, 1, doi:10.21105/joss.00024

Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, Publications of the Astronomical Society of the Pacific, 125, 306

Fraser, W. C., Brown, M. E., Morbidelli, A., Parker, A., & Batygin, K. 2014, The Astrophysical Journal, 782, 100

Freeman, M., Philpott, L., Abe, F., et al. 2015, The Astrophysical Journal, 799, 181

Frelikh, R., Jang, H., Murray-Clay, R. A., & Petrovich, C. 2019, The Astrophysical Journal Letters, 884, L47

Gaia Collaboration, Prusti, T., De Bruijne, J., et al. 2016, Astronomy & Astrophysics, 595, A1

Gaia Collaboration, Brown, A., Vallenari, A., et al. 2018, Astronomy & astrophysics, 616, A1

Gammie, C. F. 2001, The Astrophysical Journal, 553, 174

Giles, H., Bayliss, D., Espinoza, N., et al. 2018, Monthly Notices of the Royal Astronomical Society, 475, 1809

Ginsburg, A., Sipőcz, B. M., Brasseur, C., et al. 2019, The Astronomical Journal, 157, 98

Giorgini, J., Yeomans, D., Chamberlin, A., et al. 1996, in Bulletin of the American Astronomical Society, Vol. 28, 1158

Gladman, B., Kavelaars, J., Nicholson, P. D., Loredo, T. J., & Burns, J. A. 1998, The Astronomical Journal, 116, 2042

Gladman, B., Kavelaars, J., Petit, J.-M., et al. 2001, The Astronomical Journal, 122, 1051

Goldreich, P., & Tremaine, S. 1980, The Astrophysical Journal, 241, 425

Gomes, R., Levison, H. F., Tsiganis, K., & Morbidelli, A. 2005, Nature, 435, 466

González Hernández, J. I., & Bonifacio, P. 2009, Astronomy & Astrophysics, 497, 497

Grav, T., Mainzer, A. K., Bauer, J. M., Masiero, J. R., & Nugent, C. R. 2012, The Astrophysical Journal, 759, 49

Günther, M. N., & Daylan, T. 2021, The Astrophysical Journal Supplement Series, 254, 13

Güttler, C., Blum, J., Zsom, A., Ormel, C., & Dullemond, C. 2009, Geochimica et Cosmochimica Acta Supplement, 73, A482

Guzik, P., Drahus, M., Rusek, K., et al. 2020, Nature Astronomy, 4, 53

Guzmán, V. V., Huang, J., Andrews, S. M., et al. 2018, The Astrophysical Journal Letters, 869, L48

Haffert, S., Bohn, A., de Boer, J., et al. 2019, Nature Astronomy, 3, 749

Hamers, A. S., Antonini, F., Lithwick, Y., Perets, H. B., & Portegies Zwart, S. F. 2017, Monthly Notices of the Royal Astronomical Society, 464, 688

Hands, T. O., Dehnen, W., Gration, A., Stadel, J., & Moore, B. 2019, Monthly Notices of the Royal Astronomical Society, 490, 21

Harris, A. W. 1998, Icarus, 131, 291

Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, Nature, 585, 357

Hart, W., Brown, G. M., Collins, S. M., et al. 2018, in 2018 IEEE Aerospace Conference, IEEE, 1–20

Haswell, C. A. 2010, Transiting Exoplanets (Cambridge University Press)

Hernandez, X. 2019, arXiv preprint arXiv:1901.10605

Hinkel, N. R., Timmes, F., Young, P. A., Pagano, M. D., & Turnbull, M. C. 2014, The Astronomical Journal, 148, 54

Ho, A. Y., Ness, M. K., Hogg, D. W., et al. 2017, The Astrophysical Journal, 836, 5

Holman, M. J., & Payne, M. J. 2016a, The Astronomical Journal, 152, 94

—. 2016b, The Astronomical Journal, 152, 80

Holman, M. J., Payne, M. J., & Pál, A. 2019, Research Notes of the AAS, 3, 160

Holman, M. J., & Wisdom, J. 1993, The Astronomical Journal, 105, 1987

Holman, M. J., Kavelaars, J., Grav, T., et al. 2004, Nature, 430, 865

Holtzman, J. A., Hasselquist, S., Shetrone, M., et al. 2018, The Astronomical Journal, 156, 125

Howard, A. W., Johnson, J. A., Marcy, G. W., et al. 2010, The Astrophysical Journal, 721, 1467

Howell, S. B., Sobeck, C., Haas, M., et al. 2014, Publications of the Astronomical Society of the Pacific, 126, 398

Huang, C., Wu, Y., & Triaud, A. H. 2016, The Astrophysical Journal, 825, 98

Huang, C. X., Burt, J., Vanderburg, A., et al. 2018, The Astrophysical Journal Letters, 868, L39

Hubbard, W. B., Hunten, D., Dieters, S., Hill, K., & Watson, R. 1988, Nature, 336, 452

Hunter, J. D. 2007, Computing in Science & Engineering, 9, 90

Hurvich, C. M., & Tsai, C.-L. 1989, Biometrika, 76, 297

Isella, A., Benisty, M., Teague, R., et al. 2019, The Astrophysical Journal Letters, 879, L25

Isella, A., Huang, J., Andrews, S. M., et al. 2018, The Astrophysical Journal Letters, 869, L49

Jewitt, D. 2003, Earth, Moon, and Planets, 92, 465

—. 2009, The Astronomical Journal, 137, 4296

—. 2018, The Astronomical Journal, 155, 56

Jewitt, D., Luu, J., Rajagopal, J., et al. 2017, The Astrophysical Journal Letters, 850, L36

Jewitt, D. C., Trujillo, C. A., & Luu, J. X. 2000, The Astronomical Journal, 120, 1140

Johansen, A., & Youdin, A. 2007, The Astrophysical Journal, 662, 627

Johansen, A., Youdin, A., & Mac Low, M.-M. 2009, The Astrophysical Journal Letters, 704, L75

Jones, R. L., Jurić, M., & Ivezić, Ž. 2016, in IAU Symposium, Vol. 318, Asteroids: New Observations, New Models, ed. S. R. Chesley, A. Morbidelli, R. Jedicke, & D. Farnocchia, 282–292

Jurgenson, C., Fischer, D., McCracken, T., et al. 2016, in Ground-based and Airborne Instrumentation for Astronomy VI, Vol. 9908, International Society for Optics and Photonics, 99086T

Kaiser, N., Aussel, H., Burke, B. E., et al. 2002, in Survey and Other Telescope Technologies and Discoveries, Vol. 4836, International Society for Optics and Photonics, 154–164

Kammer, J. A., Becker, T. M., Retherford, K. D., et al. 2018, The Astronomical Journal, 156, 72

Kant, I. 1755, Universal Natural History and Theory of the Heavens

Karim, M. T., & Mamajek, E. E. 2016, Monthly Notices of the Royal Astronomical Society, stw2772

Kavelaars, J., Holman, M., Grav, T., et al. 2004, Icarus, 169, 474

Kennedy, G. M., Matrà, L., Facchini, S., et al. 2019, Nature Astronomy, 3, 230

Keppler, M., Benisty, M., Müller, A., et al. 2018, Astronomy & Astrophysics, 617, A44

Knutson, H. A., Fulton, B. J., Montet, B. T., et al. 2014, The Astrophysical Journal, 785, 126

Korth, J., Csizmadia, S., Gandolfi, D., et al. 2019, Monthly Notices of the Royal Astronomical Society, 482, 1807

Kos, J., Bland-Hawthorn, J., Freeman, K., et al. 2018, Monthly Notices of the Royal Astronomical Society, 473, 4612

Kraft, R. P. 1967, The Astrophysical Journal, 150, 551

Kuchner, M. J., Faherty, J. K., Schneider, A. C., et al. 2017, The Astrophysical Journal Letters, 841, L19

Lagrange, J. L. 1873, Œuvres de Lagrange: Mémoires extraits des recueils de l'Académie des sciences de Paris et de la Classe des sciences mathématiques et physiques de l'Institut de France, Vol. 6 (Gauthier-Villars)

Lai, D. 2012, Monthly Notices of the Royal Astronomical Society, 423, 486

Lamy, P. L., Toth, I., Fernandez, Y. R., & Weaver, H. A. 2004, The sizes, shapes, albedos, and colors of cometary nuclei, ed. M. C. Festou, H. U. Keller, & H. A. Weaver, 223

Laplace, P.-S. 1796, Exposition du Systéme du Monde

Laughlin, G., & Batygin, K. 2017, Research Notes of the AAS, 1, 43

Leiva, R., Sicardy, B., Camargo, J., et al. 2017, The Astronomical Journal, 154, 159

Levison, H., Olkin, C., Noll, K., & Marchi, S. 2017, European Planetary Science Congress, 11, EPSC2017

Levison, H., et al. 2016, in Lunar and Planetary Science Conference, Vol. 47, 2061

Levison, H. F., Bottke, W. F., Gounelle, M., et al. 2009, Nature, 460, 364

Levison, H. F., Morbidelli, A., Tsiganis, K., Nesvornỳ, D., & Gomes, R. 2011, The Astronomical Journal, 142, 152

Levison, H. F., Morbidelli, A., VanLaerhoven, C., Gomes, R., & Tsiganis, K. 2008, Icarus, 196, 258

Lightkurve Collaboration, Cardoso, J. V. d. M., Hedges, C., et al. 2018, Astrophysics Source Code Library, ascl

Lin, D. N., Bodenheimer, P., & Richardson, D. C. 1996, Nature, 380, 606

Lin, D. N., & Papaloizou, J. 1986, The Astrophysical Journal, 309, 846

Lin, Y., & Ogilvie, G. I. 2017, Monthly Notices of the Royal Astronomical Society, 468, 1387

Lindegren, L., Hernández, J., Bombrun, A., et al. 2018, Astronomy & Astrophysics, 616, A2

Lockhart, M., Person, M. J., Elliot, J., & Souza, S. P. 2010, Publications of the Astronomical Society of the Pacific, 122, 1207

Long, F., Pinilla, P., Herczeg, G. J., et al. 2018, The Astrophysical Journal, 869, 17

Lord, P., Tilley, S., Oh, D. Y., et al. 2017, in Aerospace Conference, 2017 IEEE, IEEE, 1–11

Lovis, C., & Fischer, D. 2010, Exoplanets, 27

LSST Science Collaboration, Abell, P. A., Allison, J., et al. 2009, arXiv e-prints, arXiv:0912.0201

Lucy, L., & Sweeney, M. 1971, The Astronomical Journal, 76, 544

Luhman, K. 2013, The Astrophysical Journal Letters, 767, L1

Lykawka, P. S., Horner, J., Jones, B. W., & Mukai, T. 2009, Monthly Notices of the Royal Astronomical Society, 398, 1715

Macintosh, B., Graham, J., Barman, T., et al. 2015, Science, 350, 64

Mainzer, A., Bauer, J., Grav, T., et al. 2011, The Astrophysical Journal, 731, 53

Mainzer, A. K., Bauer, J. M., Cutri, R. M., et al. 2016, NASA Planetary Data System, 247, EAR

Males, J. R., Skemer, A. J., & Close, L. M. 2013, The Astrophysical Journal, 771, 10

Malhotra, R., Volk, K., & Wang, X. 2016, The Astrophysical Journal Letters, 824, L22

Mamajek, E. 2017, Research Notes of the AAS, 1, 21

Manara, C. F., Morbidelli, A., & Guillot, T. 2018, Astronomy & Astrophysics, 618, L3

Mankovich, C., Marley, M. S., Fortney, J. J., & Movshovitz, N. 2019, The Astrophysical Journal, 871, 1

Marois, C., Zuckerman, B., Konopacky, Q. M., Macintosh, B., & Barman, T. 2010, Nature, 468, 1080

Martin, R. G., & Lubow, S. H. 2017, The Astrophysical Journal Letters, 835, L28

Marzari, F., Scholl, H., Murray, C., & Lagerkvist, C. 2002, Asteroids III, 1, 725

Mashchenko, S. 2019, Monthly Notices of the Royal Astronomical Society, 489, 3003

McGlynn, T. A., & Chapman, R. D. 1989, The Astrophysical Journal, 346, L105

McKee, C. F., Parravano, A., & Hollenbach, D. J. 2015, ApJ, 814, 13

McKinney, W., et al. 2010, in Proceedings of the 9th Python in Science Conference, Vol. 445, Austin, TX, 51–56

McLaughlin, D. 1924, The Astrophysical Journal, 60

McNeill, A., Mommert, M., Trilling, D., Llama, J., & Skiff, B. 2019, The Astrophysical Journal Supplement Series, 245, 29

Meech, K. J., Weryk, R., Micheli, M., et al. 2017, Nature, 552, 378

Meisner, A., Bromley, B., Kenyon, S., & Anderson, T. 2018, The Astronomical Journal, 155, 166

Micheli, M., Farnocchia, D., Meech, K. J., et al. 2018, Nature, 559, 223

Millholland, S., & Laughlin, G. 2017, The Astronomical Journal, 153, 91

Millholland, S., Wang, S., & Laughlin, G. 2017, The Astrophysical Journal Letters, 849, L33

Moore, J. M., McKinnon, W. B., Cruikshank, D. P., et al. 2018, Geophys. Res. Lett., 45, 8111

Morbidelli, A., Levison, H. F., Tsiganis, K., & Gomes, R. 2005, Nature, 435, 462

Morello, G., Waldmann, I. P., Tinetti, G., et al. 2014, The Astrophysical Journal, 786, 22

Moro-Martín, A. 2018, The Astrophysical Journal, 866, 131

—. 2019, The Astrophysical Journal Letters, 872, L32

Moro-Martín, A., Turner, E. L., & Loeb, A. 2009, The Astrophysical Journal, 704, 733

Müller, T., Kiss, C., Alí-Lagoa, V., et al. 2018, Icarus

Murray, C. D., & Dermott, S. F. 1999, Solar System Dynamics (Cambridge University Press)

Myhrvold, N. 2018a, Icarus, 303, 91

—. 2018b, Icarus, 314, 64

—. 2018c, arXiv preprint arXiv:1812.06516

Nagasawa, M., & Ida, S. 2011, The Astrophysical Journal, 742, 72

Nagasawa, M., Ida, S., & Bessho, T. 2008, The Astrophysical Journal, 678, 498

Naoz, S., Farr, W. M., Lithwick, Y., Rasio, F. A., & Teyssandier, J. 2011, Nature, 473, 187

Naoz, S., Farr, W. M., & Rasio, F. A. 2012, The Astrophysical Journal Letters, 754, L36

Napier, K., Gerdes, D., Lin, H. W., et al. 2021, The Planetary Science Journal, 2, 59

Ness, M., Hogg, D. W., Rix, H.-W., Ho, A. Y., & Zasowski, G. 2015, The Astrophysical Journal, 808, 16

Ness, M., Hogg, D. W., Rix, H.-W., et al. 2016, The Astrophysical Journal, 823, 114

Nesvorný, D. 2011, The Astrophysical Journal Letters, 742, L22

—. 2018, Annual Review of Astronomy and Astrophysics, 56, 137

Ngo, H., Knutson, H. A., Hinkley, S., et al. 2015, The Astrophysical Journal, 800, 138

Nichols, T. F., & Powers, T. R. 1964, Moonlight and Night Visibility (U.S. Army Training Center Human Research Unit)

Nielsen, E. L., De Rosa, R. J., Macintosh, B., et al. 2019, The Astronomical Journal, 158, 13

O'Brien, D. P., Morbidelli, A., & Levison, H. F. 2006, Icarus, 184, 39

Ogilvie, G. I. 2013, Monthly Notices of the Royal Astronomical Society, 429, 613

—. 2014, Annual Review of Astronomy and Astrophysics, 52, 171

Oh, D. Y., Goebel, D. M., Elkins-Tanton, L., et al. 2016, in 52nd AIAA/SAE/ASEE Joint Propulsion Conference, 4541

Oliphant, T. E. 2006, A guide to NumPy, Vol. 1 (Trelgol Publishing USA)

Olkin, C. B., Young, L. A., French, R. G., et al. 2014, Icarus, 239, 15

Ortiz, J. L., Santos-Sanz, P., Sicardy, B., et al. 2017, Nature, 550, 219

Pál, A., Molnár, L., & Kiss, C. 2018, Publications of the Astronomical Society of the Pacific, 130, 114503

Pál, A., Kiss, C., Müller, T. G., et al. 2012, Astronomy & Astrophysics, 541, L6

Pál, A., Szakáts, R., Kiss, C., et al. 2020, The Astrophysical Journal Supplement Series, 247, 26

Parker, A. H., & Kavelaars, J. 2010, Publications of the Astronomical Society of the Pacific, 122, 549

Pasachoff, J. M. 2009, Research in Astronomy and Astrophysics, 9, 613

Pascucci, I., Testi, L., Herczeg, G., et al. 2016, The Astrophysical Journal, 831, 125

Payne, M. J., Holman, M. J., & Pál, A. 2019, Research Notes of the AAS, 3, 172

Peixinho, N., Delsanti, A., & Doressoundiram, A. 2015, Astronomy & Astrophysics, 577, A35

Pepe, F., Mayor, M., Delabre, B., et al. 2000, in Optical and IR telescope instrumentation and detectors, Vol. 4008, International Society for Optics and Photonics, 582–592

Pepe, F., Affolter, M., Delabre, B., et al. 2013, ESO Messenger, 153, 6

Pérez, L. M., Benisty, M., Andrews, S. M., et al. 2018, The Astrophysical Journal Letters, 869, L50

Perryman, M. A., Lindegren, L., Kovalevsky, J., et al. 1997, Astronomy & Astrophysics, 323, L49

Petersburg, R. R., Ong, J. J., Zhao, L. L., et al. 2020, The Astronomical Journal, 159, 187

Petrovich, C. 2015, The Astrophysical Journal, 799, 27

Petrovich, C., Muñoz, D. J., Kratter, K. M., & Malhotra, R. 2020, The Astrophysical Journal Letters, 902, L5

Petrovich, C., & Tremaine, S. 2016, The Astrophysical Journal, 829, 132

Pinsonneault, M., DePoy, D., & Coffee, M. 2001, The Astrophysical Journal Letters, 556, L59

Pinte, C., van der Plas, G., Ménard, F., et al. 2019, Nature Astronomy, 1

Pitjeva, E., & Pitjev, N. 2018, Celestial Mechanics and Dynamical Astronomy, 130, 57

Pollacco, D. L., Skillen, I., Cameron, A. C., et al. 2006, Publications of the Astronomical Society of the Pacific, 118, 1407

Pollack, J. B., Hubickyj, O., Bodenheimer, P., et al. 1996, Icarus, 124, 62

Portegies Zwart, S., Torres, S., Pelupessy, I., Bédorf, J., & Cai, M. X. 2018, Monthly Notices of the Royal Astronomical Society: Letters, 479, L17

Porter, S. B., Buie, M. W., Parker, A. H., et al. 2018, The Astronomical Journal, 156, 20

Price-Whelan, A. M., Sipőcz, B., Günther, H., et al. 2018, The Astronomical Journal, 156, 123

Queloz, D., Mayor, M., Weber, L., et al. 2000, Astronomy & Astrophysics, 354, 99

Rasio, F. A., & Ford, E. B. 1996, Science, 274, 954

Raymond, S. N., Armitage, P. J., Veras, D., Quintana, E. V., & Barclay, T. 2018, Monthly Notices of the Royal Astronomical Society, 476, 3031

Read, J. 2014, Journal of Physics G: Nuclear and Particle Physics, 41, 063101

Rein, H., & Liu, S.-F. 2012, Astronomy & Astrophysics, 537, A128

Rhodes, B. C. 2011, ascl, ascl

Rice, M., & Brewer, J. M. 2020, The Astrophysical Journal, 898, 119

Rice, M., & Laughlin, G. 2019a, The Astronomical Journal, 158, 19

—. 2019b, The Astrophysical Journal Letters, 884, L22

Rice, M., Wang, S., Howard, A. W., et al. 2021, The Astronomical Journal, 162, 182

Ricker, G. R., Winn, J. N., Vanderspek, R., et al. 2015, Journal of Astronomical Telescopes, Instruments, and Systems, 1, 014003

Robinson, S. E., Laughlin, G., Bodenheimer, P., & Fischer, D. 2006, The Astrophysical Journal, 643, 484

Rogers, L., & Seager, S. 2010, The Astrophysical Journal, 712, 974

Rogers, T., & Lin, D. 2013, The Astrophysical Journal Letters, 769, L10

Rogers, T., Lin, D. N., & Lau, H. H. B. 2012, The Astrophysical Journal Letters, 758, L6

Rossiter, R. 1924, The Astrophysical Journal, 60

Rowan-Robinson, M. 2021, Monthly Notices of the Royal Astronomical Society

Roy, A. E. 2005, Orbital Motion

Saint-Pé, O., Tulet, M., Davancens, R., et al. 2017, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 10568, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 105681O

Schlaufman, K. C. 2010, The Astrophysical Journal, 719, 602

Schwab, C., Rakich, A., Gong, Q., et al. 2016, in Ground-based and Airborne Instrumentation for Astronomy VI, Vol. 9908, International Society for Optics and Photonics, 99087H

Schwarz, G., et al. 1978, Annals of Statistics, 6, 461

Schwarz, R., Funk, B., Zechner, R., & Bazsó, Á. 2016, Monthly Notices of the Royal Astronomical Society, 460, 3598

Sekanina, Z. 2019, arXiv e-prints, arXiv:1905.00935

Sheppard, S. S., & Trujillo, C. A. 2006, Science, 313, 511

—. 2010a, Science, 329, 1304

—. 2010b, The Astrophysical Journal Letters, 723, L233

Sicardy, B., Ortiz, J., Assafin, M., et al. 2011, Nature, 478, 493

Sickafoose, A., Bosh, A., Levine, S., et al. 2019, Icarus, 319, 657

Silsbee, K., & Tremaine, S. 2018, The Astronomical Journal, 155, 75

Singla, A., Chandrasekaran, B., Godfrey, P., & Maggs, B. 2015, arXiv preprint arXiv:1505.03449

Siraj, A., & Loeb, A. 2019, arXiv preprint arXiv:1904.07224

Snodgrass, C., & Jones, G. H. 2019, Nature Communications, 10, 1

Southworth, J. 2008, Monthly Notices of the Royal Astronomical Society, 386, 1644

—. 2011, Monthly Notices of the Royal Astronomical Society, 417, 2166

Souza, S. P., Babcock, B. A., Pasachoff, J. M., et al. 2006, Publications of the Astronomical Society of the Pacific, 118, 1550

Spalding, C., & Batygin, K. 2015, The Astrophysical Journal, 811, 82

Stansberry, J., Grundy, W., Brown, M., et al. 2008, The Solar System Beyond Neptune, 161

Steinpilz, T., Joeris, K., Jungmann, F., et al. 2020, Nature Physics, 16, 225

Stern, S., Weaver, H., Spencer, J., Elliott, H., et al. 2018, Space Science Reviews, 214, 77

Stern, S., Weaver, H., Spencer, J., et al. 2019, Science, 364

Suzuki, D., Bennett, D., Sumi, T., et al. 2016, The Astrophysical Journal, 833, 145

Telting, J., Avila, G., Buchhave, L., et al. 2014, Astronomische Nachrichten, 335, 41

Teske, J. K., Thorngren, D., Fortney, J. J., Hinkel, N., & Brewer, J. M. 2019, The Astronomical Journal, 158, 239

Teyssandier, J., Lai, D., & Vick, M. 2019, Monthly Notices of the Royal Astronomical Society, 486, 2265

Ting, Y.-S., Conroy, C., Rix, H.-W., & Cargile, P. 2019, The Astrophysical Journal, 879, 69

Torres, S., Cai, M. X., Brown, A., & Zwart, S. P. 2019, Astronomy & Astrophysics, 629, A139

Trilling, D. E., Robinson, T., Roegge, A., et al. 2017, The Astrophysical Journal Letters, 850, L38

Trilling, D. E., Mommert, M., Hora, J. L., et al. 2018, The Astronomical Journal, 156, 261

Trujillo, C. A., & Sheppard, S. S. 2014, Nature, 507, 471

Tsiganis, K., Gomes, R., Morbidelli, A., & Levison, H. 2005, Nature, 435, 459

Valenti, J., & Piskunov, N. 1996, Astronomy & Astrophysics Supplement Series, 118, 595

Valenti, J. A., & Fischer, D. A. 2005, The Astrophysical Journal Supplement Series, 159, 141

Van Eylen, V., Albrecht, S., Gandolfi, D., et al. 2016, The Astronomical Journal, 152, 143

VanderPlas, J., Granger, B. E., Heer, J., et al. 2018, The Journal of Open Source Software, 3(32), 1057

Veras, D., Wyatt, M. C., Mustill, A. J., Bonsor, A., & Eldridge, J. J. 2011, Monthly Notices of the Royal Astronomical Society, 417, 2104

Vick, M., Lai, D., & Anderson, K. R. 2019, Monthly Notices of the Royal Astronomical Society, 484, 5645

Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, Nature Methods, 17, 261

Viswanathan, V., Fienga, A., Gastineau, M., & Laskar, J. 2017, Notes Scientifiques et Techniques de l'Institut de Mécanique Céleste, 108

Viswanathan, V., Fienga, A., Minazzoli, O., et al. 2018, Monthly Notices of the Royal Astronomical Society, 476, 1877

Vogt, S. S., Allen, S. L., Bigelow, B. C., et al. 1994, in Instrumentation in Astronomy VIII, Vol. 2198, International Society for Optics and Photonics, 362–375

Vogt, S. S., Radovan, M., Kibrick, R., et al. 2014, Publications of the Astronomical Society of the Pacific, 126, 359

Vokrouhlickỳ, D., Nesvornỳ, D., & Dones, L. 2019, The Astronomical Journal, 157, 181

Volk, K., & Malhotra, R. 2017, The Astronomical Journal, 154, 62

Wallace, L., Hinkle, K., Livingston, W., & Davis, S. 2011, The Astrophysical Journal Supplement Series, 195, 6

Walt, S. v. d., Colbert, S. C., & Varoquaux, G. 2011, Computing in Science & Engineering, 13, 22

Wang, S., Jones, M., Shporer, A., et al. 2019, The Astronomical Journal, 157, 51

Wang, S., Winn, J. N., Addison, B. C., et al. 2021, The Astronomical Journal, 162, 50

Wang, X.-Y., Rice, M., Wang, S., et al. 2022, The Astrophysical Journal Letters, 926, L8

Waskom, M., Botvinnik, O., Hobson, P. H., et al. 2014, Seaborn, vv0.5.0, Zenodo, doi:10.5281/zenodo.12710

Weidenschilling, S. 1977, Monthly Notices of the Royal Astronomical Society, 180, 57

Weiss, L. M., Marcy, G. W., Petigura, E. A., et al. 2018, The Astronomical Journal, 155, 48

Wheatley, P. J., West, R. G., Goad, M. R., et al. 2017, Monthly Notices of the Royal Astronomical Society, 475, 4476

Winn, J. N., Fabrycky, D., Albrecht, S., & Johnson, J. A. 2010, The Astrophysical Journal Letters, 718, L145

Winn, J. N., & Fabrycky, D. C. 2015, Annual Review of Astronomy and Astrophysics, 53, 409

Winn, J. N., Howard, A. W., Johnson, J. A., et al. 2009, The Astrophysical Journal, 703, 2091

Wisdom, J. 2008, Icarus, 193, 637

Wolf, M. 1906, Astronomische Nachrichten, 170, 147

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

Wong, I., & Brown, M. E. 2015, The Astronomical Journal, 150, 174

Wright, E., Mainzer, A., Masiero, J., et al. 2018, arXiv preprint arXiv:1811.01454

Wright, E. L., Eisenhardt, P. R., Mainzer, A. K., et al. 2010, The Astronomical Journal, 140, 1868

Wu, Y., & Lithwick, Y. 2011, The Astrophysical Journal, 735, 109

Wu, Y., & Murray, N. 2003, The Astrophysical Journal, 589, 605

Wu, Y., Murray, N. W., & Ramsahai, J. M. 2007, The Astrophysical Journal, 670, 820

Yee, S. W., Petigura, E. A., & Von Braun, K. 2017, The Astrophysical Journal, 836, 77

Yoshida, F., & Nakamura, T. 2005, The Astronomical Journal, 130, 2900

Youdin, A. N., & Goodman, J. 2005, The Astrophysical Journal, 620, 459

Young, E., Young, L., Olkin, C., et al. 2011a, Publications of the Astronomical Society of the Pacific, 123, 735

Young, E., French, R., Young, L., et al. 2008, The Astronomical Journal, 136, 1757

Young, E., Becklin, E., Marcum, P., et al. 2012, The Astrophysical Journal Letters, 749, L17

Young, L., Sicardy, B., Widemann, T., et al. 2010, in Bulletin of the American Astronomical Society, Vol. 42, 982

Young, L., Howell, R., French, R., et al. 2011b, in EPSC-DPS Joint Meeting, Vol. 1341

Yu, L., Zhou, G., Rodriguez, J. E., et al. 2018, The Astronomical Journal, 156, 250

Zahn, J.-P. 1977, Astronomy & Astrophysics, 57, 383

Zakamska, N. L., Pan, M., & Ford, E. B. 2011, Monthly Notices of the Royal Astronomical Society, 410, 1895

Zhang, S., Zhu, Z., Huang, J., et al. 2018, The Astrophysical Journal Letters, 869, L47

Zhu, Z., Zhang, S., Jiang, Y.-F., et al. 2019, The Astrophysical Journal Letters, 877, L18

Zsom, A., & Dullemond, C. 2008, Astronomy & Astrophysics, 489, 931